

Wiley Series in Pure and Applied Optics • Glenn D. Boreman, Series Editor

BUILDING ELECTRO-OPTICAL SYSTEMS

Making It All Work

SECOND EDITION

PHILIP C. D. HOBBS

 **WILEY**


SITE AVAILABLE

BUILDING ELECTRO- OPTICAL SYSTEMS

BUILDING ELECTRO- OPTICAL SYSTEMS

MAKING IT ALL WORK

SECOND EDITION

Philip C. D. Hobbs
Electrooptical Innovations
Briarcliff Manor, New York

 **WILEY**

A JOHN WILEY & SONS, INC., PUBLICATION

Copyright © 2009 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Hobbs, Philip C. D.

Building electro-optical systems : making it all work / Philip C.D. Hobbs.—2nd ed.
p. cm.—(Wiley series in pure and applied optics)

Includes bibliographical references and index.

ISBN 978-0-470-40229-0 (cloth)

1. Electrooptical devices—Design and construction. I. Title.

TA1750.H63 2008

621.381'045—dc22

2008047053

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

*In memory of my father,
Gerald H. D. Hobbs
John 6:40*



We have a habit in writing articles published in scientific journals to make the work as finished as possible, to cover up all the tracks, to not worry about the blind alleys or describe how you had the wrong idea first, and so on. So there isn't any place to publish, in a dignified manner, what you actually *did* in order to get to do the work.

—Richard P. Feynman, Nobel lecture 1996

CONTENTS

Preface	xv
Acknowledgments	xix
1 Basic Optical Calculations	1
1.1 Introduction /	1
1.2 Wave Propagation /	3
1.3 Calculating Wave Propagation in Real Life /	9
1.4 Detection /	33
1.5 Coherent Detection /	34
1.6 Interferometers /	36
1.7 Photon Budgets and Operating Specifications /	38
1.8 Signal Processing Strategy /	44
2 Sources and Illuminators	52
2.1 Introduction /	52
2.2 The Spectrum /	52
2.3 Radiometry /	54
2.4 Continuum Sources /	55
2.5 Interlude: Coherence /	59
2.6 More Sources /	63
2.7 Incoherent Line Sources /	68
2.8 Using Low Coherence Sources: Condensers /	69
2.9 Lasers /	71
2.10 Gas Lasers /	72
2.11 Solid State Lasers /	73
2.12 Diode Lasers /	75
2.13 Laser Noise /	83
2.14 Diode Laser Coherence Control /	89
3 Optical Detection	91
3.1 Introduction /	91
3.2 Photodetection in Semiconductors /	92

3.3	Signal-to-Noise Ratios /	92
3.4	Detector Figures of Merit /	94
3.5	Quantum Detectors /	100
3.6	Quantum Detectors with Gain /	109
3.7	Thermal Detectors /	117
3.8	Image Intensifiers /	118
3.9	Silicon Array Sensors /	120
3.10	How Do I Know Which Noise Source Dominates? /	131
3.11	Hacks /	136
4	Lenses, Prisms, and Mirrors	145
4.1	Introduction /	145
4.2	Optical Materials /	145
4.3	Light Transmission /	149
4.4	Surface Quality /	150
4.5	Windows /	151
4.6	Pathologies of Optical Elements /	152
4.7	Fringes /	153
4.8	Mirrors /	158
4.9	Glass Prisms /	160
4.10	Prism Pathologies /	165
4.11	Lenses /	165
4.12	Complex Lenses /	171
4.13	Other Lens-like Devices /	175
5	Coatings, Filters, and Surface Finishes	180
5.1	Introduction /	180
5.2	Metal Mirrors /	182
5.3	Transmissive Optical Coatings /	184
5.4	Simple Coating Theory /	186
5.5	Absorptive Filters /	196
5.6	Beam Dumps and Baffles /	198
5.7	White Surfaces and Diffusers /	204
6	Polarization	208
6.1	Introduction /	208
6.2	Polarization of Light /	208
6.3	Interaction of Polarization with Materials /	211
6.4	Absorption Polarizers /	215
6.5	Brewster Polarizers /	216
6.6	Birefringent Polarizers /	217

6.7	Double-Refraction Polarizers /	218
6.8	TIR Polarizers /	221
6.9	Retarders /	223
6.10	Polarization Control /	226
7	Exotic Optical Components	233
7.1	Introduction /	233
7.2	Gratings /	233
7.3	Grating Pathologies /	236
7.4	Types of Gratings /	237
7.5	Resolution of Grating Instruments /	240
7.6	Fine Points of Gratings /	242
7.7	Holographic Optical Elements /	244
7.8	Retroreflective Materials /	245
7.9	Scanners /	246
7.10	Modulators /	254
8	Fiber Optics	262
8.1	Introduction /	262
8.2	Fiber Characteristics /	262
8.3	Fiber Theory /	266
8.4	Fiber Types /	272
8.5	Other Fiber Properties /	277
8.6	Working with Fibers /	281
8.7	Fiber Devices /	287
8.8	Diode Lasers and Fiber Optics /	292
8.9	Fiber Optic Sensors /	292
8.10	Intensity Sensors /	293
8.11	Spectrally Encoded Sensors /	295
8.12	Polarimetric Sensors /	298
8.13	Fiber Interferometers /	299
8.14	Two-Beam Fiber Interferometers /	300
8.15	Multiple-Beam Fiber Interferometers /	301
8.16	Phase and Polarization Stabilization /	305
8.17	Multiplexing and Smart Structures /	307
8.18	Fiber Sensor Hype /	307
9	Optical Systems	309
9.1	Introduction /	309
9.2	What Exactly Does a Lens Do? /	309
9.3	Diffraction /	319

- 9.4 Aberrations / 336
- 9.5 Representing Aberrations / 340
- 9.6 Optical Design Advice / 344
- 9.7 Practical Applications / 345
- 9.8 Illuminators / 349

10 Optical Measurements 354

- 10.1 Introduction / 354
- 10.2 Grass on the Empire State Building / 354
- 10.3 Detection Issues: When Exactly Is Background Bad? / 359
- 10.4 Measure the Right Thing / 364
- 10.5 Getting More Signal Photons / 366
- 10.6 Reducing the Background Fluctuations / 370
- 10.7 Optically Zero Background Measurements / 373
- 10.8 Electronically Zero Background Measurements / 376
- 10.9 Labeling Signal Photons / 380
- 10.10 Closure / 385

11 Designing Electro-Optical Systems 387

- 11.1 Introduction / 387
- 11.2 Do You Really Want to Do This? / 387
- 11.3 Very Basic Marketing / 393
- 11.4 Classes of Measurement / 396
- 11.5 Technical Taste / 398
- 11.6 Instrument Design / 402
- 11.7 Guiding Principles / 407
- 11.8 Design for Alignment / 410
- 11.9 Turning a Prototype into a Product / 413

12 Building Optical Systems 415

- 12.1 Introduction / 415
- 12.2 Build What You Designed / 416
- 12.3 Assembling Lab Systems / 416
- 12.4 Alignment and Testing / 421
- 12.5 Optical Assembly and Alignment Philosophy / 425
- 12.6 Collimating Beams / 426
- 12.7 Focusing / 428
- 12.8 Aligning Beams with Other Beams / 430
- 12.9 Advanced Tweaking / 434
- 12.10 Aligning Laser Systems / 439
- 12.11 Adhesives / 441

12.12	Cleaning / 443	
12.13	Environmental Considerations / 446	
13	Signal Processing	448
13.1	Introduction / 448	
13.2	Analog Signal Processing Theory / 449	
13.3	Modulation and Demodulation / 453	
13.4	Amplifiers / 462	
13.5	Departures from Linearity / 462	
13.6	Noise and Interference / 467	
13.7	Frequency Conversion / 483	
13.8	Filtering / 487	
13.9	Signal Detection / 498	
13.10	Reducing Interference and Noise / 502	
13.11	Data Acquisition and Control / 504	
14	Electronic Building Blocks	509
14.1	Introduction / 509	
14.2	Resistors / 510	
14.3	Capacitors / 512	
14.4	Transmission Lines / 522	
14.5	Transmission Line Devices / 528	
14.6	Diodes and Transistors / 530	
14.7	Signal Processing Components / 539	
14.8	Digitizers / 548	
14.9	Analog Behavior of Digital Circuits / 558	
15	Electronic Subsystem Design	560
15.1	Introduction / 560	
15.2	Design Approaches / 560	
15.3	Perfection / 569	
15.4	Feedback Loops / 571	
15.5	Signal Detectors / 577	
15.6	Phase-Locked Loops / 587	
15.7	Calibration / 590	
15.8	Filters / 592	
15.9	Other Stuff / 595	
15.10	More Advanced Feedback Techniques / 597	
15.11	Hints / 599	
15.12	Linearizing / 601	
15.13	Digital Control and Communication / 604	

- 15.14 Miscellaneous Tricks / 607
- 15.15 Bulletproofing / 607

16 Electronic Construction Techniques 612

- 16.1 Introduction / 612
- 16.2 Circuit Strays / 612
- 16.3 Stray Coupling / 617
- 16.4 Ground Plane Construction / 618
- 16.5 Technical Noise and Interference / 621
- 16.6 Product Construction / 625
- 16.7 Getting Ready / 628
- 16.8 Prototyping / 629
- 16.9 Surface Mount Prototypes / 634
- 16.10 Prototyping Filters / 637
- 16.11 Tuning, or, You Can't Optimize What You Can't See / 639

17 Digital Postprocessing 644

- 17.1 Introduction / 644
- 17.2 Elementary Postprocessing / 645
- 17.3 Dead Time Correction / 650
- 17.4 Fourier Domain Techniques / 650
- 17.5 Power Spectrum Estimation / 666
- 17.6 Digital Filtering / 671
- 17.7 Deconvolution / 675
- 17.8 Resampling / 676
- 17.9 Fixing Space-Variant Instrument Functions / 678
- 17.10 Finite Precision Effects / 680
- 17.11 Pulling Data Out of Noise / 681
- 17.12 Phase Recovery Techniques / 685

18 Front Ends 688

- 18.1 Introduction / 688
- 18.2 Photodiode Front Ends / 690
- 18.3 Key Idea: Reduce the Swing Across C_d / 692
- 18.4 Transimpedance Amplifiers / 693
- 18.5 How to Go faster / 714
- 18.6 Advanced Photodiode Front Ends / 721
- 18.7 Other Types of Front End / 731
- 18.8 Hints / 734

19 Bringing Up the System	738
19.1 Introduction /	738
19.2 Avoiding Catastrophe /	741
19.3 Debugging and Troubleshooting /	744
19.4 Getting Ready /	745
19.5 Indispensable Equipment /	748
19.6 Analog Electronic Troubleshooting /	749
19.7 Oscillations /	753
19.8 Other Common Problems /	755
19.9 Debugging and Troubleshooting Optical Subsystems /	758
19.10 Localizing the Problem /	762
 Appendix: Good Books	 769
 Index	 779

PREFACE

You are fools, to say you learn from your mistakes. I learn from the mistakes of other men.

—Otto von Bismarck

This is a book of lore. *Lore* is an old word for wisdom and knowledge. While it often refers to magic and epic poetry, what I mean by it is altogether more homely: a mixture of rules of thumb, experience, bits of theory, and an indefinable feeling for the right way to do things, a sort of technical taste. It is what makes the difference between analyzing a design once completed and coming up with a good design to fit a particular purpose. Course work and textbooks have lots of analysis but most contain no lore whatsoever.

One of the odd things about lore is that it lives in the fingers more than in the brain, like piano playing. In writing this book, I have often run up against the difference between how I do something and how I *think* I do it, or how I remember having done it. Since it's the actual lore of doing that is useful, I have where possible written or revised each section when I was actually doing that task or consulting with someone who was. I hope that this gives those sections a sense of immediacy and authenticity.

Apologia

Lore is acquired slowly through experience and apprenticeship. Beginners pester experts, who help fairly willingly, mostly because they're kept humble by stepping in potholes themselves. This mutual aid system works but is slow and unsystematic. As a beginner, I once spent nearly six months trying to get a fancy laser interferometer to work properly, a task that would now take about a week. The reason was a breakdown in the apprenticeship system—everyone consulted said “Oh, that comes with practice”—perfectly true, and by no means unsympathetic, but not too helpful. Conversations with many others in the field indicate that this sort of thing is the rule and not the exception. Time, enthusiasm, and confidence are far too precious to go wasting them like that.

This book is an attempt to provide a systematic and accessible presentation of the practical lore of electro-optical instrument design and construction—to be the book I needed as a graduate student. It is intended for graduate students at all levels, as well as practicing scientists and engineers: anyone who has electro-optical systems to build and could use some advice. Its applicability ranges from experimental apparatus to optical disc players.

The range of topics covered here is enormously broad, and I wish I were master of it all. Most of it was invented by others whose names I don't know; it's the lore of a whole field, as filtered through one designer's head. It's mostly been learned by watching and doing, or worked out with colleagues at a white board, rather than reading journal articles, so there aren't many references. For further reading, there is a list of 100 or so good books in the Appendix that should fill in the gaps.

I hope that a book like this can erect bridges between subdisciplines, prevent common mistakes, and help all those working on an instrument project to see it as a whole. So much good stuff gets lost in the cracks between physics, electrical engineering, optical engineering, and computer science, that a salvage attempt seemed justified. I apologize to those whose work has been acknowledged inadequately or whose priority has been overlooked, and hope that they can remember once needing a book like this.

Mission

Designing and constructing electro-optical instruments is without a doubt one of the most interdisciplinary activities in engineering. It makes an absorbing and rewarding career, with little danger of growing stale. On the other hand, the same interdisciplinary quality means that instrument building is a bit scary and keeps us on our toes. The very broad range of technologies involved means that at least one vital subsystem lies outside the designer's expertise, presenting a very real danger of major schedule slippage or outright failure, which may not become apparent until very late in the project.

We in electro-optics rely on whatever subset of these technologies we are familiar with, together with a combination of outside advice, collaboration, and purchased parts. Often, there are many ways of reaching the goal of a robust, working system; then the problem is where to start among a range of unfamiliar alternatives. It's like the classic computer game ADVENT: 'You are in a maze of twisty little passages, all different.' Some judicious advice (and perhaps a map left by a previous adventurer) is welcome at such times, and that's what this book is about, the lore of designing and building electro-optical instruments that work.

To have confidence in an instrument design, we really need to be able to calculate its performance ahead of time, without constructing an elaborate simulation. It is a nontrivial matter, given the current fragmented state of the literature, to calculate what the resolution and SNR of a measurement system will be before it is built. It's not that there isn't lots of information on how to calculate the performance of each lens, circuit, or computer program, but rather the complexity of the task and the very different ways in which the results are expressed in the different fields encountered. For example, what is the effect of fourth-order spherical aberration in the objective lens on the optimal band-setting filter in the analog signal processor, and then on the signal-to-noise ratio of the ultimate digital data set? Somebody on the project had better know that, and my aim is to make you that somebody.

The book is intended in the first instance for use by oppressed graduate students in physics and electrical engineering, who have to get their apparatus working long enough to take some data before they can graduate. When they do, they'll find that real-world design work has much the same harassed and overextended flavor, so in the second instance, it's intended for working electro-optical designers. It can be used as a text in a combined lecture-laboratory course aimed at graduate students or fourth-year undergraduates, and as a self-teaching guide and professional reference by working designers.

The warm reception that the first edition received suggests that despite its faults it has filled a real need. In this edition, everything has been revised, some previously over-terse sections have been expanded, and more than 100 pages' worth of new material has been added. Component lists and electronic designs have been updated where needed. Only a very few things have been dropped, owing to space constraints or component obsolescence.

Organization

Textbooks usually aim at a linear presentation of concepts, in which the stuff on page n does not depend on your knowing pages $n + 1 \dots N$. This is very valuable pedagogically, since the reader is initially unfamiliar with the material and usually will go through the book thoroughly, once, under the guidance of a teacher who is presenting information rapidly. Reference books are written for people who already have a grasp of the topic but need to find more detail or remind themselves of things dimly remembered. Thus they tend to treat topics in clumps, emphasizing completeness, and to be weak on overall explanations and on connections between topics.

Those two styles work pretty well in some subject areas, but design lore is not one of them. Its concepts aren't branched like a tree, or packed like eggs in a crate, but rather are interlinked like a fishnet or a sponge; thus a purely linear or clumped presentation of lore is all but impossible without doing violence to it. Nonetheless, to be of any use, a lore book must be highly accessible, both easy to work through sequentially and attractive to leaf through many times.

Computer scientists use the concept of locality of reference—it's a good thing if an algorithm works mainly with data near each other in storage, since it saves cache misses and page faults, but all the data have to be there, regardless. That's the way I have tried to organize this book: most of the lore on a particular topic is kept close together in the book for conceptual unity and easy reference, but the topics are presented in a sufficiently linear order that later chapters build mainly on earlier ones, and important connections are noted in both forward and backward directions.[†] A certain amount of messiness results, which (it is to be hoped) has been kept close to a minimum. This approach gives rise to one minor oddity, which is that the same instruments are considered from different angles in different chapters, so some flipping of pages is required to get the whole picture.

The book is organized into three sections: Optics; Electronics and Signal Processing; and Special Topics In Depth (Front Ends and Bringing Up the System). There is also Supplementary Material, available from the websites ftp://ftp.wiley.com/public/sci_tech_med/electrooptical and <http://electrooptical.net>, which comprises Chapter 20 on Thermal Control and chapter problems for the whole book.

The material is presented in varying levels of detail. The differences in the detail levels reflect the amount of published lore and the measured density of deep potholes that people fall into. For example, there are lots of potholes in optomechanical design, but weighty books of relevant advice fill shelf after shelf. Anyway, mechanical problems aren't usually what cause instrument projects to fail—unexamined assumptions, inexperience, and plain discouragement are. To get the job done, we talk instead about how to avoid common mistakes while coming up with something simple that works reliably.

The one big exception to this general scheme is Chapter 1. It pulls in strands from everywhere, to present the process and the rhythm of conceptual design, and so contains things that many readers (especially beginners) may find unfamiliar. Don't worry too much about the technical aspects, because there's more on all those things later in the book, as well as pointers to other sources.

A complete instrument design course based on this book would probably have to wait for a first- or second-year graduate class. Undergraduate students with a good grasp of electromagnetism, physical optics, and Fourier transforms might benefit from a

[†]Because electro-optical lore is so interconnected, useful connections that are tangential to the discussion are relegated to footnotes. An occasional polemic is found there too.

fourth-year course on optical instruments based selectively on the first ten chapters. To get the most out of such a course, the audience should be people with instruments of their own to build, either in a lab course, as a senior project, or as part of their graduate work. Because of the complicated, interdisciplinary nature of instrument building, the laboratory part of the course might best be done by teams working on an instrument project rather than individually, provided that each designer knows enough about everybody else's part to be able to explain it.

Chapter Problems

Chapter problems for the book are available on the websites listed above. Making complicated tasks intuitive is the true realm of lore—knowing the mathematical expression for the fringe pattern of a defocused beam is less useful than knowing which way to turn which knob to fix it. The most powerful method for gaining intuition is to use a combination of practical work and simple theoretical models that can be applied easily and stay close to the real physics. Accordingly, the emphasis in the problems is on extracting useful principles from theory and discussion.

Most of the problems have been taken from real design and scientific work, and so tend to be open-ended. Most students will have had a lot of theoretical training, but nowadays most will not have the skills of a Lightning Empiricist, a gimlet-eyed designer who's fast at mental rule-of-thumb calculations and who sanity checks everything by reflex. Perhaps this book can help fix that.

Errata

A certain number of errors and misconceptions—hopefully minor—are bound to creep into a book of this type, size, and scope, unfortunately. I welcome your comments and corrections, large and small: errata and omissions will be made available at ftp://ftp.wiley.com/public/sci_tech_med/electro-optical/errata2.txt and <http://electrooptical.net/www/beos2e/errata2.txt> and will be incorporated in future printings. Send e-mail to hobbs@stanfordalumni.org.

P. C. D. HOBBS

Briarcliff Manor, New York
Michaelmas (September 29), 2008

ACKNOWLEDGMENTS

To acquire lore, one needs a big sandbox and long uninterrupted stretches of time to spend there, absorbed in the play. I am forever grateful to my parents for providing that sort of environment in my growing up, and for believing in me even when only the mess was visible.

I learned most of this material through participating in the stimulating and supportive technical cultures of the places where I've been fortunate enough to study and to work: the Edward L. Ginzton Laboratory at Stanford University, Stanford, California; the Department of Physics and the Department of Geophysics & Astronomy at the University of British Columbia and Microtel Pacific Research (both in Vancouver BC) and the IBM Thomas J. Watson Research Center at Yorktown Heights, New York. I owe a special debt to IBM and to my managers there, Arthur Ciccolo, Frank Libsch, and John Mackay, for supporting this project and for generously allowing me time and resources to work on it.

I also wish to thank some of the many other gifted people who I have been privileged to have as close colleagues, teachers, and friends, particularly J. Samuel Batchelder (who first suggested I write this book), Donald M. DeCain, Kurt L. Haller, Gordon S. Kino, the late Roger H. Koch, Brian A. Murray, Martin P. O'Boyle, Marc A. Taubenblatt, Theodore G. van Kessel, and Robert H. Wolfe. Without them I'd still be stuck in one of those potholes way back along the road.

Most of all, I wish to thank my wife, Maureen, and our offspring, Bronwen, Magdalen, and Simon, for their patience and encouragement while I wrote and wrote.

P. C. D. H.

Basic Optical Calculations

An excellent plumber is infinitely more admirable than an incompetent philosopher. The society which scorns excellence in plumbing because plumbing is a humble duty and tolerates shoddiness in philosophy because it is an exalted activity will have neither good plumbing nor good philosophy. Neither its pipes nor its theories will hold water.

—John W. Gardner[†]

1.1 INTRODUCTION

Okay, we've decided to build an electro-optical system. It's going to be so great that everybody who sees it will hate us. Now comes the fun part, the planning and designing, and then the hard part, the building and testing. To design and plan, we have to have some way of knowing how our brainchild ought to behave before it is built—that is, theory.

At the conceptual stage, the measurement principle is poorly understood and many quite different schemes are suggesting themselves. To make sense of it, you need a white board and a couple of smart and experienced colleagues to bounce ideas off, plus some pointers on how to begin. The aim of this chapter is to equip you to do a conceptual instrument design on the back of an envelope. It assumes some background in optics, especially some acquaintance with plane waves and Fourier transforms.

The indispensable ingredients of a conceptual design are:

- A measurement idea
- Operational requirements (field of view, scan speed, spot size, sensitivity, etc.)
- A photon budget
- A rough optical design
- A detection strategy
- A signal processing strategy

[†]John W. Gardner, *Excellence, Can We Be Equal and Excellent Too?* Harper, New York, 1961, p. 86.

The best way to get them is through several sessions at that white board, with a lot of thought and calculation in between. (It is amazing how many people think they've finished with engineering calculations when they hand in their last exam, but that attitude is sudden death in the instrument-building business.) Once you have these, you can make a list of the main technical risks, in descending order of hairiness, and pretty soon you have a plan for how to proceed. The size of these technical risks is important—they can range from finding parts to violating laws of physics. Right at the beginning, we must decide whether the measurement is even possible, which requires more imagination than analysis. The invention of two-photon Doppler-free spectroscopy is a good example.

Example 1.1: Two-Photon Doppler-Free Spectroscopy.[†] Gas-phase spectroscopy is limited at low pressure by the random thermal motions of the gas molecules, and at high pressures by their collisions. Molecules with different speeds along the laser beam axis experience different Doppler shifts, so that their absorption features occur at different frequencies in the lab frame, leading to *Doppler broadening*. A two-photon transition involves the absorption of two photons, whose energies must sum to the transition energy. The absorption (and any resulting fluorescence) can be modulated by chopping the excitation beam. By using two excitation beams going in opposite directions, some events will involve absorption of one photon from each beam, which can occur only when both beams are unblocked by their choppers. If the modulation frequencies of the two beams are different, this part of the signal will appear at the sum and difference of the chopping frequencies. If a molecule has nonzero speed along the beam, then to leading order in V/c , it will see each beam shifted by

$$\Delta\nu_i = \frac{-\mathbf{k}_i \cdot \mathbf{v}}{2\pi}. \quad (1.1)$$

Since the two beams have oppositely directed \mathbf{k} vectors, one will be upshifted and the other downshifted by the same amount; the sum of the photon energies, $h(\nu_1 + \nu_2)$, is unshifted. Thus these mixing components are present only when the laser is tuned exactly to the rest-frame resonance frequency—they show no first-order Doppler broadening. An apparently inherent physical limit is circumvented with an extra chopper and a mirror or two; such is the power of a good measurement idea.

Once the idea is in hand, analysis is needed, to decide between competing alternatives. Such feasibility calculations are at the heart of electro-optical systems lore, and their most important fruit is a highly trained intuition, so we'll do lots of examples. The places where assumptions break down are usually the most instructive, so we'll also spend some time discussing some of the seedier areas of optical theory, in the hope of finding out where the unexamined assumptions lurk. We'll begin with wave propagation and imaging.

[†]This example is adapted from L. S. Vasilenko, V. P. Chebotaev, and A. V. Shishaev, *JETP Lett.* **3** (English translation), 161 (1970).

1.2 WAVE PROPAGATION

1.2.1 Maxwell's Equations and Plane Waves

Any self-respecting book on optics is obliged to include Maxwell's equations, which are the complete relativistic description of the electromagnetic field *in vacuo* and are the basis on which all optical instruments function. They are also useful for printing on T-shirts. Fortunately, this is a book of practical lore, and since Maxwell's equations are almost never used in real design work, we don't need to exhibit them. The most basic equation that is actually useful is the *vector wave equation*,

$$\nabla^2 \mathbf{E} - \frac{1}{c^2} \frac{\partial^2 \mathbf{E}}{\partial t^2} = 0, \quad (1.2)$$

where c is the speed of light. Most of the time we will be calculating a monochromatic field, or at least a superposition of monochromatic fields. We can then separate out the time dependence as $\exp(-i\omega t)$ and write $k = \omega/c$, leaving the vector Helmholtz equation,

$$(\nabla^2 + k^2) \mathbf{E} = 0. \quad (1.3)$$

Its simplest solution is a vector plane wave,

$$\mathbf{E}(\mathbf{x}) = \mathbf{E}_0 e^{i\mathbf{k}\cdot\mathbf{x}} \quad (1.4)$$

where the two fixed vectors are \mathbf{E}_0 , the electric field vector, which may be complex, and \mathbf{k} is the *wave vector*, whose magnitude $k = |\mathbf{k}| = \omega/c$ is called the *propagation constant*. If \mathbf{E}_0 is real, the field is said to be *linearly polarized* along \mathbf{E}_0 ; if its real and imaginary parts are the same size, so that the instantaneous \mathbf{E} rotates without changing length, the field is *circularly polarized*; otherwise, it's *elliptically polarized* (see Section 1.2.8).

Power flows parallel to \mathbf{k} in an isotropic medium, but need not in an anisotropic one, so it is separately defined as the *Poynting vector* $\mathbf{S} = \mathbf{E} \times \mathbf{H}$ (see Sections 4.6.1 and 6.3.2). In the complex notation, the one-cycle average Poynting vector is $\mathbf{S} = \text{Re}[\mathbf{E} \times \mathbf{H}^*]$.

1.2.2 Plane Waves in Material Media

So far, we have only considered propagation *in vacuo*. Electromagnetics in material media is enormously complicated on a microscopic scale, since there are $\sim 10^{22}$ scatterers/cm³. Fortunately, for most purposes their effects are well approximated by mean field theory, which smears out all those scatterers into a jelly that looks a lot like vacuum except for in a change in the propagation velocity, the E/H ratio, and some loss. A plane wave entering a material medium via a plane surface remains a plane wave, with different \mathbf{k} and \mathbf{E}_0 .

In a medium, light travels at a speed $v = c/n$. The constant n , the *refractive index*, is given by $n = \sqrt{\mu_r \epsilon_r}$, where μ_r and ϵ_r are the relative magnetic permeability and dielectric constant of the material at the optical frequency, respectively. Since μ_r is nearly always 1 in optics, $n = \sqrt{\epsilon_r}$. In addition, the material will change the *wave impedance*, $Z = E/H = \sqrt{\mu/\epsilon} = (120\pi \Omega) \sqrt{\mu_r/\epsilon_r}$. The analogy between wave impedance and transmission line impedance is a fruitful one.

In absorbing media, the refractive index is complex.[†] Assuming the medium is linear and time invariant, the temporal frequency cannot change, so k is different in the medium; the new k is $k_n = nk_0$, where k_0 is the vacuum value. We usually drop the subscripts, so k is taken to be in the medium under consideration.

There are one or two fine points deserving attention here. One is that n is not constant with ω , a phenomenon known as *dispersion*. Besides making light of different colors behave differently, this leads to distortion of a time-modulated signal. The carrier wave propagates at the *phase velocity* $v_p = \omega/k$, but the envelope instead propagates, approximately unchanged in shape, at the *group velocity* v_g , given by

$$v_g = \partial\omega/\partial k. \quad (1.5)$$

Since the carrier propagates at the phase velocity v , as an optical pulse goes along its carrier “slips cycles” with respect to its envelope; that’s worth remembering if you build interferometers. The group velocity approximation (1.5) holds for short propagation distances only, that is, when the difference Δt in the transit times of different frequency components is much less than the pulse width τ . In the opposite limit, where $\Delta t \gg \tau$, the output is a time Fourier transform of the input pulse.[‡]

The other fine point is that ϵ_r is in general a tensor quantity; there is no guarantee that the response of the material to an applied electric field is the same in all directions. In this book we’re concerned only with the linear properties of optical systems, so a tensor is the most general local relationship possible. The resulting dependence of n on polarization leads to all sorts of interesting things, such as *birefringence* and *beam walkoff*. There are in addition strange effects such as *optical activity* (also known as *circular birefringence*), where the plane of polarization of a beam rotates slowly as it propagates. We’ll talk more about these weird wonders in Chapters 4 and 6.

Aside: The Other Kind of Polarization. The dielectric constant ϵ_r expresses the readiness with which the charges in the medium respond to the applied electric field; it is given by $\epsilon_r = 1 + 4\pi\chi$, where χ is the electric susceptibility (zero for vacuum); the electric polarization \mathbf{P} is $\epsilon_0\chi\mathbf{E}$. This is a different usage than the usual optical meaning of polarization, and it’s worth keeping the two distinct in your mind.

1.2.3 Phase Matching

The two basic properties of any wave field are amplitude and phase. At any point in space-time, a monochromatic wave has a unique phase, which is just a number specifying how many cycles have gone by since time $t = 0$. Since it’s based on counting, phase is invariant to everything—it’s a Lorentz scalar, so it doesn’t depend on your frame of reference or anything else, which turns out to be a very powerful property. The requirement for phase matching at optical interfaces is the physical basis of geometrical optics. A plane wave has the unique property of being translationally invariant, meaning that if you move from one point to another, the only thing that changes is an additive

[†]Complex refractive index \tilde{n} is often quoted as $n + ik$, where n and k are real and positive, but it is conceptually simpler to leave n complex, because the Fresnel formulas and Snell’s law still work with absorbing media.

[‡]The impulse response of a linearly dispersive medium is a chirp, and the Fourier transform can be computed as the convolution of a function with a chirp. This turns out to be important in digital signal processing, where it leads to the chirp-Z transform.

phase shift (equivalent to a pure time delay). In particular, at a planar interface, moving the reference point within the plane cannot change the phase relationship between the fields on either side.

1.2.4 Refraction, Snell's Law, and the Fresnel Coefficients

If a plane wave encounters a plane interface between two semi-infinite slabs of index n_1 and n_2 , as shown in Figure 1.1, the light is partially reflected and partially transmitted—a standard problem in undergraduate electromagnetics classes. We expect the fields to consist of an incident and a reflected plane wave on the input side and a single transmitted plane wave on the output side. Phase matching at the interface requires that the tangential \mathbf{k} vectors of all the waves be the same, which reproduces the law of reflection for the reflected component and *Snell's law* for the transmitted one:

$$n_1 \sin \theta_1 = n_2 \sin \theta_2. \quad (1.6)$$

If there are m parallel planar interfaces, \mathbf{k}_{\parallel} is the same in all the layers, so since (in the j th layer) $k_j = n_j k_0$, we can use the phase matching condition to get k_{\perp} in the j th layer:

$$k_{\perp}^2 = n_j^2 k_0^2 - k_{\parallel}^2. \quad (1.7)$$

This is important in the theory of optical coatings. The continuity conditions on tangential \mathbf{E} and perpendicular \mathbf{D} across the boundary give the *Fresnel formulas* for the field amplitudes,

$$\frac{E_p''}{E_p} \equiv r_{p12} = -\frac{\tan(\theta_1 - \theta_2)}{\tan(\theta_1 + \theta_2)} = -\frac{n_2 \cos \theta_1 - n_1 \sqrt{1 - [(n_1/n_2) \sin \theta_1]^2}}{n_2 \cos \theta_1 + n_1 \sqrt{1 - [(n_1/n_2) \sin \theta_1]^2}}, \quad (1.8)$$

$$\frac{E_p'}{E_p} \equiv t_{p12} = \frac{2 \sin \theta_1 \cos \theta_2}{\sin(\theta_1 + \theta_2)} = \frac{2n_1 \cos \theta_1}{n_2 \cos \theta_1 + n_1 \sqrt{1 - [(n_1/n_2) \sin \theta_1]^2}}, \quad (1.9)$$

for light linearly polarized (i.e., \mathbf{E} lying) in the *plane of incidence*. This plane is defined by the surface normal $\hat{\mathbf{n}}$ (unrelated to n) and \mathbf{k}_{inc} . In Figure 1.1, it is the plane of




Figure 1.1. Refraction and reflection of a plane wave at a plane dielectric boundary. The angle of refraction θ_2 is given by Snell's law.

the page. For light linearly polarized perpendicular to the plane of incidence, these become

$$\frac{E''_s}{E_s} \equiv r_{s12} = -\frac{\sin(\theta_1 - \theta_2)}{\sin(\theta_1 + \theta_2)} = \frac{n_1 \cos \theta_1 - n_2 \sqrt{1 - [(n_1/n_2) \sin \theta_1]^2}}{n_1 \cos \theta_1 + n_2 \sqrt{1 - [(n_1/n_2) \sin \theta_1]^2}}, \quad (1.10)$$

$$\frac{E'_s}{E_s} \equiv t_{s12} = \frac{2 \sin \theta_2 \cos \theta_1}{\sin(\theta_1 + \theta_2)} = \frac{2n_1 \cos \theta_1}{n_1 \cos \theta_1 + n_2 \sqrt{1 - [(n_1/n_2) \sin \theta_1]^2}}. \quad (1.11)$$

The two polarizations are known as *p* and *s*, respectively. As a mnemonic, *s* polarization means that **E** *sticks* out of the plane of incidence.[†] The quantities *r* and *t* are the reflection and transmission coefficients, respectively. These *Fresnel coefficients* act on the amplitudes of the fields.[‡] The transmitted and reflected power ratios *R* and *T*, given by

$$R = |r|^2 \quad \text{and} \quad T = \frac{n_2 \cos \theta_2}{n_1 \cos \theta_1} |t|^2, \quad (1.12)$$

are known as the *reflectance* and *transmittance*, respectively.

The Fresnel coefficients have fairly simple symmetry properties; if the wave going from n_1 to n_2 sees coefficients r_{12} and t_{12} , a wave coming in the opposite direction sees r_{21} and t_{21} , where

$$\begin{aligned} r_{p21} &= -r_{p12}, & t_{p21} &= (n_1 \cos \theta_1)/(n_2 \cos \theta_2)t_{p12}; \\ r_{s21} &= -r_{s12}, & t_{s21} &= (n_2 \cos \theta_2)/(n_1 \cos \theta_1)t_{s12}. \end{aligned} \quad (1.13)$$

The symmetry expressions for t_{21} are more complicated because they have to take account of energy conservation between the two media.

1.2.5 Brewster's Angle

Especially sharp-eyed readers may have spotted the fact that if $\theta_1 + \theta_2 = \pi/2$, the denominator of (1.8) goes to infinity, so $r_p = 0$. At that angle, $\sin \theta_2 = \cos \theta_1$, so from Snell's law, $\tan \theta_1 = n_2/n_1$. This special value of θ_i is called *Brewster's angle* θ_B . Note that the transmitted angle is $\pi/2 - \theta_B$, which is Brewster's angle for going from n_2 into n_1 . Brewster angle incidence with very pure *p*-polarized light is the best existing technique for reducing reflections from flat surfaces, a perennial concern of instrument designers (see Section 4.7.3).

Laser tube windows are always at Brewster's angle to reduce the round-trip loss through the cavity. The loss in the *s* polarization due to four high angle quartz–air surfaces is nearly 40% in each direction. Regeneration in the cavity greatly magnifies this gain difference, which is why the laser output is highly polarized. Brewster angle

[†]The *s* is actually short for *senkrecht*, which is German for perpendicular. The two polarizations are also called *TE* and *TM*, for *transverse electric* and *transverse magnetic*, that is, which field is sticking out of the plane of incidence. This nomenclature is more common in waveguide theory.

[‡]There is another sign convention commonly used for the *p*-polarized case, where the incident and reflected **E** fields are taken in opposite directions, yielding a confusing sign change in (1.8). We adopt the one that makes $r_p = r_s$ at normal incidence.

incidence is also commonly used in spectroscopic sample cells, Littrow prisms, and other high accuracy applications using linearly polarized, collimated beams.

Ideally, a smooth optical surface oriented at θ_B to the incoming p -polarized beam would reflect nothing at all, but this is not the case with real surfaces. Roughness and optical anisotropy make it impossible to make every single region encounter the light beam at the same angle or with the same refractive index, so there are always residual reflections even at θ_B . Surface layers also prevent complete canceling of the reflected wave, because the two will in general have different Brewster's angles and because of the phase delay between the reflections from the top and bottom of the layer. Below the critical angle, dielectric reflections always have phase 0 or π , so there's no way to tip the surface to get rid of a phase-shifted reflection.

Aside: Fossil Nomenclature. When Malus discovered polarization (in 1808) by looking at reflections from a dielectric, he quite reasonably identified the *plane of polarization* with the plane of incidence. This conflicts with our modern tendency to fix on the \mathbf{E} field in considering polarization. There are still some people who follow Malus's convention, so watch out when you read their papers.

1.2.6 Total Internal Reflection

If $n_1 > n_2$, there exists an angle θ_C , the *critical angle*, where Snell's law predicts that $\sin \theta_2 = 1$, so $\theta_2 = \pi/2$: grazing incidence. It is given by

$$\theta_C = \arcsin(n_2/n_1). \quad (1.14)$$

Beyond there, the surds in the Fresnel formulas (1.8)–(1.11) become imaginary, so t vanishes and r sits somewhere on the unit circle (the reflectivity is 1 and the elements of \mathbf{E}' become complex).

This *total internal reflection* (TIR) is familiar to anyone who has held a glass of water, or looked up at the surface while underwater. It is widely used in reflecting prisms. There are two things to remember when using TIR: the reflection phase is massively polarization dependent and the fields extend beyond the surface, even though there is no propagating wave there. A TIR surface must thus be kept very clean, and at least a few wavelengths away from any other surface.

By putting another surface sufficiently close by, it is possible to couple light via the evanescent field, a phenomenon called *frustrated TIR* or, more poetically, *evanescent coupling*. This is the optical analogue of quantum mechanical tunneling.

The reflection phase represents a relative time delay of the propagating wave. The s -polarized wave is delayed more, because it has a larger amplitude in the evanescent region, which requires more of a phase slip between the incident and reflected waves (remember the continuity conditions). This sort of physical reasoning is helpful in keeping sign conventions straight, although it is not infallible. The phase shift δ between s and p polarizations is[†]

$$\delta = \delta_s - \delta_p = -\arctan \frac{2 \cos \theta_i \sqrt{\sin^2 \theta_i - (n_2/n_1)^2}}{\sin^2 \theta_i}. \quad (1.15)$$

[†]M. Born and E. Wolf, *Principles of Optics*, 6th ed. (corrected). Pergamon, Oxford, 1983, pp. 47–51.

1.2.7 Goos–Hänchen Shift

The angle-dependent phase shift on TIR functions much as dispersion does in the time domain, delaying different components differently. Dispersion causes the envelope of a pulse to propagate at the group velocity, which is different from the phase velocity. In the same way, the phase shift on TIR causes the envelope of a reflected beam to be shifted slightly in space from the incident beam, the *Goos–Hänchen shift*. It is less well known than the group velocity effect, mainly because the effect doesn't build up as it goes the way dispersion effects do, so the shift is small under normal circumstances, though large enough to cause image aberrations on reflection from TIR surfaces. The place it does become important is in multimode fiber optics, where a ray undergoes many, many reflections and the shift accordingly adds up.

The variation in the Goos–Hänchen shift comes from the speeding up of the wave that sticks out the most into the low index material. It is responsible for the apparently paradoxical behavior of optical fiber modes near cutoff (see Section 8.3.1). We expect high angle modes to slow down due to the decrease of k_z with angle—they spend more of their time bouncing back and forth instead of traveling down the fiber axis. In fact, they do slow down with increasing angle at first, but then speed up again as they near cutoff, when the wave sticks farther and farther out into the low index material. To leading order, the effect is the same as if the wave bounced off an imaginary surface one decay length into the low index material. Note that this does not contradict the last section; the phase shift is a delay, but the Goos–Hänchen shift makes the mode propagation anomalously fast.

1.2.8 Circular and Elliptical Polarization

What happens when the p and s components get out of phase with each other? The exponential notation, remember, is just a calculating convenience; real physical quantities always give real numbers. The instantaneous E -field strength is

$$E^2 = [\text{Re}\{E_x e^{i\omega t}\}]^2 + [\text{Re}\{E_y e^{i\omega t + \phi}\}]^2. \quad (1.16)$$

A linearly polarized monochromatic wave has an E that varies sinusoidally, passing through zero twice each cycle. When \mathbf{E} has complex coefficients, the p and s components oscillate out of phase with one another. If the two are the same size and a quarter cycle apart, the real (i.e., physical) part of the \mathbf{E} vector will spin through 2π once per cycle, without changing its length, like a screw thread. Its endpoint will traverse a circle, so that this is known as *circular polarization*. Like screws, there is right and left circular polarization, but unlike screws, the names are backwards.

If the two components are not exactly equal in magnitude, or are not exactly $\pi/2$ radians apart, the vector will still rotate, but will change in length as it goes round, tracing an ellipse. This more general case is *elliptical polarization*. Circular and linear polarizations are special cases of elliptical polarization. Elliptical polarization can also be right or left handed.

1.2.9 Optical Loss

In a lossless medium, the \mathbf{E} and \mathbf{H} fields of a propagating wave are exactly in phase with each other. Any phase difference between them is due to absorption or gain in the

material. A material with dielectric constant $\epsilon' - i\epsilon''$ has a *loss tangent* $\delta = \epsilon''/\epsilon'$. In such a material, H lags E in phase by $\frac{1}{2} \arctan \delta$.

1.3 CALCULATING WAVE PROPAGATION IN REAL LIFE

A real optical system is too complicated to be readily described in terms of vector fields, so being practical folk, we look for an appropriate sleazy approximation. We know that in a homogeneous and time-invariant medium, all propagating light waves can be decomposed into their plane wave spectra; at each (temporal) frequency, there will be a unique set of vector plane waves that combine to produce the observed field distributions. The approximations we will use are four:

1. *Scalar Optics*: Replace vector field addition with scalar addition.
2. *Paraxial Propagation*: Use an approximate propagator (the Huyghens integral) to calculate wave propagation, limiting us to beams of small cone angles.
3. *Fourier Optics*: Use a simplistic model for how individual spatial Fourier components on a surface couple to the plane wave components of the incident light.
4. *Ray Optics*: Ignore diffraction by using an asymptotic theory valid as $\lambda \rightarrow 0$.

Analytical models, within their realm of applicability, are so much superior to numerical models in intuitive insight and predictive power that it is worth sacrificing significant amounts of accuracy to get one. Numerical models have their place—but the output is just a pile of special cases, so don't use them as a crutch to avoid hard thinking.

1.3.1 Scalar Optics

If we shine an ideal laser beam (one perfectly collimated and having a rectangular amplitude profile) through a perfect lens and examine the resulting fields near focus, the result is a complete mess. There are nonzero field components along the propagation axis, odd wiggly phase shifts, and so on, due entirely to the wave and vector nature of the fields themselves. The mess becomes worse very rapidly as the sine of the cone angle θ of the beam approaches unity. The effect is aggravated by the fact that no one really knows what a lens does, in sufficient detail to describe it accurately analytically—a real system is a real mess.

For most optical systems, we don't have to worry about that, because empirically it doesn't affect real measurements much. Instead, we use *scalar optics*. Scalar optics is based on the replacement of the six components of the true vector electromagnetic field by a single number, usually thought of as being the electric field component along the (fixed) polarization axis. In an isotropic medium, the vector wave equation admits plane wave solutions whose electric, magnetic, and propagation vectors are constant in space, so that the field components can be considered separately, which leads to the *scalar Helmholtz equation*,

$$(\nabla^2 + k^2)E = 0. \quad (1.17)$$

(This rationale is nothing but a fig leaf, of course.) Any solution of (1.17) can be decomposed in a Fourier function space of plane waves, which are identified with one Cartesian




Figure 1.2. Scalar addition is a good approximation to vector addition except near high-NA foci.

component of the vector field. A scalar plane wave traveling along a direction \mathbf{k} has the form

$$\psi(\mathbf{x}) = e^{i(\mathbf{k}\cdot\mathbf{x} - \omega t)}, \quad (1.18)$$

where the vector \mathbf{k} has length $k = 2\pi/\lambda$. Conceptually, the true vector field can be built up from three sets of these.

The only difficulty with this in free space is that the field vectors of the plane waves are perpendicular to their propagation axes, so that for large numerical aperture,[†] where the propagation axes of the various components are very different, the vector addition of fields near a focus is not too well approximated by a scalar addition. Far from focus, this is not much of a worry, because the components separate spatially, as shown in Figure 1.2.

Aside: Plane Waves and δ -Functions. This separation is not entirely obvious from a plane wave viewpoint, but remember that plane waves are δ -functions in \mathbf{k} -space; that makes them just as singular in their way as δ -functions. They aren't always an aid to intuition. Of course, it is not free space propagation that provides useful data, but the interaction of light with matter; boundaries of physical interest will in general mix the different polarization components. Such mathematical and practical objections are swept under the rug.[‡]

[†]The numerical aperture (NA) of a beam is given by $NA = n \sin \theta$, where n is the refractive index of the medium and θ is the half-angle of the beam cone. By Snell's law, the numerical aperture of a beam crossing an interface between two media at normal incidence remains the same.

[‡]The actual electromagnetic boundary conditions at surfaces of interest are very complicated, and usually poorly understood, so that most of the time the inaccuracies we commit by approximating the boundary conditions are smaller than our ignorance.

Polarization and finite bandwidth effects are usually put in *by hand*. This means that we keep track of the polarization state of the beam separately and follow it through the various optical elements by bookkeeping; for frequency-dependent elements, we keep the frequency as an independent variable and integrate over it at the end. Such a procedure is inelegant and mathematically unjustified, but (as we shall see) it works well in practice, even in regimes such as high numerical aperture and highly asymmetric illumination, in which we would expect it to fail. Everyone in the optical systems design field uses it, and the newcomer would be well advised to follow this wholesome tradition unless driven from it by unusual requirements (and even then, to put up a fight).

1.3.2 Paraxial Propagation

Discussions of beam propagation, pupil functions, optical and coherent transfer functions, and point spread functions take place with reference to the plane wave basis set. There are an infinite variety of such basis sets for decomposition of solutions of the scalar Helmholtz equation, nearly none of which are actually useful. Plane waves are one exception, and the Gauss–Legendre beams are another—or would be if they quite qualified. Gaussian beams (as they are usually called) don't even satisfy the true scalar Helmholtz equation, because their phase fronts are paraboloidal rather than spherical, and because they extend to infinity in both real- and \mathbf{k} - (spatial frequency) space.

Instead, they satisfy the *slowly varying envelope equation*, also known as the *paraxial wave equation*. First, we construct a field as a product of a plane wave e^{ikz} times an envelope function $\Theta(\mathbf{x})$ that varies slowly on the scale of a wavelength. This field is plugged into the scalar Helmholtz equation, the product rule for the Laplacian operator is invoked, and the subdominant term $d^2\Theta/dz^2$ is discarded, leaving a Schrödinger-type equation for the envelope Θ , the paraxial wave equation

$$\frac{d^2\Theta}{dx^2} + \frac{d^2\Theta}{dy^2} + 2ik\frac{d\Theta}{dz} = 0. \quad (1.19)$$

A general solution to this equation for all (x, y, z) is given by the *Huyghens integral*,

$$\Theta(x, y, z) = -\frac{i}{\lambda} \iint_P \Theta(x', y', z') \frac{\exp\left[ik \frac{(x-x')^2 + (y-y')^2}{2(z-z')}\right]}{(z-z')} dx' dy', \quad (1.20)$$

where P is the $x'y'$ plane. In diffraction theory (1.20) is also known as the *Fresnel approximation*. The Huyghens integral is an example of a *propagator*, an integral operator that uses the field values on a surface to predict those in the entire space. It is slightly inconvenient to lose the explicit phase dependence on z , but that can be recovered at the end by calculating the phase of an axial ray (one traveling right down the axis of the system) and adding it in. The Huyghens kernel depends only on $\mathbf{x}-\mathbf{x}'$ and so is a convolution (see Section 1.3.8), leading naturally to a Fourier space (\mathbf{k} -space) interpretation. In \mathbf{k} -space, (1.20) is

$$\Theta(x, y, z) = \iint_{P'} U(u, v) e^{i(2\pi/\lambda)(ux+vy)} e^{-i(2\pi z/\lambda)[(u^2+v^2)/2]} dudv, \quad (1.21)$$

where P' is the uv plane and U is the plane wave spectrum of Θ at $z = 0$, which is given by

$$U(u, v) = \iint_P \Theta(x, y, 0) e^{-i(2\pi/\lambda)(ux+vy)} \frac{dx}{\lambda} \frac{dy}{\lambda} \quad (1.22)$$

The quantities u and v are the *direction cosines* in the x and y directions, respectively, and are related to the *spatial frequencies* k_x and k_y by the relations $u = k_x/k$, $v = k_y/k$. What we're doing here is taking the field apart into plane waves, propagating each wave through a distance z by multiplying by $\exp(ik_z z)$, and putting them back together to get the field distribution at the new plane. The (u, v) coordinates of each component describe its propagation direction. This is a perfectly general procedure.

Aside: Use k -Space. The real-space propagator (1.20) isn't too ugly, but the Rayleigh–Sommerfeld and Kirchhoff propagators we will exhibit in Section 9.3.2 are not easy to use in their real-space form. The procedure of splitting the field apart into plane waves, propagating them, and reassembling the new field is applicable to all these propagators, because in \mathbf{k} -space they differ only slightly (in their *obliquity factors*, of which more later). This is really the right way to go for hand calculations.

It is actually easier to spot what we're doing with the more complicated propagators, because the $\exp(ik_z z)$ appears explicitly. The Huyghens propagator ignores the $\exp(ik_z z)$ factor and uses an approximation for $\exp[iz(k_z - k)]$, which obscures what's really happening.

1.3.3 Gaussian Beams

The Gauss–Legendre beams are particular solutions to the paraxial wave equation. The general form of a zero-order Gauss–Legendre beam that travels along the z axis in the positive direction and whose phase fronts are planar at $z = 0$ is

$$\Phi(x, y, z, t) = \sqrt{\frac{2}{\pi}} \frac{1}{w(z)} \exp \left\{ i\phi(z) + (x^2 + y^2) \left[\frac{-1}{w^2(z)} + \frac{ik}{2R(z)} \right] \right\}, \quad (1.23)$$

where $R(z)$, $\phi(z)$, $w(z)$, z_R , and w_0 are given in Table 1.1. (Remember that the scalar field E is the envelope Φ multiplied by the plane wave “carrier” $e^{ikz - \omega t}$.) These parameters depend only on the beam waist radius w_0 and the wavelength λ of the light in the medium.

The envelope function is complex, which means that it modulates both the amplitude and the phase ϕ of the associated plane wave. This gives rise to the curved wavefronts (surfaces of constant phase) of focused beams, and also to the less well-known variations in $\partial\phi/\partial z$ with focal position, the Gouy phase shift.

Gaussian beams reproduce the ordinary properties of laser beams of small to moderate numerical aperture. They also form a complete set of basis functions, which means that any solution of (1.19) can be described as a sum of Gaussian beams. This useful property should not be allowed to go to the user's head; Gaussian beams have a well-defined axis, and so can only represent beams with the same axis. The number of terms required for a given accuracy and the size of the coefficients both explode as the beam axis departs from that of the eigenfunctions.

TABLE 1.1. TEM₀₀ Gaussian Beam Parameters (Beam Waist at z = 0)

Central intensity	$I_0 = 2P/(\pi w^2)$
Central intensity at the waist	$I_{0W} = 2P/(\pi w_0^2)$
Total power	$P = \pi w^2 I_0/2$
Beam waist radius (power density on axis = I_0/e^2)	$w_0 = \lambda/(\pi \text{NA})$
$1/e^2$ Power density radius	$w(z) = w_0[1 + (z/z_R)^2]^{1/2}$
3 dB Power density radius vs. $1/e^2$ radius	$w_{1/2}(z) = 0.5887 w(z)$
Radius within which $I >$ given I_{th}	$r = w/\sqrt{2} \ln^{1/2}(I_0/I_{\text{th}})$
Power included inside $r \leq w$	86.4%
99% Included power radius	$r_{99} = 1.517 w$
Fourier transform pair	$\exp[-\pi(r/\lambda)^2] \supset \exp(-\pi \sin^2 \theta)$
Separation of variables	$\exp(-\pi(r/\lambda)^2)$ $= \exp[-\pi(x/\lambda)^2] \exp[-\pi(y/\lambda)^2]$
Numerical aperture ($1/e^2$ points in k -space)	$\text{NA} = \lambda/(\pi w_0)$
Radius of curvature of phase fronts	$R(z) = z + z_R^2/z$
Rayleigh range (axial intensity 50% of peak)	$z_R = \pi w_0^2/\lambda = \lambda/(\pi(\text{NA})^2)$
Displacement of waist from geometric focus	$\Delta z \approx -z_R^2/f$
Envelope phase shift	$\phi(z) = \tan^{-1}(z/z_R)$
Equivalent projected solid angle	$\Omega'_{\text{eq}} = \pi(\text{NA})^2 = \lambda^2/(\pi w_0^2)$

In this book, as in most of practical electro-optical instrument design, only this lowest-order mode, called TEM₀₀, is needed. This mode describes the field distribution of a good quality laser beam, such as that from a HeNe or circularized diode laser.

At large z , the Gaussian beam looks like a spherical wave with a Gaussian cutoff in u and v , but for small z , it appears to be a collimated beam. The distance, called z_R or the Rayleigh range, over which the beam stays approximately collimated goes as $1/(\text{NA})^2$ —the beam waist goes as $1/\text{NA}$ and the the angular width as NA . At $z = \pm z_R$, the $1/e^2$ beam diameter has increased by a factor of $\sqrt{2}$, so that the central intensity has halved.

The Gaussian beam is a paraxial animal: it’s hard to make good ones of high NA. Its extreme smoothness makes it exquisitely sensitive to vignetting, which of course becomes inevitable as $\sin \theta$ approaches 1, and the slowly varying envelope approximation itself breaks down as the numerical aperture increases (see Example 9.8).

There are a variety of parameters of Gaussian beams which are frequently of use, some of which are summarized in Table 1.1; P is the total power in watts, I is the intensity in W/m^2 , w is the $1/e^2$ intensity radius, w_0 is the beam waist radius, z_R is the Rayleigh range, and NA is measured at the $1/e^2$ intensity points in **k**-space. Of interest in applications is the envelope phase, which shows a $\pm\pi/4$ phase shift (beyond the plane wave’s $\exp(ikz)$) over the full depth of focus (twice the Rayleigh range), so that in a focused beam it is not a good assumption that the phase is simply $\exp(ikz)$. This phase is exploited in phase contrast systems such as the Smartt interferometer.

Aside: Gaussian Beams and Lenses. When a Gaussian beam passes through a lens, it is transformed into a different Gaussian beam. For the most part, ray optics is sufficient to predict the position of the beam waist and the numerical aperture, from which the waist radius and Rayleigh range can be predicted. There are some useful invariants of

this process: for example, a Gaussian beam whose waist scans back and forth by b waist radii will be transformed into another beam whose waist scans b times the new waist radius. A position c times the Rayleigh range from the waist will image to a point c times the new Rayleigh range from the new waist. A corollary is that the number of *resolvable spots*, that is, the scan range divided by the spot diameter, is also invariant. These invariants, which are not limited to the Gaussian case, allow one to juggle spot sizes, focal positions, and scan angles freely, without having to follow them laboriously through the whole optical system.

1.3.4 The Debye Approximation, Fresnel Zones, and Fresnel Number

The plane wave decomposition of a given disturbance can be calculated from (1.22) or its higher-NA brethren in Section 9.3.6, and those work regardless of where we put the observation plane. When discussing the NA of a lens, however, we usually use a much simpler method: draw rays representing the edges of the beam and set $\text{NA} = n \sin \theta$. This sensible approach, the *Debye approximation*, obviously requires the beam to be well represented by geometric optics, because otherwise we can't draw the rays—it breaks down if you put the aperture near a focus, for instance. We can crispen this up considerably via the *Fresnel construction*.

In a spherical wave, the surfaces of constant phase are equally spaced concentric hemispheres, so on a plane, the lines of constant phase are concentric circles, corresponding to annular cones, as shown in Figure 1.3. Drawing these circles at multiples of π radians divides the plane into annular regions of positive and negative field contributions, called

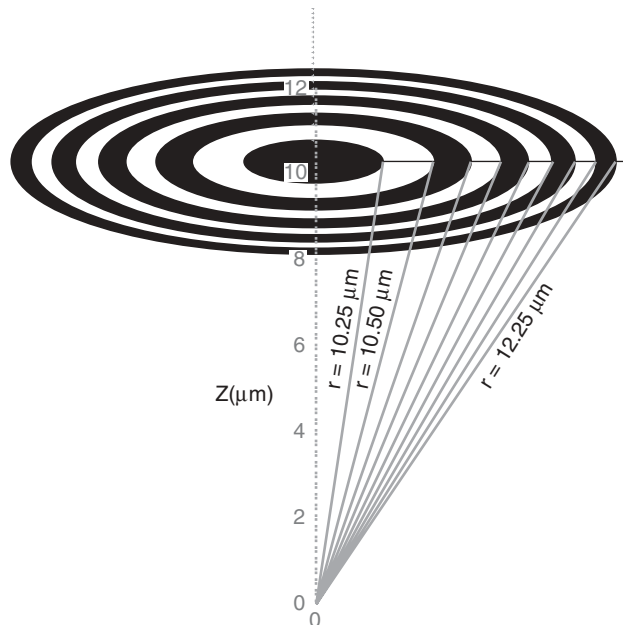


Figure 1.3. The Fresnel zone construction with $f = 10 \mu\text{m}$ and $\lambda = 0.5 \mu\text{m}$. For a plane wave, taking the phase on axis as 0, alternating rings produce positive and negative field contributions at f , so blocking alternate ones (or inverting their phases with a $\lambda/2$ coating) produces a focus at f . For a converging spherical wave, all zones produce positive contributions at the focus.

Fresnel zones. The zones are not equally spaced; for a beam whose axis is along $\hat{\mathbf{z}}$ and whose focus is at $z = 0$, the angular zone boundaries in the far field are at

$$\theta_n = \cos^{-1} \frac{1}{1 + (2n + 1)\lambda/(4f)} \quad (1.24)$$

(the equation for the N th zone center is the same, except with $2N$ instead of $(2n + 1)$).

The *Fresnel number* N is the number of these zones that are illuminated. This number largely determines the character of the beam in the vicinity of the reference point—whether it is dominated by diffraction or by geometric optics. The Debye approximation is valid in the geometric limit, that is, $N \gg 1$.

Taking r to be the radius of the illuminated circle and applying a couple of trigonometric identities to (1.24) gets us a quadratic equation for N , the number of annular zone centers falling inside r . Assuming that $N\lambda \ll f$, this simplifies into

$$N = \frac{r^2}{\lambda z}, \quad (1.25)$$

which due to its simplicity is the usual definition of Fresnel number.

In small Fresnel-number situations, the focus is displaced toward the lens from its geometric position and diffraction is important everywhere, not just at the focus. The number of resolvable spots seen through an aperture of radius r is $N/2$.

Example 1.2: Gaussian Beams and Diffraction. For small numerical apertures, the position of the beam waist does not coincide with the geometric focus, but is closer. This somewhat counterintuitive fact can be illustrated by considering a beam 40λ in radius, with a lens of $10^4\lambda$ focal length placed at its waist. The lens changes the Gaussian beam parameters, as we can calculate. Geometrically, the incoming beam is collimated, so the focus is $10^4\lambda$ away, but in reality the Rayleigh range of the beam is $40(\pi/4)$ spot diameters, or 1260λ . This is only $1/8$ of the geometric focal length, so the lens makes only a small perturbation on the normal diffractive behavior of the original beam. At the geometric focus, $N = 40^2/10^4 = 0.16$, so the total phase change due to the lens is only $\pi/6$ across the beam waist.

1.3.5 Ray Optics

We all know that the way you check a board for warpage is by sighting along it, because light in a homogeneous medium travels in straight lines. The departures of light from straight-line propagation arise from nonuniformities in the medium (as in mirages) and from diffraction. Most of the time these are both small effects and light can be well described by *rays*, thought of as vanishingly thin pencil beams whose position and direction are both well defined—the usual mild weirdness exhibited by asymptotic theories. Ray optics does not require the paraxial approximation, or even scalar waves.

In the absence of diffraction (i.e., as $\lambda \rightarrow 0$), the direction of propagation of a light beam in an isotropic medium is parallel to the gradient of the phase[†] $\nabla\phi$ (see

[†]M. Born and E. Wolf, *Principles of Optics*, 6th ed. (corrected). Pergamon, Oxford, 1983, p. 112.

Section 9.2.3). This means that a beam whose phase fronts are curved is either converging or diverging (see Section 9.2.2) and that rays can be identified with the normals to the phase fronts. Rays are the basis of elementary imaging calculations, as in the following example.

Example 1.3: Imaging with a Camera Lens. As a simple example of the use of ray optics, consider using a 35 mm camera to take a head-and-shoulders portrait of a friend. For portraits, the most pleasing perspective occurs with a camera-to-subject distance of a few feet, 4 feet (1.3 m) being about optimal. What focal length lens is required?

The film frame is 24 by 36 mm in size and the outline of a human head and shoulders is about 400 by 500 mm. Thus the desired magnification is 24/400, or 0.06. The rules of thin-lens optics are:

1. Rays passing through the center of the lens are undeviated.
2. Rays entering parallel to the axis pass through the focus.
3. The locus of ray bending is the plane of the center of the lens.

All of these rules can be fixed up for the thick-lens case (see Section 4.11.2). The similar triangles in Figure 1.4 show that the magnification M is

$$M = \frac{d_i}{d_o} \quad (1.26)$$

and elementary manipulation of the geometric identities shown yields

$$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f} \quad (1.27)$$

and

$$s_o s_i = f^2. \quad (1.28)$$




Figure 1.4. Portraiture with a 35 mm camera.

Using (1.26) and (1.27), we find that

$$f = \frac{Md_o}{1 + M} \quad (1.29)$$

or 73.5 mm. Since the optimum distance is only a rough concept, we can say that a portrait lens for 35 mm photography should have a focal length of around 70 to 80 mm.

We can associate a phase with each ray, by calculating the phase shift of a plane wave traversing the same path. In doing this, we have to ignore surface curvature, in order that the wave remain plane. A shorthand for this is to add up the distances the ray traverses in each medium (e.g., air or glass) and multiply by the appropriate values of k .

1.3.6 Lenses

In the wave picture, an ideal lens of focal length f transforms a plane wave $e^{ik(ux+vy)}$ into a converging spherical wave, whose center of curvature is at (uf, vf) . It does so by inserting a spatially dependent phase delay, due to propagation through different thicknesses of glass. In the paraxial picture, this corresponds to a real-space multiplication by

$$L(x, y : f) = \exp \left[\frac{i\pi}{\lambda f} (x^2 + y^2) \right]. \quad (1.30)$$

Example 1.4: A Lens as a Fourier Transformer. As an example of how to use the Huyghens propagator with lenses, consider a general field $\Theta(x, y, -f)$ a distance f behind a lens whose focal length is also f . The operators must be applied in the order the fields encounter them; here, the order is free-space propagation through f , followed by the lens's quadratic phase delay (1.30) and another free-space propagation through f . Using (1.20) twice, the field becomes

$$\Theta(x, y, +f) = \frac{i\lambda}{f} \int_{-\infty}^{\infty} d\frac{x'}{\lambda} \int_{-\infty}^{\infty} d\frac{y'}{\lambda} e^{-i(2\pi/\lambda f)(xx'+yy')} \Theta(x', y', -f), \quad (1.31)$$

which is a pure scaled Fourier transform. Thus a lens performs a Fourier transform between two planes at $z = \pm f$. If we put two such lenses a distance $2f$ apart, as shown in Figure 1.5, then the fields at the input plane are reproduced at the output plane, with a Fourier transform plane in between. The image is inverted, because we've applied two forward transforms instead of a forward ($-i$) followed by a reverse ($+i$) transform, so $(x, y) \rightarrow (-x, -y)$.

If we put some partially transmitting mask at the transform plane, we are blocking some Fourier components of Θ , while allowing others to pass. Mathematically, we are multiplying the Fourier transform of Θ by the amplitude transmission coefficient of the mask, which is the same as convolving it with the Fourier transform of the mask, appropriately scaled. This operation is called *spatial filtering* and is widely used.

The Fourier transforming property of lenses is extremely useful in both theory and applications. Perhaps surprisingly, it is not limited to the paraxial case, as we will see below.




Figure 1.5. Spatial filtering.

1.3.7 Aperture, Field Angle, and Stops

Not every ray that enters an optical system will make it out the other side. At most locations in a system, there is no sharp boundary. At a given point in space, light going in some directions will make it and that going in other directions will not; similarly for a given angle, there may be part of the system where it can pass and part where it cannot. However, each ray that fails to make it will wind up hitting some opaque surface. The surface that most limits the spatial field of a ray parallel to the axis is called the *field stop* and that which limits the angular acceptance of a point on the axis most, the *aperture stop*. At these surfaces, the boundary between blocked and transmitted components is sharp. These locations are shown in Figure 1.6. It is common to put the aperture stop at a Fourier transform plane, since then all points on the object are viewed from the same range of angles. Optical systems image a volume into a volume, not just

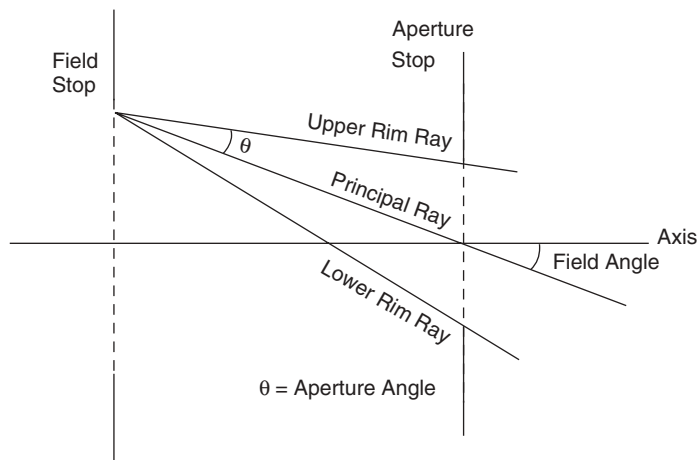


Figure 1.6. Definitions of aperture and field angle.

a plane into a plane, so the stops can't always be at the exact image and transform planes.

Aside: Vignetting. Aperture and field are defined in terms of axial points and axial rays. There's no guarantee that the aperture and field are exactly the same for other points and other directions. Rays that get occluded somewhere other than the field or aperture stops are said to have been *vignetted*. Vignetting isn't always bad—it's commonly used to get rid of badly aberrated rays, which would degrade the image if they weren't intercepted.

When a laser beam hits the edge of an aperture, it is also loosely termed vignetting, even when it does happen at one of the stops.

1.3.8 Fourier Transform Relations

Fourier transforms crop up all the time in imaging theory. They are a common source of frustration. We forget the transform of a common function or can't figure out how to scale it correctly, and what was a tool becomes a roadblock. This is a pity, because Fourier transforms are both powerful and intuitive, once you have memorized a couple of basic facts and a few theorems. In an effort to reduce this confusion, here are a few things to remember. Following Bracewell, we put the factors of 2π in the exponents and write $g \supset G$ and $G = \text{F}g$ for “ g has transform G .”

$$G(f) = \int_{-\infty}^{\infty} g(x)e^{-i2\pi fx} dx, \quad (1.32)$$

$$g(x) = \int_{-\infty}^{\infty} G(f)e^{i2\pi fx} df. \quad (1.33)$$

A side benefit of this is that we work exclusively in units of cycles. One pitfall is that since for a wave traveling in the positive direction, the x and t terms in the exponent have opposite signs, so it is easy to get mixed up about forward and inverse transforms. Physicists and electrical engineers typically use opposite sign conventions.

Useful Functions. The Heaviside unit step function $U(x)$ is 0 for $x < 0$ and 1 for $x > 0$. The derivative of $U(x)$ is the Dirac δ -function, $\delta(x)$. Sinc and jinc functions come up in connection with uniform beams: $\text{sinc}(x) = \sin(\pi x)/(\pi x)$ and $\text{jinc}(x) = J_1(2\pi x)/(\pi x)$. Even and odd impulse pairs $\mathbf{II}(x) = [\delta(x - \frac{1}{2}) + \delta(x + \frac{1}{2})]/2$ and $\mathbf{I}_1(x) = [\delta(x + \frac{1}{2}) - \delta(x - \frac{1}{2})]/2$ have transforms $\cos(\pi f)$ and $i \sin(\pi f)$, respectively.

Conjugacy. Conjugate variables are those that appear multiplied together in the kernel of the Fourier transform, such as time in seconds and frequency in hertz. In optical Fourier transforms, the conjugate variables are x/λ and u , which is as before the direction cosine of the plane wave component on the given surface, that is, $u = k_x/k$.

Convolution. A convolution is the mathematical description of what a filter does in the real time or space domain, namely, a moving average. If $g(x)$ is a given data stream and $h(x)$ is the impulse response of a filter (e.g., a Butterworth lowpass electrical filter, with x standing for time):

$$h(x) * g(x) = \int_{-\infty}^{\infty} h(\xi)g(x - \xi)d\xi = \int_{-\infty}^{\infty} g(\xi)h(x - \xi)d\xi. \quad (1.34)$$

The second integral in (1.34) is obtained by the transformation $u \rightarrow x - \xi$; it shows that convolution is commutative: $g * h = h * g$. Convolution in the time domain is multiplication in the frequency domain,

$$\mathcal{F}(h * g) = HG, \quad (1.35)$$

where capitals denote transforms, for example, $G(f) = \mathcal{F}(g(x))$. This makes things clearer: since multiplication is commutative, convolution must be too. A lot of imaging operations involve convolutions between a *point spread function* (impulse response) and the sample surface reflection coefficient (coherent case) or reflectance (incoherent case). The Huyghens propagator is also a convolution. Note that one of the functions is flipped horizontally before the two are shifted, multiplied, and integrated. This apparently trivial point in fact has deep consequences for the phase information, as we'll see in a moment. The convolution theorem is also very useful for finding the transform of a function, which looks like what you want, by cobbling together transforms that you know (see Example 1.5).

Symmetry. By a change of variable in the Fourier integral, you can show that $g(-x) \supset G(-f)$, $g^*(x) \supset G^*(-f)$, $g^*(-x) \supset G^*(f)$, and (if g is real) $G(-f) = G^*(f)$.

Correlation and Power Spectrum. The cross-correlation $g \star h$ between functions g and h is the convolution of $g(x)$ and $h^*(-x)$: $g \star h = g(x) * h^*(-x) \supset GH^*$. This can also be shown by a change of variables.

An important species of correlation is the autocorrelation, $g \star g$, whose transform is $GG^* = |G|^2$, the power spectrum. The autocorrelation always achieves its maximum value at zero (this is an elementary consequence of the Schwarz inequality) and all phase information about the Fourier components of g is lost.

Equivalent Width. We often talk about the width of a function or its transform. There are lots of different widths in common use; 3 dB width, $1/e^2$ width, the Rayleigh criterion, and so on. When we come to make precise statements about the relative widths of functions and their transforms, we talk in terms of equivalent width or sometimes autocorrelation width. The equivalent width of a function is

$$w_e(g) = \frac{\int_{-\infty}^{\infty} g(x') dx'}{g(0)} = \frac{G(0)}{g(0)}. \quad (1.36)$$

It is obvious from this that if g and G are nonzero at the origin, the equivalent width of a function is the reciprocal of that of its transform. It is this relationship that allows us to say airily that a 10-wavelength-wide aperture has an angular spectrum 0.1 rad wide.

Functions having most of their energy far from zero are not well described by an equivalent width. For example, if we move the same aperture out to $x = 200\lambda$, it will have a very large equivalent width (since $g(0)$ is very small), even though the aperture itself hasn't actually gotten any wider. Such a function is best described either by quoting its *autocorrelation width*, which is the equivalent width of the autocorrelation $g \star g$, or by shifting it to the origin. (We commonly remove the tilt from a measured wavefront, which is equivalent to a lateral shift of the focus to the origin.) Autocorrelations always achieve their maximum values at zero. Since the transform of the autocorrelation is the

power spectrum, the autocorrelation width is the reciprocal of the equivalent width of the power spectrum.

Shifting. Given $g(x)$, then shifting the function to the right by x_0 corresponds to subtracting x_0 from the argument. If $g(x)$ is represented as a sum of sinusoids, shifting it this way will phase shift a component at frequency f by $f x_0$ cycles:

$$\mathcal{F}(g(x - x_0)) = e^{-i2\pi f x_0} G(f). \quad (1.37)$$

Scaling. A feature 10λ wide has a transform 0.1 rad wide (in the small-angle approximation). Making it a times narrower in one dimension without changing its amplitude makes the transform a times wider in the same direction and a times smaller in height, without changing anything in the perpendicular direction:

$$g(at) \supset \frac{1}{|a|} G\left[\frac{f}{a}\right]. \quad (1.38)$$

You can check this by noting that the value of the transform at the origin is just the integral over all space of the function.

Integrating and Differentiating. For differentiable functions, if $g \supset G$, then

$$\frac{dg}{dx} \supset i2\pi f G, \quad (1.39)$$

which is easily verified by integrating by parts. If g is absolutely integrable, then

$$\int_{-\infty}^x g \, dx' \supset \frac{G}{i2\pi f} + K\delta(f), \quad (1.40)$$

where K is an arbitrary integration constant. It follows from (1.39) that the derivative of a convolution is given by

$$\frac{d}{dx}(h * g) = h * \frac{dg}{dx} = g * \frac{dh}{dx} \supset i2\pi f GH. \quad (1.41)$$

Power Theorem. If we compute the central value of the cross-correlation of g and h , we get the odd-looking but very useful *power theorem*:

$$\int_{-\infty}^{\infty} dx \, g(x)h^*(x) = \int_{-\infty}^{\infty} df \, G(f)H^*(f) \quad (1.42)$$

(e.g., think of g as voltage and h as current). With the choice $g = h$, this becomes *Rayleigh's theorem*,[†]

$$\int_{-\infty}^{\infty} dx |g(x)|^2 = \int_{-\infty}^{\infty} df |G(f)|^2, \quad (1.43)$$

which says that the function and its transform have equal energy. This is physically obvious when it comes to lenses, of course.

[†]The same relationship in Fourier series is Parseval's theorem.

Asymptotic Behavior. Finite energy transforms have to fall off eventually at high frequencies, and it is useful to know how they behave as $f \rightarrow \infty$. A good rule of thumb is that if the n th derivative of the function leads to delta functions, the transform will die off as $1/f^n$. You can see this by repeatedly using the formula for the transform of a derivative until you reach delta functions, whose transforms are asymptotically constant in amplitude.

Transform Pairs. Figure 1.7 is a short gallery of Fourier transform pairs.

Example 1.5: Cobbling Together Transforms. In analyzing systems, we often need a function with certain given properties, but don't care too much about its exact identity, as long as it is easy to work with and we don't have to work too hard to find its transform. For example, we might need a function with a flat top, decreasing smoothly to zero on

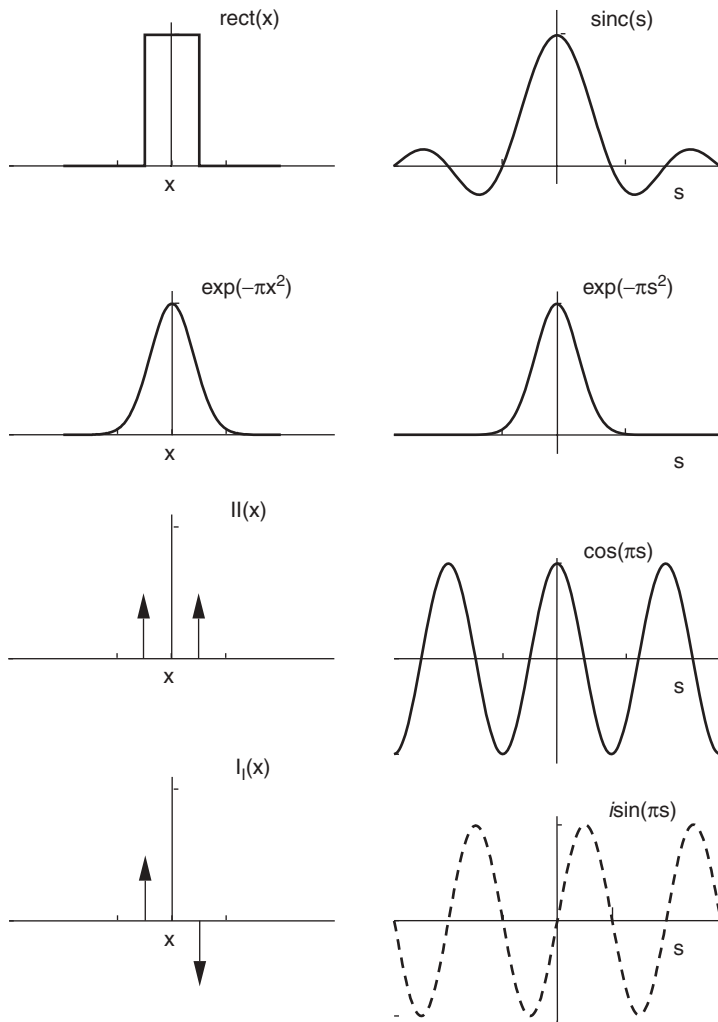


Figure 1.7. A pictorial gallery of Fourier transform pairs. Bracewell has lots more.

both sides, to represent a time gating operation of width t_g followed by a filter. The gross behavior of the operation does not depend strongly on minor departures from ideal filtering, so it is reasonable to model this function as the convolution of $\text{rect}(t/t_g)$ with a Gaussian:

$$m(t) = \exp[-\pi(t/\tau)^2] * \text{rect}\left[\frac{t}{t_g}\right], \quad (1.44)$$

whose transform is

$$M(f) = \tau t_g e^{-\pi(f\tau)^2} \text{sinc}(ft_g). \quad (1.45)$$

One can write $m(t)$ as the difference of two error functions, but the nice algebraic properties of convolutions make the decomposed form (1.44) more useful.

1.3.9 Fourier Imaging

We have seen that a lens performs a Fourier transform between its front and back focal planes, and that in \mathbf{k} -space, the propagation operator involves Fourier decomposing the beam, phase shifting the components, and reassembling them. There is thus a deep connection between the imaging action of lenses and Fourier transforms. Calculating the behavior of an imaging system is a matter of constructing an integral operator for the system by cascading a series of lenses and free-space propagators, then simplifying. Nobody actually does it that way, because it can easily run to 20th-order integrals. In a system without aberrations, we can just use ray optics to get the imaging properties, such as the focal position and numerical aperture, and then use at most three double integrals to get the actual fields, as in Example 1.4.

Most of the time, we are discussing imaging of objects that are not self-luminous, so that they must be externally illuminated. Usually, we accept the restriction to *thin objects*—ones where multiple scattering can be ignored and the surface does not go in and out of focus with lateral position. The reasoning goes as follows: we assume that our incoming light has some simple form $E_{\text{in}}(x, y)$, such as a plane wave. We imagine that this plane wave encounters a surface that has an amplitude reflection coefficient $\rho(x, y)$, which may depend on position, but not on the angle of incidence, so that the outgoing wave is

$$E_{\text{out}}(x, y) = E_{\text{in}}(x, y)\rho(x, y), \quad (1.46)$$

and then we apply the Huyghens integral to E_{out} .

Small changes in height (within the depth of focus) are modeled as changes in the phase of the reflection coefficient. Since different plane wave components have different values of k_z , we apply a weighted average of the k_z values over the pupil function. The breakdown of this procedure due to the differences in $k_z z$ becoming comparable to a cycle gives rise to the limits of the depth of focus of the beam. We ignore the possibility that the height of the surface might be multiple-valued (e.g., a cliff or overhang) and any geometric shadowing.

A very convenient feature of this model, the one that gives it its name, is the simple way we can predict the angular spectrum of the scattered light from the Fourier transform of the sample's complex reflection $\rho(x, y)$. The outgoing wave in real space is the product ρE_{in} , so in Fourier space,

$$E_{\text{out}}(u, v) = E_{\text{in}}(u, v) * P(u, v) \quad (1.47)$$

where $\rho(x/\lambda, y/\lambda) \supset P$. The real power of this is that $E_{\text{out}}(u, v)$ is also the angular spectrum of the outgoing field, so that we can predict the scattering behavior of a thin sample with any illumination we like.

If E_{in} is a plane wave, $E_{\text{in}}(x, y) = \exp[i2\pi(u_{\text{in}}x + v_{\text{in}}y)/\lambda]$, then its transform is very simple: $E_{\text{in}}(u, v) = \delta(u - u_{\text{in}})\delta(v - v_{\text{in}})$. Convolution with a shifted delta function performs a shift, so

$$E_{\text{out}}(u, v) = E_{\text{in}}(u - u_{\text{in}}, v - v_{\text{in}}). \quad (1.48)$$

The angular spectrum of E_{out} is the spatial frequency spectrum of the sample, shifted by the spatial frequency of the illumination—the spatial frequencies add. In an imaging system, the spatial frequency cutoff occurs when an incoming wave with the largest positive u is scattered into the largest negative u the imaging lens can accept.[†] Since $u^2 + v^2 \leq (\text{NA})^2$, if the NAs of the illumination and the collecting lenses are equal, the highest spatial frequency an imaging system can accept is $2 \text{NA}/\lambda$.

The conceptual deficiencies of this procedure are considerable, even with thin objects. It works fine for large holes punched in a thin plane screen, but for more complicated objects, such as transparent screens containing phase objects (e.g., microscope slides), screens with small features, or nearly anything viewed in reflection, the approximations become somewhat scallier. The conceptual problem arises right at the beginning, when we assume that we know *a priori* the outgoing field distributions at the boundary.

There is no real material that, even when uniform, really has reflection or transmission coefficients independent of angle and polarization at optical frequencies, and the situation is only made worse by material nonuniformity and topography. This and the scalar approximation are the most problematic assumptions of Fourier optics; paraxial propagation is a convenience in calculations and not a fundamental limitation (see Section 9.3.5).

1.3.10 The Pupil

As anyone who has ever been frustrated by an out-of-focus movie knows, the image plane of an optical system is rather special and easily missed. Some other special places in an optical system are less well known. The most important of these is the *pupil*, which is an image of the aperture stop. If we look into the optical system from the object side, we see the *entrance pupil*. Looking from the image side, we see the *exit pupil*. By moving from side to side, we can locate the position in space of a pupil by how it moves in response. (This is the same way we tell how far away anything appears.)

There's nothing magical about pupils, although they are talked about in terms that may confuse newcomers—they really are just places in an optical system, which can be imaged where you want them and otherwise manipulated just as a focal plane can.

The aperture stop is usually put at the Fourier transform plane, to avoid nonuniform vignetting. The field distribution at a pupil then is the Fourier transform of that at the object or an image, appropriately scaled and with an obliquity correction. Equivalently, the field function at the transform plane is a scaled replica of the far-field diffraction

[†]There's nothing special about the choice of axes, so the limiting resolution might be different along y or at other angles.

pattern of the object, as derived using the Huyghens integral (1.20). In an imaging system, the propagator is a convolution in space, so the imaging properties are controlled by the illumination pattern and detector sensitivity function at the transform plane. Since the transform plane is usually at the pupil, these are loosely called *pupil functions*, and are two-dimensional versions of the complex frequency response of an electronic system.[†] (They are still called pupil functions even when the transform plane is not at the pupil. Laziness is the father of invention.)

Aside: Perspective. The center of the entrance pupil (or really, of the Fourier transform plane in the object space) is the center of perspective. If you’re trying to make a panoramic view using an image mosaic, you’ll want both foreground and background objects to have the same perspective—because otherwise, the positions of the joints between mosaic elements would have to be different depending on the distance. You can accomplish this by rotating the camera around its center of perspective.

1.3.11 Connecting Wave and Ray Optics: *ABCD* Matrices

This section could be subtitled “How to combine optical elements without drowning in multiple integrals.” In an optical system consisting of lenses, mirrors, and free-space propagation, it is possible to model the paraxial imaging properties by means of very simple transformation matrices, one for each element or air space, which are multiplied together to form a combined operator that models the entire system. Here we shall discuss the 2×2 case, appropriate for axially symmetric systems or for systems of cylindrical lenses whose axes are aligned. Generalization to 4×4 matrices is straightforward but more laborious.

In the small-angle approximation (where $\sin \theta \approx \theta$), a ray at height x above the optical axis and propagating at an angle θ measured counterclockwise from the optical axis is represented by a column vector $(x, \theta)^T$, and it transforms as

$$\begin{bmatrix} x \\ \theta \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x' \\ \theta' \end{bmatrix}, \quad (1.49)$$

where the matrix $abcd$ is the ordered product of the *ABCD* matrices of the individual elements. Let’s do an example to see how this works.

Example 1.6: Deriving the *ABCD* Matrix for a Thin Lens. In the ray tracing section, we saw that a thin lens brings all rays entering parallel to the axis and that a ray passing through the center of the lens is undeviated. We can use these facts to derive the *ABCD* matrix for a thin lens, as shown in Figure 1.8. The undeviated central ray, $(0, \theta)^T$ is unchanged, so element *B* must be zero and element *D* must be 1. The ray parallel to the axis, $(1, 0)^T$, remains at the same height immediately following the lens, so that element

[†]The analogy depends on the Debye approximation, so the exponential in/exponential out property of linear systems doesn’t hold as accurately in Fourier optics as in most circuits, but it’s still pretty good if the Fresnel number is high.




Figure 1.8. Action of a lens, for deriving its $ABCD$ matrix.

TABLE 1.2. $ABCD$ Matrices for Common Operations


Free-space propagation through distance z	$\begin{bmatrix} 1 & z \\ 0 & 1 \end{bmatrix}$
Thin lens of focal length f	$\begin{bmatrix} 1 & 0 \\ -1/f & 1 \end{bmatrix}$
Magnification by M	$\begin{bmatrix} M & 0 \\ 0 & 1/M \end{bmatrix}$
Fourier transform	$\begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$

A is also 1. However, it is bent so as to cross the axis at f , so element C must be $-1/f$. Thus a thin lens has an $ABCD$ matrix given in Table 1.2.

Optical layouts conventionally have light going from left to right, whereas matrix multiplication goes right to left. Thus we have to write the matrix product backwards: the $ABCD$ matrix of the first element encountered by the beam goes at the right, with subsequent operations left-multiplying it in succession, as in Figure 1.9.

It is straightforward to extend this formalism to small deviations from given angles of incidence, for example, oblique reflection from a spherical mirror; when doing that, however, excellent drawings are required to avoid confusion about just what is going on.

Example 1.7: Portraiture Calculation Using $ABCD$ Matrices. Example 1.3 demonstrated how to find elementary imaging parameters such as magnification and focal length rapidly using the thin-lens rules on rays passing through the center of the lens and rays passing through the focus. Let us follow the path of a more general paraxial ray using $ABCD$ matrices. We note first that the light from the object propagates through $d_o = 1300$ mm of free space, then a thin lens of focal length $f = 73.5$ mm, and finally another free-space propagation through a distance d_i . The column vector representing the



$$\begin{aligned}
 \begin{bmatrix} 1 & d_i \\ 0 & 1 \end{bmatrix} & \begin{bmatrix} 1 & 0 \\ -1/f & 1 \end{bmatrix} \begin{bmatrix} 1 & d_o \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1-d_i/f & d_o + d_i - d_o d_i/f \\ -1/f & 1 - d_o/f \end{bmatrix} \\
 & = \begin{bmatrix} -d_i/d_o & 0 \\ -1/f & -d_o/d_i \end{bmatrix} = \begin{bmatrix} -M & 0 \\ -1/f & -1/M \end{bmatrix}
 \end{aligned}$$

Figure 1.9. Imaging geometry with ray vectors and $ABCD$ matrices: rays $(h, \theta)^T$ are successively multiplied by $ABCD$ matrices corresponding to free space d_o , a lens of focal length f , and free space d_i . For the imaging condition, $1/d_o + 1/d_i = 1/f$, which makes the last three equalities true.

ray must be acted on by the matrix operators (written in reverse order as already noted):

$$\begin{aligned}
 \begin{bmatrix} x' \\ \theta' \end{bmatrix} &= \begin{bmatrix} 1 & d_i \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/f & 1 \end{bmatrix} \begin{bmatrix} 1 & d_o \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ \theta \end{bmatrix} \\
 &= \begin{bmatrix} 1 - d_i/f & d_o + d_i - d_i d_o/f \\ -1/f & 1 - d_o/f \end{bmatrix} \begin{bmatrix} x \\ \theta \end{bmatrix} \\
 &= \begin{bmatrix} 1 - 0.0136d_i & 1300 + 18.68d_i \\ -0.0136 & -16.68 \end{bmatrix} \begin{bmatrix} x \\ \theta \end{bmatrix}.
 \end{aligned} \tag{1.50}$$

Comparing the algebraic form of the matrix product in (1.50) to the prototypes in Table 1.2, it is apparent that the combination of a lens plus free space on either side behaves as a Fourier transformer (scaled by a magnification of $-f$) when $d_o = d_i = f$. Furthermore, the imaging condition demands that all rays leaving an object point coincide at the same image point; this means that b , the (1, 2) element of the matrix, must be zero, which reproduces (1.27). These sorts of considerations are very valuable for more complex systems, where the thin-lens ray picture is cumbersome.

Useful as these are, matrix multiplication is not sufficiently powerful to model such elementary operations as the addition of a thin prism of angle ϕ and index n , which requires adding an angle $\Delta\theta = (n - 1)\phi$. These matrix operators ignore wave effects and are completely unable to cope with absorbing or partially scattering objects such as beamsplitters and diffraction gratings. While these can of course be put in by hand, a more general operator algebra is desirable, which would take account of the wave nature of the light and model light beams more faithfully.

Nazarathy and Shamir[†] have produced a suitable operator algebra for Fourier optics. The key simplification in use is that they have published a multiplication table for these operators, which allows easy algebraic simplification of what are otherwise horrible high order multiple integrals. These transformations are in principle easy to automate and could be packaged as an add-on to symbolic math packages. This algebra takes advantage of the fact that commonly encountered objects such as lenses, gratings, mirrors, prisms, and transparencies can be modeled as operator multiplication in the complex field representation, which (as we have seen earlier) many cannot be so modeled in the ray representation.

Another way of coming at this is to use $ABCD$ matrices for the operator algebra, and then convert the final result to a Huyghens integral. In the paraxial picture, an axisymmetric, unaberrated, unvignetted optical system consisting of lenses and free space can be expressed as a single $ABCD$ matrix, and any $ABCD$ matrix with $d \neq 0$ can be decomposed into a magnification followed by a lens followed by free space:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 1 & z \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/f & 1 \end{bmatrix} \begin{bmatrix} M & 0 \\ 0 & 1/M \end{bmatrix}, \quad (1.51)$$

where

$$M = \frac{1}{d}, \quad z = \frac{b}{d}, \quad f = -\frac{1}{cd}. \quad (1.52)$$

Element a does not appear because that degree of freedom is used up to ensure that the determinant of the matrix is unity, as required by the conservation of phase space volume. A magnification by M corresponds to the integral operator

$$\Theta(x, y) = \frac{1}{M} \iint dx' dy' \Theta(x', y') \delta \left[x' - \frac{x}{M} \right] \delta \left[y' - \frac{y}{M} \right]. \quad (1.53)$$

Identifying these matrix operators with the corresponding paraxial integral operators (1.20), (1.30), and (1.53), we can construct the equivalent integral operator to a general $ABCD$ matrix with $d \neq 0$:

$$\Theta(x, y) = \frac{-i}{zM\lambda} \iint dx' dy' \Theta \left[\frac{x'}{M}, \frac{y'}{M} \right] \exp \left[\frac{i\pi}{\lambda} \left(\frac{(x - x')^2 + (y - y')^2}{z} + \frac{x'^2 + y'^2}{f} \right) \right]. \quad (1.54)$$

[†]M. Nazarathy and J. Shamir, First-order optics—a canonical operator representing lossless systems. *J. Opt. Soc. Am.* **72**, 356–364 (March 1982).

This transformation is simple, and it can save a *lot* of ugly integrals. The special case where $d = 0$ corresponds to a lens, followed by a scaled Fourier transform:

$$\begin{bmatrix} a & b \\ -1/b & 0 \end{bmatrix} = \begin{bmatrix} M & 0 \\ 0 & 1/M \end{bmatrix} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/f & 1 \end{bmatrix}. \quad (1.55)$$

The constraint that $c = -1/b$ keeps the determinant 1, as before. Here the parameters are $M = -b$, $f = -b/a$ so that the equivalent integral in the wave picture is

$$\Theta(x, y) = \frac{-iM}{\lambda} \iint dx' dy' \Theta(Mx', My') \exp \left[\frac{i\pi M}{\lambda} \left(xx' + yy' + \frac{x'^2 + y'^2}{f} \right) \right]. \quad (1.56)$$

These two equivalences allow wave and ray descriptions of the same optical system to be freely interchanged, which is very convenient in calculations.

The problem of offsets, both in position and angle, can be dealt with by using the augmented vectors $[x, \theta, 1]^T$ and 3×3 matrices. A general element producing a transformation $ABCD$ and then adding an offset $[\Delta x, \Delta \theta]^T$ is then

$$\begin{bmatrix} a & b & \Delta x \\ c & d & \Delta \theta \\ 0 & 0 & 1 \end{bmatrix}. \quad (1.57)$$

This is especially useful in getting some idea of the allowable tolerances for wedge angle, tilt, and decentration; a lens of focal length f , decentered by a distance d , adds an angular offset $\Delta \theta = d/f$. Similarly, a window of thickness t and index n , whose axis is α degrees off the normal, looks like a free-space propagation of t/n with a spatial offset of $\Delta x = \alpha t(n - 1)$. In the integral representation, offsets are modeled as convolutions with shifted δ -functions; a shift of ξ is a convolution with $\delta(x - \xi)$.

Aside: Complex ABCD Matrices and Diffraction. Siegman[†] shows that a Gaussian amplitude apodization can be modeled using the $ABCD$ matrix for a thin lens with an imaginary focal length. This isn't magic, it's just that a thin lens is a multiplication by an imaginary parabolic exponential, $I(x) = \exp[i\pi x^2/(\lambda f)]$, so a Gaussian of $1/e^2$ radius w , $A(x) = \exp(-x^2/w^2)$, might be said mathematically to be a lens of focal length $i\pi w^2/\lambda$.[‡] Thus by making the first (rightmost) $ABCD$ matrix a Gaussian aperture,

$$\begin{bmatrix} 1 & 0 \\ -i\lambda/(\pi w^2) & 1 \end{bmatrix}, \quad (1.58)$$

you can carry the beam radius right through the $ABCD$ calculation, including converting it to a Helmholtz integral. This makes it simple to find the beam waist, for instance, and if you're building interferometers with very small diameter beams, allows you to calculate the phase front matching at the beam combiner.

[†]A. E. Siegman, *Lasers*. University Science Books, Mill Valley, CA 1986, pp. 786–797.

[‡]Note that we're using field amplitudes and not intensity here.

1.3.12 Source Angular Distribution: Isotropic and Lambertian Sources

A light source (such as the Sun) whose output is independent of direction is said to be *isotropic*; there's no special direction. Light inside an integrating sphere (Section 5.7.7) is also isotropically distributed because there's no special direction. When there's a surface involved, though, things change, on account of obliquity. If you shine a light on a perfectly matte-finished surface, its surface looks equally bright no matter what angle you look from. If you tilt it, it gets foreshortened by the perspective, but if you looked at it from a distance through a drinking straw, you wouldn't be able to tell from its brightness whether it was tilted or not. A surface that passes the drinking-straw test is said to be *Lambertian*.

If you replace your eye with a photodiode, each drinking-straw patch of surface contributes the same amount of photocurrent. As the angle increases, the patches get longer like evening shadows, so $\cos\theta$ times fewer patches will fit on the surface of the source. Another way to put this is that the total projected area of the source goes down like the cosine of the angle of incidence, so the detected photocurrent will be multiplied by the *obliquity factor* $\cos\theta$. Obliquity factors come in just about everywhere—often disguised as $\hat{\mathbf{n}} \cdot \nabla\psi$ —and sometimes they seem mysterious, but all that's really going on is that shadows get longer in the evening.

1.3.13 Solid Angle

Waves expand as they propagate, but as the ray model predicts, a given wave's angular spread is asymptotically constant as $R \rightarrow \infty$. A plane angle is measured between two straight lines, so its measure doesn't depend on how far out you go. A cone has the same property in three dimensions, leading to a natural generalization, *solid angle*. The measure of a plane angle is the arc length cut out by the angle on the unit circle, so we define the solid angle of a cone to be the area it cuts out of the unit sphere. (Note that the cone need not be circular in cross section, or convex in outline, or even be a single glob—it just needs to have a shape that's independent of distance from the vertex.) This area is of course

$$\Omega = \iint_{\text{cone}} \sin\theta \, d\theta \, d\phi, \quad (1.59)$$

where θ is the polar angle (measured from the surface normal) and ϕ is the azimuth (angle in the horizon plane). In optics, we're normally calculating the flux in or out of a surface, so we have to worry about obliquity. Once again, obliquity is nothing deep or difficult to understand—when a beam of light hits a surface at an angle θ off normal, the illuminated patch is stretched out by $\sec\theta$, just as afternoon shadows of vertical objects lengthen as $\tan\theta$. Mathematically, the outward flux through each bit of surface $d\mathbf{A}$ is $\mathbf{P} \cdot d\mathbf{A}$. It simplifies matters if we fold the obliquity into the quoted solid angle, so we usually work with the *projected solid angle* Ω' , where

$$\Omega' = \iint_{\text{cone}} \sin\theta \cos\theta \, d\theta \, d\phi. \quad (1.60)$$

To crispen this idea up, consider a circular angular pattern of half-angle ψ around the surface normal, that is, one that covers the angular disc $\theta < \psi$, $0 \leq \phi < 2\pi$. Its solid angle is $\Omega = 2\pi(1 - \cos\psi) = \pi(\psi^2 - \psi^4/12 + \dots)$ and its projected solid angle is

$\Omega' = \pi \sin^2 \psi = \pi(\psi^2 - \psi^4/3 + \dots)$. Conveniently, if $n = 1$ then $\Omega' = \pi(\text{NA})^2$, which is a useful and easily remembered rule.

A Lambertian surface (one that has no preferred direction) emits into π steradians (the projected solid angle of a hemisphere). Optics folk tend to be loose about the distinction between Ω and Ω' , but it isn't hard to keep straight—if the emitter or receiver is a surface, there's obliquity to worry about, so use Ω' ; if not (e.g., as in gas spectroscopy) use Ω . In the usual low-NA situations, the two are equivalent for practical purposes. There are two cautions to keep in mind: first, be careful if the surface isn't flat—it's the angle between the local surface normal and the light rays that matters. Second, both solid angle and obliquity are far-field concepts, so near a focus we have to use the plane wave decomposition of the field to get the right answer.

1.3.14 Étendue: How Much Light Can I Get?

The first thing that an optical system has to be able to do is transmit light. Apart from solar telescopes, electro-optical systems are limited at least some of the time by how much light they can emit, collect, or detect. Figuring out how much you have and how much more you can get is the aim of radiometry. In Section 1.3.11, we saw that a given light beam can be focused into a smaller area, but only at the price of increasing its numerical aperture. Since $\sin \theta$ cannot exceed unity, a given beam cannot be focused arbitrarily tightly. Some beams can be focused better than others; for example, a beam from an incandescent bulb cannot be focused as tightly as one from a laser. The difference is in their degree of spatial coherence.

The spatial coherence of a beam is a measure of how well its different components (Fourier or real-space) stay in phase with each other. This is revealed by how deep the interference fringes are when different components are made to interfere with one another, as in Young's slit experiment (there's more on this in Section 2.5.4). The theory of imaging with partially coherent light is discussed by Goodman, Born and Wolf, and others and is beyond the scope of this book. As a practical matter, we usually want spatial coherence low enough to eliminate fringes in a full-field (i.e., not scanning) imaging system and high enough not to limit our ability to focus it on our area of interest. The *coherence area* of an optical field gives an idea of how far apart the slits can be and still have interference.

Conversely, one important attribute of an optical system is how well it can cope with low coherence sources. To transmit the most light from such sources, the system needs both a large area and a large angular acceptance. The figure of merit for this attribute is called the *étendue* and is given by

$$E = n^2 A \Omega', \quad (1.61)$$

where A is the clear area and n is the refractive index of the medium in which the projected solid angle Ω' is measured. It's usually just written $A\Omega'$, which assumes that $n = 1$, but we'll carry the n along explicitly. For on-axis circular pupils (the usual case), $E = n^2 A \pi (\text{NA})^2$. This quantity is invariant under magnification, which increases A while decreasing Ω' proportionately, and under refraction. Étendue is a purely geometric property, which explicitly neglects the transmittance of the optical system. This is fine as long as this is reasonably uniform up to the edges of A and Ω' . It is less useful with systems whose transmittance is a strong function of angle, high-index dielectric

interfaces. The *useful* étendue is not preserved on passing through a succession of such elements, so the transmittance must be expressed as a function of position and angle, and carried along mathematically. Étendue is related to the statistical mechanics notion of phase space volume, and the conservation of étendue is the optical analogue of the conservation of phase space volume by adiabatic processes. (The *ABCD* matrices of Section 1.3.11 are all unitary, which is the paraxial version of this.)

With a given low coherence source, any two lossless optical systems with the same étendue will pass the same total optical power, if the source is matched to their characteristics with an appropriate magnification. Any mismatch will reduce the power actually transmitted. A corollary is that the étendue of a system stays the same if you send the light back through the other way. The étendue of an optical system cannot be larger than that of its poorest component, and can easily be worse due to mismatch. This is worth keeping in mind, for example, in the choice of polarizing prisms; types relying on total internal reflection (such as the Glan–Taylor) have much smaller acceptance angles than those relying on double refraction (such as Wollastons), so a bigger prism can have a smaller étendue.

Example 1.8: Étendue and Mismatch. Consider coupling sunlight into a $100\times$, 0.95 NA microscope objective ($f = 2$ mm, FOV diameter = $100\ \mu\text{m}$). If we shine sunlight (9 mrad angular diameter) in the pointy end, we get an effective $n^2A\Omega'$ of $\pi(0.005\ \text{cm})^2[\pi(4.5\ \text{mrad})^2] = 5 \times 10^{-9}\ \text{cm}^2 \cdot \text{sr}$. If we turn it around, the étendue is unaltered, but we get the 6 mm diameter back element instead. The angular acceptance on the exit pupil is a few degrees, so we don't lose anything, and the effective $n^2A\Omega'$ goes up by a factor of 3600 to $1.8 \times 10^{-5}\ \text{cm}^2 \cdot \text{sr}$ —and the source is still mismatched.

1.3.15 What Is “Resolution”?

The classical definitions of Rayleigh and Sparrow specify that two point-like objects of equal brightness are resolved if their images are separated by a defined multiple of the diameters of their diffraction discs. This definition is reasonably adequate for photographic detection, where the photon statistics do not significantly reduce the precision of the measurement.

With modern detectors, it is impossible to specify the resolution of an optical system when signal-to-noise considerations are absent. For example, for a two-point object, one can model the image as the sum of two Airy patterns, whose locations and intensities are parameters. By fitting the model to the observed data, the positions and intensities of the two sources can be extracted. With a high enough signal-to-noise ratio and a sufficiently accurate knowledge of the exact imaging characteristics of our systems, there is no clear limit to the two-point resolution of an optical system, as defined in this way. Optical lithography is another example where the “resolution limit” has repeatedly turned out not to be where it was expected, largely on account of the very high contrast of photoresist and, recently, phase shift masks and computational mask design.

What we really mean by *resolution* is the ability to look at an object and see what is there, in an unambiguous way that does not depend on our choice of model. This model-independent imaging property does degrade roughly in line with Rayleigh and Sparrow, but it is a much more complicated notion than simple two-point resolution. Most of the disagreement surrounding the subject of resolution is rooted here.

1.4 DETECTION

To calculate what an instrument will detect, we need to know how to model the operation of a photodetector. Fortunately, this is relatively simple to do, providing that the detector is reasonably uniform across its sensitive area. From a physical point of view, all detectors convert optical energy into electrical energy, and do it in a *square-law* fashion—the electrical power is proportional to the square of the optical power, with a short time average through the detector’s impulse response. Throughout the rest of this chapter, we will normalize the scalar field function ψ so that the (paraxial) power function $\psi\psi^*$ has units of watts per square meter.

A general square-law detector with an input beam $\psi(\mathbf{x})$ and a responsivity \mathcal{R} will produce an output signal S given by

$$S(t) = \mathcal{R} \iint \langle \psi(\mathbf{x}, t)^* \hat{\mathbf{n}} \cdot \nabla \psi(\mathbf{x}) / k \rangle d^2x \quad (1.62)$$

which for small NA is

$$S(t) = \mathcal{R} \iint \langle |\psi(\mathbf{x}, t)|^2 \rangle d^2x, \quad (1.63)$$

where angle brackets denote time averaging through the temporal response of the detector and the integral is over the active surface of the detector. The gradient $\nabla\psi$ is parallel to the local direction of propagation (see Section 9.2.3) and the dot product supplies the *obliquity factor*, as we saw in Section 1.3.13. If the detector is seriously nonuniform, the responsivity becomes a function of \mathbf{x} , so $\mathcal{R}(\mathbf{x})$ must be put under the integral sign.

The square-law behavior of detectors has many powerful consequences. The first is that all phase information is lost; if we want to see phase variations, we must convert them to amplitude variations before the light is detected. Furthermore, provided that no light is lost in the intervening optical components (watching vignetting especially), the detector can in principle be placed anywhere in the receiving part of the optical system, because the time averaged power will be the same at all positions by conservation of energy. This has great practical importance, because we may need to use a small detector in one situation, to minimize dark current or ambient light sensitivity, and a large one in another, to prevent saturation and attendant nonlinearity due to high peak power levels. The small detector can be put near focus and the large one far away. This freedom applies mathematically as well; provided once again that no additional vignetting occurs, (1.63) can be applied at an image, a pupil, or anywhere convenient.[†] This becomes very useful in interferometers.

As in all interactions of light with matter, the surface properties of the detector and their variation with position, polarization, and angle of incidence are important. Fortunately, detector manufacturers endeavor to make their products as easy to use as possible, so that the worst nonuniformities are eliminated, and in addition, by the time the light gets to the detector, its numerical aperture is usually reduced sufficiently that obliquity factors and dependence on overall polarization are not too serious. As usual, they can be put in by hand if needed, so we’ll continue to use the scalar model and neglect these other effects.

There are a fair number of head-scratchers associated with square-law detection. We’ll talk more about it in Section 3.3.

[†]This is exact and not a Debye approximation.

1.5 COHERENT DETECTION

1.5.1 Interference

An interferometer is nothing more than a device that overlaps two beams on one detector, *coherently*, rather than combining the resulting photocurrents afterwards, *incoherently*. Coherent addition allows optical phase shifts between the beams to give rise to signal changes. In many applications the two beams are different in strength and the weaker one carries the signal information. Consequently, they are often referred to as the signal and local oscillator (LO) beams, by analogy with superheterodyne radios. Coherent detection gives the optical fields the chance to add and subtract before the square law is applied, so that the resulting photocurrent is

$$\begin{aligned} i(t) &= \mathcal{R} \iint_{\text{det}} |\psi_{\text{LO}}(\mathbf{x})e^{-i(\omega_{\text{LO}}t + \phi_{\text{LO}}(\mathbf{x}, t))} + \psi_S(\mathbf{x})e^{-i(\omega_S t + \phi_S(\mathbf{x}, t))}|^2 dA \\ &= i_{\text{LO}} + i_S + i_{\text{AC}}, \end{aligned} \quad (1.64)$$

assuming that the beams are polarized identically (if there are signal beam components in the orthogonal polarization state, they add in intensity, or equivalently in i). The individual terms are

$$i_{\text{LO}} = \mathcal{R} \iint_{\text{det}} d^2x \psi_{\text{LO}} \psi_{\text{LO}}^*, \quad (1.65)$$

$$i_S = \mathcal{R} \iint_{\text{det}} d^2x \psi_S \psi_S^*, \quad (1.66)$$

and

$$\begin{aligned} i_{\text{AC}} &= 2\mathcal{R} \operatorname{Re} \left\{ \iint_{\text{det}} d^2x \psi_{\text{LO}} \psi_S^* \right\} \\ &= 2\mathcal{R} \operatorname{Re} \left\{ \exp(-i \Delta\omega t) \iint_{\text{det}} |\psi_{\text{LO}}(\mathbf{x})| |\psi_S(\mathbf{x})| \exp(i \Delta\phi(\mathbf{x}, t)) dA \right\}. \end{aligned} \quad (1.67)$$

The first two terms, i_{LO} and i_S , are the photocurrents the two beams would generate if each were alone. The remaining portion is the *interference term*. It contains information about the relative phases of the optical beams as a function of position. The interference term can be positive or negative, and if the two beams are at different optical frequencies it will be an AC disturbance at their difference frequency $\Delta\omega$. If the two beams are superposed exactly and have the same shape (i.e., the same relative intensity distributions, focus, and aberrations), ψ_{LO} and ψ_S differ only by a common factor, so the interference term becomes

$$i_{\text{AC}} = 2\sqrt{i_{\text{LO}}i_S} \cos(\Delta\omega t + \phi). \quad (1.68)$$

Aside: Fringe Visibility. Looking at the light intensity on the detector (or on a sheet of paper), we can see a pattern of light and dark fringes if the light is sufficiently coherent.

These fringes are not necessarily nice looking. For laser beams of equal strength, they will go from twice the average intensity to zero; for less coherent sources, the fringes will be a fainter modulation on a constant background. The contrast of the fringes is expressed by their visibility V ,

$$V = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \quad (1.69)$$

which we'll come back to in Section 2.5.4 in the context of coherence theory.

1.5.2 Coherent Detection and Shot Noise: The Rule of One

Application of coherent detection to improve the signal-to-noise ratio is covered in Section 3.11.7. There are three key observations to be made here: coherent detection is extremely selective, preserves phase information, and provides noiseless signal amplification.[†] These three properties give it its power. If the two beams are exactly in phase across the entire detector, the amplitude of the interference term is twice the square root of the product of the two DC terms:

$$i_{AC}(\text{peak}) = 2\sqrt{i_{LO}i_S}. \quad (1.70)$$

If i_S is much weaker than i_{LO} , this effectively represents a large amplification of ψ_S . The amplification is noiseless—the LO shot noise is

$$i_{N\text{shot}} = \sqrt{2ei_{LO}}, \quad (1.71)$$

which is exactly the rms value of i_{AC} when i_S is 1 electron per second (the noise current is down by $\sqrt{2}$ due to the ensemble average over 2π phase). Thus with $\eta = 1$, a signal beam of 1 photon/s is detectable at 1σ confidence in 1 s in a 1 Hz bandwidth, which is a remarkable result—bright-field measurements can be made to the same sensitivity as dark-field measurements.

This leads us to formulate the Shot Noise Rule of One: *One* coherently added photon per *One* second gives an AC measurement with *One* sigma confidence in a *One* hertz bandwidth. (See Sections 1.8.1 and 13.1 for more on AC versus DC measurements.)

Aside: Photons Considered Harmful. Thinking in terms of photons is useful in noise calculations but pernicious almost everywhere else—see Section 3.3.2.

1.5.3 Spatial Selectivity of Coherent Detection

If the phase relationship is not constant across the detector, fringes will form, so the product $E_{LO}E_s^*$ will have positive and negative regions; this will reduce the magnitude of the interference term. As the phase errors increase, the interference term will be reduced more and more, until ultimately it averages out to nearly zero. This means that a coherent detector exhibits gain only for signal beams that are closely matched to the LO beam, giving the effect of a matched spatial filter plus a noiseless amplifier.

[†]A. V. Jelalian, *Laser Radar Systems*. Artech House, Boston, 1992, pp. 33–41.

Another way to look at this effect is to notionally put the detector at the Fourier transform plane, where the two initially uniform beams are transformed into focused spots. A phase error that grows linearly across the beam (equally spaced fringes) corresponds to an angular error, which in the transform plane means that the two focused spots are not concentric. As the phase slope increases, the spots move apart, so that their overlap is greatly reduced. Ultimately, they are entirely separate and the interference term drops to zero. Mathematically these two are equivalent, but physically they generally are not.

If the detector is placed at a focused spot, the local photocurrent density can be so large as to lead to pronounced nonlinearity; this is much less of a problem when the beams fill a substantial fraction of the detector area. On the other hand, two spots that do not overlap will not give rise to any interference term whatsoever, which is not in general true of two broad beams exhibiting lots of interference fringes; even if the broad beams are mathematically orthogonal, small variations in sensitivity across the detector will prevent their interference pattern from averaging to exactly zero.

It is hard to say exactly how serious this effect is in a given case, as it depends strongly on the details of the sensitivity variations. Sharp, strong variations (e.g., vignetting) will give rise to the largest effects, while a smooth center-to-edge variation may do nothing at all noticeable. If the application requires >40 dB (electrical) selectivity between beams at different angles, consider changing focus to separate them laterally, or relying on baffles or spatial filters as well as fringe averaging.

1.6 INTERFEROMETERS

1.6.1 Two-Beam Interferometers

Two-beam interferometers implement the scheme of Section 1.5 in the simplest way: by splitting the beam into two with a partially reflecting mirror, running the two through different paths, and recombining them. Figure 1.10 shows the heavy lifters of the interferometer world, the Michelson and Mach–Zehnder. Mach–Zehnders are more common in technological applications, because the light goes in only one direction in each arm, so it's easier to prevent back-reflections into the laser. On the other hand, a Michelson is the right choice when robust alignment is needed, because one or both mirrors can be replaced by corner cubes (don't tell anyone, but if the cubes are offset from the beam axis,




Figure 1.10. Workhorse two-beam interferometers: (a) Michelson and (b) Mach–Zehnder. The compensator plate in (a) more or less eliminates the effects of dispersion, which will smear out the white-light fringes otherwise, and also reduces the effect of finite aperture.

that's really a skinny Mach–Zehnder). Example 1.12 shows an intensive use of a corner cube type interferometer. Michelsons are a bit easier to align, because autocollimation (sending the beam back on itself) is an easy criterion to use.

An interferometer is intrinsically a four-port device; light is steered between the output ports by interference. If the two beams are perfectly coherent with one another, the output powers P_{O+} and P_{O-} from the two output ports are

$$P_{O\pm} = P_1 + P_2 \pm 2\sqrt{P_1 P_2} \cos \phi, \quad (1.72)$$

where P_1 and P_2 are the split beam powers and ϕ is the phase angle between them. The sign difference comes from beam 1 being reflected and beam 2 transmitted going into port + and vice versa for port –.

1.6.2 Multiple-Beam Interferometers: Fabry–Perots

If instead of splitting the light into multiple paths, we just take two partially reflecting mirrors and put them next to each other, parallel, we get a *Fabry–Perot* (F-P) interferometer. The multiple reflections give Fabry–Perots much greater selectivity for a given size, at the expense of far greater vulnerability to mirror errors and absorption. We'll go through the math in Section 5.4, but the upshot is that a plane-mirror F-P whose mirrors have reflectance R and are spaced d apart in a medium of index n has a total transmission

$$T_{F-P} = \frac{1}{1 + \frac{4R}{(1-R)^2} \sin^2(nk_0 d \cos \theta)}, \quad (1.73)$$

where θ is the angle of incidence of the beam on the mirrors (inside the medium). This obviously consists of spikes at multiples of $\Delta\nu = 1/(2nd)$, the *free spectral range* (FSR). The FWHM of the peaks is FSR/\mathcal{F} , where \mathcal{F} is the *finesse*. Although \mathcal{F} is nominally $(\pi\sqrt{R})/(1-R)$, it is really a measured quantity, because mirror flatness errors are usually the limiting factor. If the mirrors have rms error δ in waves, that limits the finesse to

$$F_{\max} < 1/(2\delta), \quad (1.74)$$

which is a pretty serious limitation most of the time—achieving a finesse of 100 requires mirror accuracy and alignment precision of better than $\lambda/200$. Real F-Ps have a peak T less than 1 (sometimes a lot less), and their total reflectance R_{F-P} is less than $1 - T_{F-P}$. Recently, fabrication precision and coating quality have advanced to the point where finesse values of 2×10^5 or higher can be obtained, at a price.

Inside a F-P, the field is enhanced a great deal; the easiest way to calculate it is to notice that the forward and backward propagating fields inside the cavity are nearly equal, and that the transmitted power has to be T times the forward power. In a perfect F-P, that means that if P_{inc} is coming in, the forward power inside is $T_{F-P}/(1-R) \cdot P_{\text{inc}}$.

1.6.3 Focused-Beam Resonators

Fabry–Perots can have variable or fixed spacing; a fixed F-P is called an *etalon*. Etalons can be tuned over a small range by tipping them, but the finesse drops pretty fast when you do that since the N th reflections start missing each other completely.

The highest finesse F-Ps are not plane-mirror devices, but rather more like laser resonators; as the finesse goes up, even small amounts of diffraction become an increasing difficulty. They need careful matching of the incoming wave to the spherical wave cavity mode, which is a big pain; fortunately, single-mode fiber coupled ones are available—buy that kind if you possibly can. Otherwise, not only do you face critical matching problems, but the minor pointing instability of the laser will turn into noise that cannot easily be removed. Fibers have their problems, but very high finesse focused F-Ps are much worse.

Aside: Confocal Cavities and Instability. It might seem that the ideal optical resonator would be a confocal cavity, where the two mirrors' centers of curvature coincide at the center of the cavity. This is not so, at least not for lasers. Such a cavity is a portion of a single sphere, and there is no special direction in a sphere—any direction is as good as any other, hence a confocal resonator has no stable axis. A tilt of ϵ radians in one mirror produces a shift of the resonator axis of $\delta\theta \approx \epsilon[L/(\Delta L)]$, where L is the distance between the mirror vertices and ΔL is the distance between their foci—which goes to ∞ as $\Delta L \rightarrow 0$. The NA of the resonant mode depends on how far off confocal the cavity is—resonant wavefronts will coincide with the cavity mirror surfaces, so $\text{NA} = 0$ for planar mirrors, $\text{NA} = 1$ for confocal mirrors, and in between, the NA can be backed out from the equation for $R(z)$ in Table 1.1. (Should ΔL be positive or negative?)

1.7 PHOTON BUDGETS AND OPERATING SPECIFICATIONS

1.7.1 Basis

Photons are like money: a certain number are needed for the job at hand, and they're easier to lose than to gain back. Thus the idea of a budget applies to photons as to finances, but it is more complicated in that not all photons are useful—as though we had to budget a mixture of green and purple dollars. A photon budget is an accounting of where photons come from, where they go, and how many are expected to be left by the time they are converted to electrons by the detector. Also like the other kind, people sometimes don't even expect to be on budget; they settle for “this was the best we could do, but I'm not sure what was the problem.” In the author's experience, it is possible to achieve an SNR within 3 dB of budget almost always, and 1 dB most of the time. Don't give up, this theory stuff really works.

On the other hand, the budget must be realistic too. Don't try to measure anything in a bandwidth of less than 10 Hz, unless you have lots of cheap graduate students, and remember that you need a decent signal-to-noise ratio to actually do anything. Sensitivity limits are frequently given as *noise equivalent power* (NEP) or *noise equivalent temperature difference* (NE Δ T or NETD), meaning the amount of signal you need in order to have a signal-to-noise ratio of 1, or equivalently a confidence level of 1σ (68%). Don't let this convince you that an SNR of 1 is useful for anything, because it isn't. Generally for any reasonable measurement you need an SNR of at least 20 dB—even to do a single go/no-go test with a reasonable false call probability, you'll need at least 10 or 15 dB ($3\text{--}5\sigma$). Tolerable images need at least 20 dB SNR; good ones, about 40 dB. Just how large an SNR your measurement has to have is a matter of deepest concern, so it should be one of the first things on the list. Don't rely on a rule of thumb you don't understand fully, including this one. There's lots more on this topic in Section 13.6.

Arriving at a photon budget and a set of operational specifications is an iterative process, as the two are inseparably linked. As the concepts are simple, the subtlety dwells in the actual execution; we will therefore work some real-life examples. The first one will be a shot-noise-limited bright-field system; the second, a background-limited dark-field system, and the third, an astronomical CCD camera. Before we start, you'll need to know how to combine noise sources (see Section 13.6.7) and to think in decibels.

Aside: Decibels. One skill every designer needs is effortless facility with decibels. There are just two things to remember: first, *decibels always measure power ratios, never voltage*. $G(\text{dB}) = 10 \log_{10}(P_2/P_1)$ (we'll drop the subscript from now on and use "ln" for natural log). That formula with a 20 in it is a convenience which applies only when the impedances are identical, for example, a change in the noise level at a single test point. If you don't remember this, you'll start thinking that a step-up transformer has gain. Second, you can do quick mental conversions by remembering that a factor of two is 3 dB (since $\log_{10}(2) \approx 0.3010$), a factor of 10 is 10 dB, and ± 1 dB is $1.25\times$ or $0.8\times$ (since $10^{0.1} \approx 1.259$). For example, if you're looking at the output of an amplifier, and the noise level changes from 10 mV to 77 mV, that's about an 18 dB change: you get 80 from 10 by multiplying by 10 and dividing by 1.25 (20 dB – 2 dB, remembering that dB measure power), or by multiplying by 2 three times ($6 + 6 + 6 = 18$).

Example 1.9: Photon Budget for a Dual-Beam Absorption Measurement. One way of compensating for variations in laser output is to use two beams, sending one through the sample to a detector and the other one directly to a second detector for comparison. A tunable laser (e.g., Ti:sapphire or diode) provides the light. The desired output is the ratio of the instantaneous intensities of the two beams, uncontaminated by laser noise, drift, and artifacts due to etalon fringes or atmospheric absorption. For small absorptions, the beams can be adjusted to the same strength, and the ratio approximated by their difference divided by a calibration value of average intensity,

$$\epsilon = \frac{n_{\text{sig}}(\lambda, t)}{n_{\text{comp}}(\lambda, t)} - 1 \approx \frac{n_{\text{sig}}(\lambda, t) - n_{\text{comp}}(\lambda, t)}{\langle n_{\text{comp}}(\lambda) \rangle}, \quad (1.75)$$

where n_{sig} and n_{comp} are the photon flux (s^{-1}) in the beams. The total electrical power in the signal part is

$$P_{\text{sig}} = [\eta e(n_{\text{sig}} - n_{\text{comp}})]^2 R_L. \quad (1.76)$$

In the absence of other noise sources, shot noise will set the noise floor:

$$i_{N\text{shot}} = e\sqrt{2\eta(n_{\text{comp}} + n_{\text{sig}})}. \quad (1.77)$$

For 1 mW per beam at 800 nm and $\eta = 1$, this amounts to a dynamic range (largest electrical signal power/noise electrical power) of 150 dB in 1 Hz, or a 1σ absorption of 3 parts in 10^8 . Real spectrometers based on simple subtraction are not this good, due primarily to laser noise and etalon fringes. Laser noise comes in two flavors, intensity and frequency; it's treated in Section 2.13.

Laser noise cancelers use subtraction to eliminate intensity noise and actually reach this shot noise measurement limit (see Sections 10.8.6 and 18.6.3). When frequency noise is a problem (e.g., in high resolution spectroscopy) we have to stabilize the laser.

Frequency noise also couples with the variations in the instrument's T versus λ to produce differential intensity noise, which in general cannot be canceled well. If the instrument's optical transmittance is T and the laser has a (one-sided) FM power spectrum $S(f_m)$ (which we assume is confined to frequencies small compared to the scale of T 's variation), FM-AM conversion will contribute rms current noise i_{Nfa} :

$$i_{Nfa}(f) = n\eta e S(f) \frac{dT}{d\nu}. \quad (1.78)$$

If this noise source is made small, and the absorption can be externally modulated, for example, by making the sample a molecular beam and chopping it mechanically, the shot noise sensitivity limit can be reached fairly routinely. Note that this is not the same as a measurement accuracy of this order; any variations in T or drifts in calibration will result in a multiplicative error, which, while it goes to zero at zero signal, usually dominates when signals are strong.

Example 1.10: Photon Budget for a Dark-Field Light Scattering System. Many systems (e.g., laser Doppler anemometers) detect small amounts of light scattered from objects near the waist of a beam. Consider a Gaussian beam with $P = 0.5$ mW and 0.002 NA at 633 nm. From Table 1.1, the 3 dB beam radius at the waist is $70 \mu\text{m}$, and the central intensity is $2P\pi(\text{NA})^2/\lambda^2 = 3.1 \times 10^4$ W/m², which is 1.0×10^{23} photons/m²/s. A sufficiently small particle behaves as a dipole scatterer, so that the scattered light intensity is zero along the polarization axis, and goes as the sine squared of the polar angle θ . It will thus scatter light into 2π steradians, so if the total scattering cross section of the particle is σ_{tot} , the averaged scattered flux through a solid angle Ω (placed near the maximum) will be $F = 1.0 \times 10^{23} \sigma_{\text{tot}} \Omega / (2\pi)$. (Larger particles exhibit significant angular structure in their scattering behavior, with a very pronounced lobe near the forward direction.) A 1 cm diameter circular detector at a distance of 5 cm from the beam waist will subtend a projected solid angle of $\Omega' = \pi(\text{NA})^2 = \pi(0.5/5)^2 \approx 0.03$. A particle crossing the beam waist will thus result in $N = 5 \times 10^{21} \sigma_{\text{tot}}$ photons/s.

If the detector is an AR-coated photodiode ($\eta \approx 0.9$) with a load resistor $R_L = 10$ M Ω , then in a time t (bandwidth $1/(2t)$ Hz), the 1 Hz rms Johnson noise current $(4kTB/R)^{1/2}$ is $(2kT/10^7)^{1/2}$, which is 0.029 pA or 1.8×10^5 electrons/s (which isn't too good).

From Section 13.6.15, we know that in order to achieve a false count rate R of 1 in every 10^6 measurement times, we must set our threshold at approximately 5.1 times the rms amplitude of the additive Gaussian noise (assuming that no other noise source is present). Thus a particle should be detectable in a time t if σ_{tot} exceeds

$$\sigma_{\text{min}} = \frac{1.8 \times 10^5 (5.1)}{5 \times 10^{21} \eta \sqrt{t}} = \frac{1.8 \times 10^{-16}}{\eta \sqrt{t}}, \quad (1.79)$$

where η is the quantum efficiency. A particle moving at 1 m/s will cross the $140 \mu\text{m}$ 3-dB diameter of this beam in $140 \mu\text{s}$, so to be detectable it will need a scattering cross section of at least 1.6×10^{-14} m², corresponding to a polystyrene latex (PSL) sphere ($n = 1.51$) of about $0.2 \mu\text{m}$ diameter.

Now let's sanity-check our assumptions. For the circuit to respond this fast, the time constant $C_{\text{det}} R_L$ must be shorter than $100 \mu\text{s}$ or so, which limits the detector capacitance to 10 pF, an unattainable value for such a large detector. Even fast detectors are generally limited to capacitances of 50 pF/cm², so assuming the detector capacitance is actually

100 pF, we cannot measure pulses faster than 1 to 2 ms with our system as assumed. There are circuit tricks that will help considerably (by as much as $600\times$, see Section 18.4.4), but for now we'll work within this limit. If we can accept this speed limitation, and the accompanying $\sim 10\times$ decrease in volumetric sampling rate, we can trade it in for increased sensitivity; a particle whose transit time is 2 ms can be detected at $\sigma_{\min} = 4 \times 10^{-15} \text{ m}^2$. If the particle is much slower than this, its signal will start to get down into the $1/f$ noise region, and non-Gaussian noise such as popcorn bursts will start to be important. Already it is in danger from room lights, hum, and so on.

If the detector is a photon-counting photomultiplier tube (PMT, $\eta \approx 0.2$), the noise is dominated by the counting statistics, and the technical speed limitation is removed (see Section 3.6.1). However, PMTs have a certain rate of spurious *dark counts*. Dark counts generally obey Poisson statistics, so if we assume a mean rate $N_{\text{dark}} = 200 \text{ Hz}$, then in a $140 \mu\text{s}$ measurement, the probability of at least one dark count is $200 \times 140 \mu\text{s} \approx 0.028$. We are clearly in a different operating regime here, where the fluctuations in the dark count are not well described by additive Gaussian noise as in the previous case. From Section 13.6.16, the probability of a Poisson process of mean rate λ per second producing exactly M counts in t seconds is

$$P(M) = \frac{(\lambda t)^M e^{-\lambda t}}{M!}. \quad (1.80)$$

If we get 200 photons per second, then the probability that a second photon will arrive within $140 \mu\text{s}$ of the first one is $(0.028)e^{-0.028} \approx 0.027$, so we expect it to happen $200 \times 0.027 \approx 5.5$ times a second. Two or more additional photons will arrive within $140 \mu\text{s}$ of the first one about 0.076 times per second, and three or more additional photons only 0.0007 times per second, so if we require at least four counts for a valid event, the false count rate will be one every 20 minutes or so, which is usually acceptable. The limiting value σ_{\min} is then

$$\sigma_{\min} \approx \frac{4}{5 \times 10^{21} \eta t} \approx \frac{8 \times 10^{-22}}{\eta t}, \quad (1.81)$$

which is about $3 \times 10^{-17} \text{ m}^2$, nearly three orders of magnitude better than the photodiode, and sufficient to detect a PSL sphere of $0.08 \mu\text{m}$ (alas for particle counters, the signal goes as a^6). We can use the same Poisson statistics to predict the probability of detection, that is, $P(\geq 4 \text{ photons})$ as a function of σ_{tot} .

Besides detector capacitance, such measurements are often limited by the shot noise of light from the background, from imperfect beam dumps, or from molecular Rayleigh scatter (as in the blue sky), so that our photon budget is not a sufficient theoretical basis for the planned measurement. More detail is available in Chapter 10. In addition, we have here assumed a very simple deterministic model for signal detection in noise; any event whose nominal detected signal is greater than our threshold is assumed to be detected. This assumption is unreliable for signals near the threshold, and is dealt with in a bit more detail in Section 13.6.15. Finally, we assumed that our noise was limited by the use of a boxcar averaging function of width t , equal to the 3 dB width of the pulse. Even with *a priori* knowledge of the arrival time of the particle, this is not the optimal case; if the particles are moving faster or slower than we anticipate, a fixed averaging window may be far from optimal. This is the topic of Section 13.8.10.

Example 1.11: Photon Budget for an Astronomical CCD Camera. An astronomical telescope gathers photons from sources whose brightness cannot be controlled. It also gathers photons from Earth, due to aurora, meteors, cosmic ray bursts, scattered moonlight, and human artifacts such as street lights. Extraneous celestial sources such as the zodiacal light and the Milky Way are also important. It represents an interesting signal detection problem: the signals that it tries to pull out of random noise are themselves random noise. The inherent noisiness of the signal is of less consequence in the optical and infrared regions than in the radio region. The brightness of sky objects is quoted in logarithmic relative *magnitudes*, with respect to standard spectral filters. A look in Allen’s *Astrophysical Quantities* (affectionately known as “AQ”) reveals that in the “visible” (*V*) filter band, centered at 540 nm in the green, a very bright star such as Arcturus or α Centauri has a magnitude $m_V = 0$, and that such an object produces a total flux at the top of the atmosphere of about 3.8 nW/m^2 , of which about 80% makes it to the Earth’s surface. A first magnitude star ($m_V = 1.0$) is 100 times brighter than a sixth magnitude star, which is about the limit of naked-eye detection in a dark location. A one-magnitude interval thus corresponds to a factor of $100^{0.2} \approx 2.512$ in brightness.[†]

Even in a dark location, the night sky is not perfectly dark; its surface brightness in the *V* band is about $400 \text{ nW/m}^2/\text{sr}$, which at 2.3 eV/photon is about 1×10^{10} photons/s/m²/sr. An extended object such as a galaxy will be imaged over several resolution elements of the detector, whereas a point-like object such as a star will ideally be imaged onto a single detector element. Without adaptive optics, the turbulence of the atmosphere during the (long) measurement time limits the resolution to the size of a *seeing disc* of diameter $0.25''$ (arc seconds) on the best nights, at the best locations, with $1''$ being more typical, and $3''$ – $5''$ being not uncommon in backyard observatories. Thus with enough pixels, even a stellar object produces an extended image, and the SNR will be limited by the fluctuations in the sky brightness in a single pixel. Each pixel subtends the same solid angle Ω on the sky, so the mean photoelectron generation rate per pixel is

$$n_{\text{tot}} = \eta \Omega A (J_{\text{sky}} + J_{\text{star}}), \quad (1.82)$$

where A is the area of the telescope objective, η is the quantum efficiency, n is in electrons/s, and J is the photon angular flux density in photons/s/m²/sr.

There are two classes of noise source in the measurement: counting statistics of the sky photoelectrons and of the dark current, which go as $(nt)^{1/2}$, and the (fixed) readout noise q_{RO} . The electrical SNR thus is

$$\text{SNR} = \frac{(J_{\text{star}} t \eta \Omega A)^2}{\Delta n_{\text{RO}}^2 + n_{\text{dark}} t + t \eta \Omega A (J_{\text{star}} + J_{\text{sky}})}. \quad (1.83)$$

CCD pixels have a fixed *full well* capacity B , ranging from about 5×10^4 electrons for the ones used in a camcorder to 10^6 for a scientific device. Thus the maximum exposure time for a perfect device is equal to eB/i_{sky} , which is on the order of 1 week, so that the

[†]The magnitude scale goes back to the ancient Greeks—the numerical factor is weird because the log scale was bolted on afterwards and tweaked to preserve the classical magnitudes while being reasonably memorable.

well capacity is not a limitation for dim objects. The noise-equivalent photon flux density (NE Δ J) is the number of photons per second required to achieve an SNR of 1 (0 dB) in the measurement time. Setting the SNR in (1.83) to 1 yields a quadratic equation for J_{star} , whose solution is

$$\text{NE}\Delta J_{\text{star}} \approx \frac{1}{\eta\Omega A} \sqrt{\frac{n_{\text{dark}} + \eta\Omega A J_{\text{sky}}}{t} + \frac{\Delta q_{\text{RO}}^2}{t^2}}. \quad (1.84)$$

A 100 cm telescope with a 40% central obstruction (due to the secondary mirror) has an area $A = 0.66 \text{ m}^2$, and an angular resolution of about 0.14 arc seconds. A commercially available V filter has a transmittance of about 0.72 at the nominal peak.[†] Assuming that the telescope itself has an efficiency of 0.8, due primarily to absorption in uncoated aluminum primary and secondary mirrors, and is used with a cooled, back-surface illuminated CCD with $\eta = 0.85$ and $B = 1.8 \times 10^5$, the end-to-end quantum efficiency of the telescope system is $\eta \approx 0.5$. A good cooled CCD has a dark current of around 1 electron per second per pixel, and rms readout noise of about 5 electrons. With this apparatus,

$$\text{NE}\Delta J \approx \frac{3}{\Omega} \sqrt{\frac{1 + 3.3 \times 10^{11} \Omega}{t} + \frac{5^2}{t^2}}. \quad (1.85)$$

This is dominated by readout noise at short exposures. At longer times, dark current fluctuations dominate for very small values of Ω but sky background fluctuations dominate for larger Ω . The largest SNR improvement comes from increasing integration time, as shown in Figure 1.11a–d which are pictures of the galaxy M100 taken by Dennis di Cicco using an 11 inch Schmidt–Cassegrain telescope at $f/6.2$ and an SBIG ST-6 CCD camera.[‡]

With an exposure time of an hour, and a 3×3 arc second pixel size (favored by many amateurs, whose CCD budgets don't run to as many pixels as the professionals'), J_{min} is 2.2×10^9 , which corresponds to a surface brightness of -28 magnitudes per square arc second, which is how this is usually quoted in astronomical circles.

If we were merely users of this instrument, that would be that. However, we are the designers, so we have complete flexibility to trade off the operating parameters. What should Ω be? For detection of faint objects, we set Ω as large as is consistent with adequate image resolution, so that the photoelectrons dominate the dark current. For stellar objects, such as star clusters, we want Ω small, since if all the light from each star is going on one pixel anyway, shrinking Ω reduces the sky light while keeping the dark current and signal current the same. For a pixel size of 0.22 arc seconds, J_{min} is 1.1×10^{11} , which means that with a 1 hour exposure on a really great seeing night, we could detect a star of magnitude 25.6 with 3σ confidence. (We've left out the noise contributed by the calibration process, which may be significant if there aren't enough calibration frames—see Section 3.9.19.)

[†]Optec Corp, *Model PFE-1 Technical Manual*.

[‡]Michael V. Newberry, *The signal to noise connection. CCD Astronomy*, Summer 1994. (Copyright © 1994 by Sky Publishing). Reprinted with permission.




Figure 1.11. CCD image of M100: (a) 1 s integration, (b) 10 s integration, (c) 100 s integration, and (d) 1000 s integration.

1.8 SIGNAL PROCESSING STRATEGY

Once a photocurrent leaves the detector, the work of the optical system is done and the maximum SNR has been fixed by the photon budget. The signal processing job is to avoid screwing this up by needlessly rejecting or distorting the signal, on one hand, or by admitting noise unnecessarily, on the other. Signal processing operates under physical constraints, having to do with maximum allowable time delays, and practical ones such as complexity and cost. Most of the time, the inherently highly parallel data coming in as light is converted to a serial electrical channel, so that electrical wave filtering is appropriate.

1.8.1 Analog Signal Processing

The first bit of signal processing is the detector front end amplifier, in which should be included any summing or subtracting of photocurrents, for example, in a differential position measurement using a split detector. As we saw in Example 1.10, a bad front end can hobble the system by adding noise and having too narrow a bandwidth with a

given detector. Most designers are uncomfortable designing front ends, and many wind up crippling themselves by buying a packaged solution that isn't appropriate to the problem; there's detailed help available in Chapters 3 and 18.

In a baseband (near DC) measurement, the next part is usually filtering and then digitizing (AC signals will usually need frequency conversion and detection too). The filter should roll off by at least $6N$ dB (where N is the number of bits) at the Nyquist limit of the digitizer (see Section 17.4.4) and should reject strong spurious signals (e.g., 120 Hz from room lights) that would otherwise overwhelm subsequent stages or require more bits in the digitizer. The bandwidth of the filter should be the same as that of the signal. Although some of the signal is lost, maximum SNR occurs when the filter's frequency response is identical with the signal's power spectrum. Filters wider than that allow more noise in than necessary, and narrower ones clip off too much signal. Some filters have much better time responses than others. Have a look at the Bessel and equiripple group delay filters of Section 15.8.4. To get the best SNR with pulses, use a matched filter (Section 13.8.10).

We need to be able to convert time-domain to frequency-domain specifications. Remember that a 1 s averaging window corresponds to a bandwidth of 0.5 Hz at DC. The reason is that in the analytic signal picture negative frequencies get folded over into positive ones. However, that same 1 s window and the same folding, applied to an AC measurement, gives 0.5 Hz both above and below the (positive frequency) carrier, a total of 1 Hz. The result is that a baseband (near-DC) measurement has half the noise bandwidth[†] of an AC measurement with the same temporal response. The resulting factor of 3 dB may be confusing.

The digitizing step requires care in making sure that the dynamic range is adequate; an attractively priced 8 bit digitizer may dominate the noise budget of the whole instrument. For signals at least a few ADUs[‡] in size, an ideal digitizer contributes additive *quantization noise* of $1/\sqrt{12}$ ADU to the signal, but real A/Ds may have as much as several ADUs of technical noise and artifacts (up to ~ 100 ADUs for $\Delta\Sigma$ ADCs, see Section 14.8.3), so check the data sheet. Converter performance will degrade by 1–3 bits' worth between DC and the Nyquist frequency. Bits lost here cannot be recovered by postprocessing, so be careful to include a realistic model of digitizer performance in your conceptual design.[§]

1.8.2 Postprocessing Strategy

With the current abundance of computing power, most measurements will include a fair bit of digital postprocessing. This may take the form of simple digital filtering to optimize the SNR and impulse response of the system, or may be much more complicated, such as digital phase-locked loops or maximum likelihood estimators of signal properties in the presence of statistically nonstationary noise and interference. In general, the difference between a simplistic postprocessing strategy and an optimal one is several decibels; this

[†]Noise bandwidth is the equivalent width of the power spectrum of the filter (see Section 13.2.5). If we put white noise into our filter, the noise bandwidth is the output power divided by the input noise power spectral density (in W/Hz), corrected for the passband insertion loss of the filter.

[‡]An ADU (analog-to-digital converter unit) is the amount of signal required to cause a change of 1 in the least significant bit of the converter.

[§]High-end oscilloscopes are a partial exception—they overcome timing skew and slew-dependent nonlinearity by calibrating the daylights out of themselves. They cost \$100k and are only good to 6 bits.

may seem trivial, but remember that a 3 dB improvement in your detection strategy can sometimes save you half your laser power or $2/3$ of the weight of your collection optics. (As an aside, calling it postprocessing doesn't mean that it can't be happening in real time.)

1.8.3 Putting It All Together

Let's catch our breath for a moment. We began this chapter with the aim of learning how to do a complete feasibility calculation for an optical instrument, which we should now be fully equipped to do. We have covered a lot of ground in a short time, so don't be discouraged if it takes a while for it all to slot together. It becomes easier with practice, and standing around with one or two smart colleagues doing this stuff on a white board is the most fun of any technical activity. To sum up, a conceptual design goes like this:

1. *Write down what you know.* Get a handle on the basic physics, figure out which is the largest effect operating, and estimate how much smaller the next biggest one is.
2. *Think up a measurement principle.* This may be routine, or may take a lot of imagination. Use it to take a rough cut at your photon budget. From that, estimate, for example, the laser power required, the size of the detection area needed, and so on.
3. *Simplify the analysis.* Use limiting cases, but estimate where they will break down.
4. *Make a very rough optical design.* Choose the NAs, wavelength, working distances, and so on. Check that it makes physical and economic sense.
5. *Guess a detection and signal processing strategy.* One might choose a large germanium photodiode operating at zero bias, followed by a somewhat noisy front end amplifier, a Butterworth lowpass filter, and a fixed threshold. Watch out for sources of systematic error and drift (e.g., etalon fringes or spectral changes in the source).
6. *Make a detailed photon budget.* See if your scheme will do the job, without unrealistically difficult steps.[†] If so, you have a measurement. If not, try step 5 again. If no amount of background reduction, low noise circuitry, or signal processing will help, go back to step 2 and think harder.
7. *Check it for blunders.* Do it all over again a different way, and using scaling arguments to make sure the power laws are all correct. If you go ahead with this instrument, a lot will be riding on the correctness of your calculation—don't scrimp on time and effort here. If you have one or two colleagues who are difficult to convince, try it out on them, to see if they can poke any holes in your logic.

Remember that what will get you is not misplacing a factor of Boltzmann's constant—that's large and obvious—but rather the factors of order 1 that you haven't thought out carefully. These are things like using the peak value when the rms is what's relevant, or forgetting that when you multiply two similar peaks together the result is $\sqrt{2}$ times narrower, or assuming that the bandwidth of a 1 s averaging window is 1 Hz (see

[†]As in the light bulb spectrometer of Example 17.9.




Figure 1.12. The ISICL sensor is an alignment-insensitive scanning interferometer for finding submicron particles in hostile places such as plasma chambers.

Section 13.2.5). A few of these taken together can put you off by factors of 10 or more—very embarrassing.

Estimating and controlling systematic errors is one of the high arts of the designer, since they don't obey any nice theorems as random errors sometimes do. For now we'll just try to keep the optical system simple and the processing modest and linear. Getting too far away from the measurement data is a ticket to perdition.

Let's do an extended example that illustrates many of the concepts of this chapter.

Example 1.12: Conceptual Design of a Shot-Noise-Limited, Scanning Interferometer.

The ISICL (in situ coherent lidar) system[†] detects submicron ($>0.25 \mu\text{m}$) contaminant particles in plasma chambers and other semiconductor processing equipment. As shown in Figure 1.12, it works by scanning a weakly focused laser beam around inside the chamber through a single window, and detecting the light backscattered from particles.

Backscatter operation is very difficult. Between the strong plasma glow and the stray laser light bouncing off the back wall of the chamber, it must deal with very bright stray light—thousands of times worse than that seen by an ordinary dark-field detector, and around a million times brighter than a nominally detectable particle. Coherent detection is probably the only feasible measurement strategy. With a coherent detector, a particle crossing near the focus of the beam gives rise to a short ($3 \mu\text{s}$) tone burst, whose duration is the transit time of the beam across the particle and whose carrier frequency is the Doppler shift from its radial motion. These tone bursts are filtered and amplified, then compared with a threshold to determine whether a particle event has occurred. Since

[†]P. C. D. Hobbs, ISICL: in situ coherent lidar for particle detection in semiconductor processing equipment. *Appl. Opt.* **34**(9), 1579–1590 (March 1995). Available at <http://electrooptical.net/www/isicl/isiclAO.pdf>.

the phase of the tone burst is random, the highest peak may be positive or negative, so bipolar thresholds are used.

The sensitivity of an instrument is expressed by the minimum number of photons it requires in order to detect a particle, which is a function of the confidence level required. A false alarm rate of 10^{-11} in the measurement time corresponds to about 1 false count per day with bipolar thresholds in a 1 MHz bandwidth. From Section 13.6.15, the false count rate depends only on the measurement bandwidth and the ratio α of the threshold to the rms noise voltage. A false count rate of 10^{-11} per inverse bandwidth requires $\alpha = 7.1$.

In an AC measurement, the shot noise level is equivalent to a single coherently detected noise photon in the measurement time, so it might seem that 7 scattered photons would be enough. Coherent detectors detect the amplitude rather than the intensity of the received light, however, so to get a signal peak 7.1 times the rms noise current, we actually need $7.1^2 \approx 50$ photons per burst.

We have a first estimate of how many photons we need, so let's look at how many we expect to get. For a particle in the center of the sensitive region, the received power is the product of the transmit beam flux density times the differential scattering cross section $\partial\sigma/\partial\Omega$ of the particle, times the detector projected solid angle Ω_d . Working in units of photons is convenient, because the relationship between photon counts and shot noise means that electrical SNR is numerically equal to the received photon count per measurement time. This works independently of the signal and filter bandwidth. Initially we ignore the losses imposed by the matched filter.

For a Gaussian transmit beam of power P at wavelength λ , focused at a numerical aperture NA, Table 1.1 gives the photon flux density at the beam waist as

$$J(P, \lambda, \text{NA}) = \frac{2\pi(\text{NA})^2 P \lambda}{hc}. \quad (1.86)$$

Assuming the scattered field is constant over the detector aperture, Table 1.1 predicts that the effective detector solid angle is

$$\Omega_d = \pi(\text{NA})^2 \quad (1.87)$$

and so the expected number of photons available per second is

$$\langle n_0 \rangle = \frac{2\pi^2(\text{NA})^4 P \lambda}{hc} \frac{\partial\sigma}{\partial\Omega}. \quad (1.88)$$

Not all of these will be detected, due to imperfect efficiency of the optics and the detector. The quantum efficiency η of the detector is the average number of detection events per incident photon on the detector; it's always between 0 and 1. (Photodetectors generally give one photoelectron per detection event, so $\eta\langle n_0 \rangle$ is the number of photoelectrons before any amplification.) A good quality, antireflection coated silicon photodiode can have $\eta \approx 0.95$ in the red and near IR, but there are in addition front surface reflections from windows and other optical losses. A receiver whose end-to-end efficiency is over 0.5 is respectable, and anything over 0.8 is very good. A value of 0.9 can be achieved in systems without too many elements, but not cheaply. (We

should also multiply by the square of the Strehl ratio to account for aberrations; see Example 9.6.)

The SNR can be traded off against measurement speed, by the choice of scanning parameters and filter bandwidths. Narrowing the bandwidth improves the time-averaged SNR but smears pulses out. In a pulsed measurement, the optimal filter is the one that maximizes the received SNR at the peak of the pulse. For pulses in additive white Gaussian noise, the optimum filter transfer function is the complex conjugate of the received pulse spectrum (this filter is not necessarily the best in real life, but it's an excellent place to start—see Section 13.8.10). Such a matched filter imposes a 3 dB signal loss on a Gaussian pulse.[†] However, the measurement detects threshold crossings, and for the weakest detectable signals, the peaks just cross the threshold. Thus this 3 dB is made up by the factor of $\sqrt{2}$ voltage gain from the peak-to-rms ratio, so that the minimum detectable number of photons (in the deterministic approximation) is still α^2 .

We can estimate the stray light in the following way. Assume a detector requiring 50 photons for reliable particle identification, operating in backscatter with a 100 mW laser at 820 nm and $NA = 0.008$. Assume further (optimistically) that the back wall of the chamber is a perfectly diffuse (Lambertian) reflector, so that the light is scattered into π sr. The incident beam has 4×10^{17} photons/s, so that the backscattered stray light has average brightness 1.3×10^{17} photons/s/sr; the detector solid angle is $\pi(NA)^2 \approx 0.0002$ sr, so the total detected signal due to stray light is about 2.6×10^{13} photons per second. Assuming that a particle takes $3 \mu\text{s}$ to yield 50 photons (1.7×10^7 photons/s), the stray light is 10^6 times more intense than the signal from a nominally detectable particle. What is worse, the signal from the back wall exhibits speckles, which move rapidly as the beam is scanned, giving rise to large (order-unity) fluctuations about the average stray light intensity. The size of the speckles (and hence the bandwidth of the speckle noise) depends on the distance from the focus to the chamber wall and on the surface finish of the wall. Peak background signals are generally much larger than this average.

The Doppler frequency shift in the detected signal due to a particle traveling with velocity \mathbf{v} encountering incident light with wave vector \mathbf{k}_i and scattering it into a wave with wave vector \mathbf{k}_s is

$$f_d = \mathbf{v} \cdot (\mathbf{k}_s - \mathbf{k}_i) / (2\pi). \quad (1.89)$$

For a system operating in backscatter, $\mathbf{k}_s - \mathbf{k}_i \approx -2\mathbf{k}_i$. At 830 nm, a particle moving axially at 50 cm/s will give rise to a tone burst whose carrier frequency is about 1.22 MHz. This is the nominal maximum particle axial velocity for ISICL.

The remaining engineering problems center on the choice of detection bandwidths and the accurate estimation of the shot noise level in the presence of large signals due to particles. In the present case, where the Doppler shift may be large compared to the transit time bandwidth, the peak frequency of the received spectrum of a given pulse is not known in advance. The optimal signal processing system depends on the range of particle velocities expected. In quiet plasma chambers, where most particles orbit slowly within well-defined traps, the maximum expected velocity may be as low as 5 cm/s,

[†]M. Skolnik, ed., *Radar Handbook*, 2nd ed. McGraw-Hill, New York, 1990, pp. 3.21–3.23.

whereas in an environment such as a rapidly stirred fluid tank or the roughing line of a vacuum pump, the velocity range may be much greater. The scan speed of the beam focus through the inspected volume is much higher (about 20 m/s).

With carrier frequencies ranging from 0 to 1.22 MHz, the Doppler bandwidth is much larger than the transit time bandwidth (150 kHz for a 3 μ s burst FWHM), so that it is inefficient to perform the thresholding operation in a single band. In the present system, four bands are used to cover the Doppler bandwidth. This is frequently best in low NA systems like ISICL, where the focus is many wavelengths across.

In a thresholding operation, it is essential to set a high enough threshold that the sensor does not report erroneously high particle counts, possibly resulting in needless down time for the processing tool being monitored. At the same time, it is economically important to use the available laser power as efficiently as possible; at the time, the laser used in this sensor cost over \$1200, so that (loosely speaking) too-high thresholds cost \$400 per decibel. The signal processing strategy is to set separate bipolar thresholds for each band, using an automatic thresholding circuit. This circuit exploits the accurately known noise amplitude statistics to servo on the false counts themselves and ignore signal pulses, however large they may be. In radar applications, this is known as a CFAR (constant false alarm rate) servo; the technologies employed are quite different, however, since a radar system can look at a given target many times, and its noise is very non-Gaussian. The ISICL false alarm rate (FAR) tracker can accurately servo the FAR at a level much below the true count rate in most applications.

Figure 1.13 shows a typical tone burst from a 0.8 μ m diameter PSL sphere, together with cursors showing the predicted peak-to-peak voltage for a particle passing through the center of the sensing volume. For a particle exactly in focus, the photon flux predicted




Figure 1.13. Measured tone burst caused by a 0.8 μ m polystyrene latex (PSL) sphere crossing the beam near focus. The horizontal cursors are the predicted peak-to-peak value of the output. The bandwidth of the signal shown is several times wider than the actual measurement bandwidth, which prevents distortion of the tone burst but increases the noise background.

by (1.88) is converted to a signal current i_{AC} by (1.68), and to a voltage by multiplying by the known current to voltage gain (*transimpedance*) of the front end and any subsequent amplifiers. Because of the known relationship between the signal size and the shot noise in a coherent detector, we can check the ratio of the rms noise to the signal easily as well (the aberration contribution is calculated in Example 9.6). The measured system parameters are $NA = 0.0045$, $P = 90$ mW, $\lambda = 820$ nm, $\eta = 0.64$. Taking into account the addition of the noise and signal, the error is less than 20% (1.6 dB electrical), indicating that the theory correctly predicts the signal size and SNR.

Sources and Illuminators

And God saw the light, that *it* was good: and God divided the light from the darkness.
—Genesis 1:4 (Authorized Version)

2.1 INTRODUCTION

In order to make an optical measurement, we need a light source. In some systems, the only source required is ambient illumination or the luminosity of the object itself. More often, though, the instrument must supply the light. Instrument designers are less likely to neglect the light source than the detector, but still it is often chosen without proper consideration of the pain and suffering its deficiencies may cause one, or without regard to the special attributes of an alternative.

This chapter deals with light sources and illumination systems, dwelling on their strengths and weaknesses in applications, rather than on the physics of how they work. We stick with the mainstream choices: gas, solid state, and diode lasers, tungsten bulbs, arc lamps, and LEDs. There are other interesting sources, useful in special cases, many of which are discussed in the OSA Handbook.

2.2 THE SPECTRUM

Dye laser catalogs usually have a nice chart showing the electromagnetic spectrum from the infrared (IR), through the surprisingly narrow visible range, to the ultraviolet (UV), short UV, and vacuum UV (VUV). Electrodynamics doesn't alter its character between regions, but the interaction of light and matter does change systematically with wavelength.

2.2.1 Visible Light

The visible part of the spectrum is conventionally taken to be the octave from 400 to 800 nm. Any measurement that can be done with visible light should be.[†] A visible light

[†]There are rare exceptions, such as IR laser based intraocular measurements, where the use of visible light would be very uncomfortable, or solar blind UV photomultipliers for seeing dim UV sources in normal daylight.




Figure 2.1. Relative response function of human cone cells versus wavelength. Note the almost complete overlap of the red and green pigments, leading to good wavelength discrimination in the yellow and less elsewhere. Note: The blue channel has been normalized to the same height as the others, but it's much smaller in reality.

system is enormously easier to align, use, and find parts for than any other kind. If lasers are involved, it's also somewhat safer for your eyes, because of the blink reflex and because you can see the stray beams and block them all before they leave your setup.

The wavelength resolution of human color perception is wildly variable, due to the pronounced overlap of the sensitivities of the retinal pigments, as shown in Figure 2.1.[†]

2.2.2 Ultraviolet

The UV begins at about 400 nm and extends to the border of X-rays at a few nanometers. Near UV (320–400 nm) is generally safe and convenient, because lenses still work. Deep UV (180–320 nm) is less safe, and lens materials are generally scarce there and may be darkened by exposure. The most miserable is vacuum UV (below about 180 nm), where O_2 dissociates to produce ozone, and room air and everything else becomes strongly absorbing. Beware of insidious skin and corneal burns from bright deep-UV sources.

Ultraviolet photons are energetic enough to cause ground state electronic transitions in gases and solids, so every material known becomes strongly absorbing somewhere in the UV. Finding optical materials becomes very difficult around 200 nm, another reason

[†]The blue channel is about 10 times less sensitive, so the light-adapted eye has its peak sensitivity at 550 nm, with its 50% points at 510 and 610 nm.

for using vacuum systems. This same property makes solid state UV detectors difficult to build; photocathodes are the usual choice in the VUV.

2.2.3 Infrared

The IR starts at about 800 nm, although since the human eye is nearly blind beyond 700, the point at which light becomes invisible depends on its brightness. Out to about 2.5 μm is the near infrared, defined roughly as the region in which glass lenses work and decent photodiodes are available. From there to about 10 μm is the mid-IR, where radiation from room temperature objects is still not a critical limitation, and beyond 10 μm almost to 1 mm is the far-IR, where everything is a light bulb, including the detector itself.

The infrared contains the fundamental vibrational frequencies of molecular bonds (e.g., C—H), so most of the interesting spectroscopy happens there. Infrared spectral lines are often combinations of rotational and vibrational transitions, whose frequencies add and subtract; it is more convenient to quote such systems in terms of frequency than wavelength. The usual frequency unit is not hertz, but wave numbers,[†] that is, how many waves fit in one centimeter. Fortunately, the speed of light is a round number to three decimal places, so the rough conversion that $1\text{ cm}^{-1} \approx 30.0\text{ GHz}$ is convenient for rapid calculation. Interesting spectroscopy happens between 100 cm^{-1} and about 5000 cm^{-1} (100 μm to 2 μm). Note the near-total failure of interesting spectroscopy and good detectors to overlap. Lots of work has gone into that problem.

2.3 RADIOMETRY

In order to be able to choose between illuminators, we need some way to compare them. Measurement and comparison of illumination conditions is the aim of radiometry. Radiometry is an almost completely isolated subfield of optics. This is a bit strange, since all designers need to know how much power crosses a given surface element into some solid angle, in some frequency interval, and most of the time they do. It sounds simple. It would *be* simple, too, if it weren't for the names.

Radiometric nomenclature is important in fields such as remote sensing and architecture, but because it's such a horrible mishmash, it remains a foreign language to most instrument designers, engineers as well as physicists. For one thing, the names have no mnemonic value of their own, and preexisting terms that everybody understands (*intensity* and *brightness*) have been redefined to mean something completely different; for another, there's an impenetrable thicket of redundant units in various subfields. Ideally, technical language provides a concise method of communicating difficult concepts precisely, and radiometric and photometric nomenclature is an excellent example of how *not* to do that.

We'll eschew the footcandles per fortnight and stick with something the author can cope with: $L_{\overline{\nu}}(\nu, \theta)$, the power emitted per square meter per steradian per hertz.[‡] This is the most natural description of a low coherence source such as a light bulb. Its official name is something out of Edgar Allan Poe—*spectral radiance*—and it has a welter of

[†]The official name for the inverse centimeter is the kayser, which nobody uses.

[‡]The subscript Greek *nu* is easily confused with subscript ν , which means something else, so we write it with an overbar, $\overline{\nu}$, instead.

other names and symbols. Everything else is an integral of the spectral radiance over one or more of the quantities in the denominator: frequency, solid angle, area. You can find charts and tables and explanations of this in several sources. One thing to remember is that radiometrists like to do things in per-wavelength units instead of the per-frequency units that most physicists and engineers tend to prefer, leading to mysterious factors of ν^2 turning up on account of the chain rule.

Unless you need to communicate with diehard radiometrists, and don't care that few others will understand, stick with SI units written out in full. Use flux for total power through a surface, flux density for power per unit area, and brightness or radiance for flux density into unit solid angle. Watts per square meter per steradian per hertz are widely understood, internationally standardized, and won't get you into trouble, even if you can't tell a metrelambert from an apostilb.[†]

Photometry is similar, except squishier: the aim there is to compare illumination conditions not in W/Hz, but by what a standardized human eye defines as brightness. Since human eyes have different spectral sensitivities under different lighting conditions (rods vs. cones), this is a bit fraught, so it is solved by fiat: there are two published curves for the scotopic (dark adapted) and photopic (light adapted) response of the eye, and you just multiply L_{ν} by that and integrate over ν to get the visual brightness or *luminance*, measured in lumens/m²/sr. The photopic curve peaks at 550 nm, and is down 50% at 510 and 610 nm. The unit of “photometric power” is the lumen, corresponding roughly with the watt; for historical reasons 1 W at 552 nm (540 THz) is equivalent to 683 lumens, probably the largest prime number ever used for unit conversion.

The reason we can't completely ignore all this is that making accurate measurements of optical power, flux, fluence, and so on is *hard*, which is why there are people who specialize in it. A decent DVM costing \$200 will measure voltage and current to an accuracy of 0.1% or so over a dynamic range of 120 dB. Nothing remotely comparable exists for optical power measurements: $\pm 1\%$ is quite respectable even for narrow signal level ranges. The skill of radiometry is in making decent measurements, not in expressing them in units.[‡]

2.4 CONTINUUM SOURCES

A basic distinction between sources is whether their outputs are predominantly in a spectral continuum or in isolated spectral lines. Lasers and low pressure gas discharge sources have narrow line spectra, high pressure discharge tubes (such as arc lamps and flashlamps) produce a mix of broadened lines and continuum, and incandescent objects produce continuum only. Within each class, there are distinctions as to how wide the spectral features are. The groupings are somewhat rough-and-ready, but nonetheless useful.

In practical applications, the other main distinction is spatial coherence—basically how much of your source light you can corral into your optical system, and how small a spot you can focus it into. Piping laser beams around is easy, because the beam is so well behaved; all the light can be pointed in the same direction. Other sources are not so tractable. Large diameter thermal sources cannot be focused down as tightly as

[†]They're the same thing.

[‡]One of the most heroic feats of radiometric calibration to date is the Cosmic Background Explorer (COBE) satellite—see Kogut et al., *Astrophysical Journal* **470**(2), 653–673 (1996).

small diameter ones. Since we're always trying either to get more photons or to reduce the source power (to save cost, weight, and electric power), extremely low coherence sources such as large arc lamps or tungsten bulbs present problems in illuminator design.

2.4.1 Black Body Radiators

Electromagnetic radiation from an object in thermal equilibrium with its surroundings is called *black body radiation*. Since the closest thing to pure black body radiation is generated experimentally by putting a small hole into a hot cavity, the name may seem odd; however, it is the same as the radiation emitted by a purely theoretical object whose coupling to the electromagnetic field is perfect—every photon hitting it is absorbed. More subtly, by the second law of thermodynamics, it must have the maximum emittance as well. If this were not so, a black body would spontaneously heat up or cool down in isothermal surroundings, since it would absorb every photon incident but emit more or fewer. Thus a black body forming one wall of an isothermal cavity must emit a radiation spectrum exactly the same as that which fills the cavity. Since resonator effects in a large cavity are slight, they can modify the black body's spectrum only slightly (asymptotically, not at all), so a black body radiates cavity radiation regardless of where it is placed. This is a remarkable result.

Such an object would indeed appear black when cold, and many objects have emission spectra that are reasonably well approximated by that of a black body. Among these approximate black bodies are the Sun and tungsten bulbs. The ratio of the total power emitted by some real surface at a given wavelength to that predicted from an ideal black body is known as its *spectral emissivity*, or emissivity for short.[†] It depends strongly on surface composition, angle of incidence, finish, and coatings, but tends to vary slowly with frequency. The ratio of total power absorbed to total power incident is the absorptivity, but this term is almost never used, since it is identical to the emissivity, by the same thermodynamic argument we used before. Landau and Lifshitz, *Statistical Physics Part I*, present an excellent discussion of this topic. We're going to do everything in per-frequency units, for conceptual unity and calculating convenience.

The power spectral density curve of black body radiation peaks at a frequency given by the Wien displacement law,

$$h\nu_{\text{peak}} = 2.8214k_B T, \quad (2.1)$$

where, as usual, h is Planck's constant, ν is frequency, k_B is Boltzmann's constant,[‡] and T is absolute temperature. It has a FWHM of about two octaves, in round figures. If the spectral density is quoted in per-wavelength units, the peak is slightly shifted toward the blue. (Why?) Although the position of the peak moves, black body curves from different temperatures never cross; heating up a black body makes it glow more brightly at all wavelengths (this is also obvious from the second law, if reflective colored filters are admitted).

The formulas for black body radiation tend to be a bit confusing, since different authors quote formulas for different things: exitance from a semi-infinite black body,

[†]Emissivity is formally the average value over some spectral range, but since the emissivity changes slowly with frequency and is anyway a function of temperature, the distinction is not usually emphasized in real life.

[‡]This is one of the few places in this book where Boltzmann's constant k can be confused with the propagation constant k , so we need the inelegant subscript B .

cavity energy density per unit frequency, or per unit wavelength, and so on. The most fundamental is the energy density per hertz of the cavity radiation,

$$e_0(\nu, T) = \frac{2hn^3\nu^3}{c^3\{\exp[h\nu/(k_B T)] - 1\}}, \quad (2.2)$$

where n is the refractive index of the (uniform and isotropic) medium filling the cavity.

There is no special direction inside the cavity, so the energy is equally distributed over all propagation directions. The rate at which a component with wave vector \mathbf{k} leaves the cavity through a patch of area dA is $c\mathbf{k}\cdot d\mathbf{A}/k$, where the vector $d\mathbf{A}$ is dA times the outward directed unit surface normal $\hat{\mathbf{n}}$. Thus there is a cosine dependence on the angle of incidence θ , and the spectral radiance leaving the cavity is

$$L_{\nu}(\nu, T) = \frac{c}{n} \cos\theta e_0(\nu, T) = \frac{2hn^2\nu^3 \cos\theta}{c^2\{\exp[h\nu/(k_B T)] - 1\}} \quad (2.3)$$

which is the most useful all around black body formula. Integrated over all ν , this yields the Stefan–Boltzmann formula for the total power emitted from unit surface area of a black body into unit solid angle,

$$P_{\text{tot}}(T) = \sigma n^2 T^4 \cos\theta, \quad (2.4)$$

where σ is Stefan’s constant, approximately 1.8047×10^{-8} W/m²/sr/K⁴. A real object cannot (as we saw) radiate any more than this, and most will radiate quite a bit less; not everything is black, after all.

The cosine dependence of black body radiation arose, remember, because there was no special direction. A source with this cosine dependence is said to be *Lambertian*, and broad area sources with Lambertian or near-Lambertian angular dependence are said to be *diffuse*.

2.4.2 Radiance Conservation and the Second Law of Thermodynamics

In imaging with thermal light, it’s important to keep in mind a fundamental limitation: *no passive system can form an output whose surface radiance is higher than that of the source*. We’ve seen this as conservation of phase space volume. This can be shown rigorously from Maxwell’s equations, and for paraxial optics it follows from the *ABCD* matrix for an imaging system: the image-side NA goes as $1/M$. For our purposes we’ll use a thought experiment and the second law of thermodynamics.

Consider two reservoirs, initially at the same temperature, and well insulated except for a sufficiently small heat engine (SSHE) running off any temperature difference from R_{hot} to R_{cold} . Drill a hole in each reservoir’s insulation, and place a conceptual highly asymmetric optical system (CHAOS) in between (also suitably insulated). Add well-silvered baffles so that only light that will make it through the CHAOS gets out of the reservoir insulation.

Call the $A\Omega'$ products looking into each end of the CHAOS $A\Omega'_{\text{hot}}$ and $A\Omega'_{\text{cold}}$. Then the net radiative energy transfer from hot to cold will be

$$\dot{Q} = \frac{\sigma}{\pi} (T_{\text{hot}}^4 A\Omega'_{\text{hot}} - T_{\text{cold}}^4 A\Omega'_{\text{cold}}), \quad (2.5)$$

assuming Lambertian radiation (which should be right, since it's cavity radiation we're considering).

If $A\Omega_{\text{hot}} < A\Omega_{\text{cold}}$, the hot reservoir will spontaneously heat up until $\dot{Q} = 0$, which will happen when $T_{\text{hot}}/T_{\text{cold}} = (A\Omega_{\text{hot}}/A\Omega_{\text{cold}})^{0.25}$, and the SSHE will run forever as a perpetual motion machine. If we allow the CHAOS to contain lossless reflective optical filters and amend (2.5) appropriately, the same argument shows that radiance conservation applies for each wavelength independently, not just in aggregate. If the two ends are immersed in different media, then there's a factor of n^2 that has to be added to Stefan's law on each end, and the math comes out the same: there's no temperature drop created.

2.4.3 Tungsten Bulbs

Tungsten bulbs are excellent optical sources for many purposes. They are quiet, stable, cheap, broadband, and reasonably bright. Although their radiance (in $\text{W}/\text{m}^2/\text{sr}$) is far lower than a laser's, they can put out a lot of photons cheaply. Their electrical to optical conversion efficiency is excellent—75% or better—provided that the application does not require throwing away major fractions of their output.

The primary problem with tungsten bulbs is their relatively short operating life, which is limited by the evaporation of the tungsten filament, with the subsequent formation of hot spots, leading to filament breakage. The rate of tungsten evaporation goes as $e^{-10500/T}$, while the brightness of the bulb is proportional to T^4 , which makes the trade-off between lifetime and brightness pretty sharp; you really can't push tungsten beyond about 3300 K. This isn't quite as bad as it sounds, because hot tungsten's emissivity is around 0.45 in the visible but only 0.15–0.2 in the IR, so that you get proportionately more visible photons than you expect, and the bulb color looks about 100 K hotter than it is (maximum color temperature ≈ 3400 K).

The trade-off between lifetime and brightness can be exploited in appropriate circumstances: a tungsten bulb run at 10% higher than rated voltage will have three times shorter life and about a third higher output. Alternately, running it at 10% lower than rated voltage will give five times longer life at about 30% less output.

Evaporated tungsten collects on the bulb envelope, which produces progressive dimming with time, as well as some spectral shift. The evaporation rate can be reduced by reducing the bulb temperature or by increasing the pressure of the gas filling the envelope, which retards the diffusion of tungsten vapor away from its source. Thermal losses and safety considerations put a practical limit on the gas pressure.

Tungsten-halogen bulbs offer a partial solution to the tungsten loss problem, through a clever regenerative mechanism that keeps the envelope clean. A bit of iodine is introduced into a small, very strong quartz envelope around the filament. Evaporating tungsten combines with the iodine to form tungsten iodide, which has a low boiling point. The envelope is hot enough to prevent tungsten iodide from condensing on it, so that the halide stays in the gas phase until it encounters the filament and is redeposited as metallic tungsten. Unfortunately, the redeposition does not take place selectively on the hot spots, so the filament does not heal itself. The longer life of halogen bulbs is due to the high gas pressure inside, which by slowing vapor diffusion *does* cause selective redeposition.

Running a bulb on AC helps its long-term stability, since electromigration of the filament metal partly cancels out when the current changes direction frequently.

Temperature variations over a cycle give rise to intensity variations. These aren't that small—asymptotically, their power spectrum goes as $1/f^2$ (1 pole). Square wave

AC drive helps by keeping the power dissipation constant over a cycle. Ripple in the dissipated power will cause thermal forcing of mechanical oscillations of the filament. You can't readily simulate this, so be sure to test it if you're planning anything fancy.

The other major problem with tungsten bulbs (as with other low coherence sources) is how to keep the source radiance as high as possible as the light passes through the bulb envelope and condenser. More on that soon.

2.4.4 Glow Bulbs and Globars

Occasionally one runs across these odd ceramic light bulbs. They provide nice wide, uniform diffuse thermal emission in the IR, but are not very bright as the ceramic seldom goes much above 1000 K. They are primarily useful in FTIR spectrometers. Good ones have low thermal mass so you can do chopping measurements by pulsing the light source (see Section 10.9.1).

2.5 INTERLUDE: COHERENCE

In various places we've already encountered the idea of coherence, which is basically a measure of how well different parts of an electromagnetic field stay in phase with each other. It's time to be a bit more systematic about what we mean by that. It's a big subject, though, and is explained well in J. W. Goodman's *Statistical Optics*. We're just going to dip our toes here.[†]

Coherence theory is easiest to understand in the context of two famous interferometers: Young's slits (actually pinholes) for spatial coherence, and the Michelson interferometer for temporal. Both of these experiments involve sampling the fields at different points in space-time, and looking at the time-averaged interference term. The result of the experiments is a fringe pattern; the visibility of the fringes expresses the degree of coherence between the fields at the two points. This is an appealingly intuitive definition of coherence, and what's more, it is powerful enough for all optical purposes, provided the averaging time is a small fraction of a cycle of the highest signal frequency we care about.

Thermal sources such as tungsten bulbs are said to be *spatially incoherent*; the phase and amplitude of the light emitted from a point on the surface are independent of that from points more than a wavelength or so away. As a result, the intensity of the light at any point in the system can be calculated by finding the intensity due to a single arbitrary source point, and integrating over the source, without regard for the interference of light from different points on the source.

The light from a tungsten bulb is also *temporally incoherent*, meaning that the phase and amplitude of light from the same source point at some time t is independent of what it was at $t - \tau$, for sufficiently large τ (several femtoseconds for a 3000 K bulb).

These properties may be substantially modified by spatial- and frequency-selective filters, condensers, and even free-space propagation. For example, starlight is temporally as incoherent as tungsten light, but spatially highly coherent, due to the small angular size of the star as seen from Earth. A nearby star would have an angular size of a few

[†]This discussion assumes some familiarity with interferometers. Have a look at Section 1.6, or for more basic stuff, try Hecht and Zajac or Klein and Furtak.

nanoradians. The angular size of an incoherent (thermal) source determines its spatial coherence; the divergence of the beam is equal to the angular size of the source. (Why?)

The basic quantity in coherence theory is the *complex degree of coherence* γ , which is the normalized statistical cross-correlation of the optical field at two points (\mathbf{P}_1, t) and ($\mathbf{P}_2, t + \tau$),

$$\gamma_{12}(\tau) \equiv \frac{\Gamma_{12}(\tau)}{\sqrt{\Gamma_{11}(0)\Gamma_{22}(0)}}, \quad (2.6)$$

where Γ_{12} is the usual ensemble-averaged statistical cross-correlation (see Section 13.5),

$$\Gamma_{12}(\tau) \equiv \langle \psi(\mathbf{P}_1, t)\psi(\mathbf{P}_2, t + \tau) \rangle. \quad (2.7)$$

A pure temporal coherence discussion sets $\mathbf{P}_1 = \mathbf{P}_2$ and usually drops the subscripts.

Given a screen with unresolved pinholes at \mathbf{P}_1 and \mathbf{P}_2 , the fringe pattern at some point \mathbf{x} on the far side of the screen from the source is

$$I(\mathbf{x}) = I_1(\mathbf{x}) + I_2(\mathbf{x}) + 2\sqrt{I_1(\mathbf{x})I_2(\mathbf{x})}\text{Re} \left\{ \gamma_{12} \left[\frac{|\mathbf{x} - \mathbf{P}_1| - |\mathbf{x} - \mathbf{P}_2|}{c} \right] \right\}. \quad (2.8)$$

Temporal coherence can be increased by a narrowband filter, at the price of throwing most of the light away; you can make fringes using light with a frequency range $\Delta\nu$ if the time delay τ between components is less than $1/\Delta\nu$ (if $\tau > 1/\Delta\nu$ you get fringes in the spectrum instead). The envelope of the fringes in a Michelson interferometer will decay as the path difference is increased, so that we can define coherence time more rigorously as the equivalent width of the squared modulus of the fringe envelope,

$$\tau_c \equiv \int_{-\infty}^{\infty} |\gamma(\tau)|^2 d\tau. \quad (2.9)$$

The *coherence length* sounds like a spatial parameter but is in fact defined as $c\tau_c$. The real spatial parameter is the *coherence area*, A_c , the area over which a given field's phase stays highly correlated. It is analogously defined as the two-dimensional (2D) equivalent width in the $\xi\eta$ plane, where $\mathbf{P}_2 - \mathbf{P}_1 = (\xi, \eta, 0)$.

$$A_c \equiv \iint_{-\infty}^{\infty} |\gamma_{12}(0)|^2 d\xi d\eta. \quad (2.10)$$

Young's fringes also fall off in strength as the pinholes are moved apart, but measurements must be made in the plane of zero path length difference in order that temporal coherence effects not enter. In practice, $\Delta\nu$ must be small enough that at least a few fringes are visible for us to be able to measure the fringe visibility as a function of pinhole separation, but mathematically this can be swept under the rug.

Coherence theory predicts how these properties change and what their effects are. It also provides a fairly rigorous conceptual framework in which to describe them. It is a huge topic, which is really beyond our scope. It is treated well in Goodman and in Born and Wolf.

2.5.1 Speckle

Light scattered from a rough surface undergoes wholesale modification of its plane wave spectrum. A single plane wave gets scattered into a range of angles that is characteristic of the surface. At each angle, the light from different locations interferes together, producing strongly modulated, random-looking fringes called *speckle*. Speckle is a three-dimensional (3D) interference pattern, and the characteristic size of the speckles in radians gives the characteristic length scale of the roughness in wavelengths. Speckle is a particular problem in laser-based full-field measurements and when using diffusers (see Section 5.7.11). Although it's usually a nuisance, speckle does contain information about the position and shape of the surface; this information can be extracted by electronic speckle pattern interferometry (ESPI, frequently called TV interferometry).

When low coherence light is scattered from a rough surface, its angular spectrum is modified in a fairly simple way; usually it just spreads out some more, and the scattered radiance remains reasonably smooth in both position and angle. The lack of speckle is a consequence of both spatial and temporal incoherence; in each propagation direction, temporal incoherence (i.e., wide optical bandwidth with no special phase relationships, as in short pulses) smears the fringes out in the space domain (due to the path differences), and spatial incoherence makes the phase of the interference pattern at each wavelength vary at optical frequencies, so that no fringe contrast can be observed in any reasonable measurement.

2.5.2 Imaging Calculations with Partially Coherent Light

Besides the two limiting cases of fully coherent and incoherent sources, there's a broad range of partially coherent cases, where the source linewidth and angular subtense are big enough to notice but too small to dominate. Most of the time, partially coherent imaging is done by starting with an incoherent source such as a light bulb. There are lots of books and papers telling how to predict the results of a partially coherent imaging system, but for most of us, getting through that mathematical machinery is a big time investment. Another way, less general but conceptually much easier, is to replace the incoherent source by a collection of monochromatic point sources, calculate the (coherent) image from each one, and then add all the resulting intensities by integrating over the source distribution in space and frequency. Integrating the intensities expresses the fact that the expectation value of the interference term between two incoherent sources (or two parts of a single one) is 0. Even from a single point, light of two different frequencies produces fringes that move at optical frequency, averaging to 0 (but see below). The interference term in (1.67) is thus 0, and the total photocurrent from the detector is the sum of those from each source point.

If your source is inherently partially coherent (e.g., starlight coming through a turbulent atmosphere), this procedure breaks down and you have to use the big mathematical guns.

2.5.3 Gotcha: Coherence Fluctuations at Finite Bandwidth

One coherence effect that we're going to run into again and again is intensity fluctuations due to the interference of light with a delayed copy of itself. This effect is often overlooked, but it limits the SNR of measurements a lot of the time; it's normally what sets the ultimate SNR of laser-based fiber sensors, for example, and Section 19.1.1 has a very sad story about what can happen when it bites. It's true that the DC photocurrent from

an interferometer whose path difference is $\gg 1/\Delta\nu$ will be the sum of the intensities of the two beams, but that doesn't mean that the instantaneous noise current is just the sum of their noises.

The autocorrelation is an ensemble- or long-time-averaged quantity, whereas we're actually detecting the instantaneous intensity, including the instantaneous interference term, with only the short-time averaging due to our finite measurement bandwidth.

Whenever we have two beams adding together, they do interfere, because at any instant the total energy arriving at the detector goes as the integral of E^2 ; all that happens at τ_C is that the *DC component* of the interference becomes small. So where does the interference signal go?

Temporal coherence limits usually arise from the gradual building up of phase fluctuations as Δt grows. When this happens, the interference contribution from each point doesn't go away, it just gets turned into noise spread out over a wide bandwidth, about $2/\tau_C$. A light bulb has a bandwidth of two octaves, centered around 700 nm, so its temporal bandwidth is 600 THz and its coherence time is 2 fs. It takes a lot of noise power to fill that up.

Furthermore, since these AC fluctuations are uncorrelated across the surface of an incoherent source, they average to zero pretty well too if the source and detector are big enough. On the other hand, if you build an interferometer using a spatially coherent source of limited temporal coherence, it will give you fringes, and this is true whether you meant it to be an interferometer, or just have a stray fringe or two.

In a single-mode fiber measurement, for example, we get no spatial averaging whatever, so we have to cope with the whole interference term, changed into random noise, thundering into our detector. If there is no delayed wave to interfere with, it's no problem, but since the effect is gigantic and the intrinsic SNR of optical measurements is so high, it doesn't take much delayed light to reduce the SNR catastrophically. The detection problem for FM noise in a system with multiple paths of different delays is treated in Section 15.5.12; here we'll just wave our arms a bit.

Consider a Michelson interferometer using a well-collimated, 820 nm diode laser. It is VHF modulated (see Section 2.14.3) so that $\Delta\nu \approx 1$ THz ($\Delta\lambda \approx 2.2$ nm), but has a single transverse mode. Its spatial coherence is then nearly perfect. Recall from Section 1.5 that the instantaneous intensity reaching the photodetector varies as $I_1 + I_2 \pm 2 \cos \phi \sqrt{I_1 I_2}$ (see Section 1.6), where ϕ is the instantaneous phase difference between the two beams. Because ϕ is constant across the detector, spatial averaging of the fringes does not reduce the fringe contrast in the detected signal. For $\Delta t \gg 1/\Delta\nu$, the variance of ϕ , $\langle |\phi - \langle \phi \rangle|^2 \rangle \gg (2\pi)^2$, so ϕ modulo 2π is a uniformly distributed random variable. The noise probability density is therefore that of the cosine (peaked at the edges). The noise current variance is then $2I_1 I_2$, spread out over a frequency range from 0 to about $\Delta\nu$. The resulting 1 Hz current noise i_n as a fraction of the peak interference term is on the order of

$$\frac{i_n}{\sqrt{i_1 i_2}} \sim \frac{2}{\sqrt{\Delta\nu}}, \quad (2.11)$$

which can easily dwarf the shot noise. Using our 1 THz wide diode laser example, two 1 mW beams at 820 nm will generate fluctuation noise on the order of $2 \text{ nW/Hz}^{1/2}$ (optical), so that the detected noise will be nearly 40 dB (electrical) greater than the shot noise. This can be a pretty big shock when you're expecting the interference signal to go away completely. The same diode laser with the VHF modulation turned off might have $\Delta\nu \approx 10$ MHz, in which case the measurement will be much quieter for small

path differences, but much noisier for large ones; for $\Delta t > 100$ ns, i_n will be more like $0.4 \mu\text{W}/\text{Hz}^{1/2}$, an absolutely grotesque figure—the photocurrent noise will be 86 dB over the shot noise.

Even for path differences well within $1/\Delta\nu$, this effect can be very large; at 1% of the coherence length ($3 \mu\text{m}$ path difference here), the RMS phase wobble is on the order of 0.01 radian, which in the 1 THz instance would bring the fluctuation noise near but not below the shot noise level. Note that this noise current grows as i , as signal does, not \sqrt{i} as shot noise does; increasing the laser power will make it relatively worse. We see that even if $\Delta\nu$ is huge compared with the measurement bandwidth, these seemingly unimportant fluctuations may be the dominant noise source. How's that for an insidious gotcha?

One final addendum: it is not unusual for a very small-amplitude side mode in a diode laser to have very significant mode partition noise, and it gets out of phase with the carrier really easily in the presence of dispersion.

2.5.4 Measuring Laser Noise in Practice

Okay, so there are lots of noise sources to worry about. How do we measure the actual noise? In two steps: first, calibrate the detection setup; and second, measure the noise.

For calibration, take whatever setup you have and shine a battery-powered, incandescent-bulb flashlight on it. Move the flashlight in and out to make the photocurrent roughly the same as in your laser measurement. This will give you a beautifully calibrated white noise source (the shot noise of the photocurrent) of spectral density $i_N(1 \text{ Hz}) = \sqrt{2eI_{\text{DC}}}$, delivered right to your photodiode. Measuring the noise level in the dark (circuit noise) and with the flashlight (circuit noise + shot noise) will give you the gain versus frequency and noise versus frequency of your measurement system. This will allow you to compare the noise of the laser to the shot noise.

To do the actual measurement, the best bet is a spectrum analyzer. You can do this with a fancy digitizing scope, by taking the discrete Fourier transform of the output, squaring it, and averaging many runs together. (In a DFT power spectrum with only one run, the standard deviation is equal to the mean, so to get good results you have to average quite a few spectra or take the RMS sum over many adjacent frequency bins—see Section 17.5).

Alternatively, you can get a filter whose bandwidth coincides with your measurement bandwidth, and just measure the total power that gets through it, using a sufficiently fast RMS-DC converter (the author often uses an old HP3400A RMS voltmeter, which has a 10 MHz bandwidth). The noise power is the time-averaged total power of the signal, minus the DC power. Since the noise and the DC are uncorrelated, it doesn't matter whether you subtract the DC before or after taking the RMS. You can't do a good job with just a scope and no filter, because you won't know the bandwidth accurately enough.

Really work at relating the laser noise to the shot noise, because if you don't, you won't know how well or badly you're doing.

2.6 MORE SOURCES

2.6.1 LEDs

Light-emitting diodes (LEDs) are fairly narrowband continuum sources, which are nowadays ubiquitous because of their low cost and good properties. Their lifetimes are so long

(hundreds of thousands of hours, without abuse) that you'll probably never wear one out. Their other properties are intermediate between those of tungsten bulbs and lasers; their spectral width is about 5–20% of their center frequency (say, 50–100 nm full width for a 600 nm LED) and their emitting areas are smallish (100 μm diameter or so). Their moderately wide bandwidth means that their temporal coherence length is extremely small, about 3–10 microns, which is often very useful where some spectral information is needed but etalon fringes must be avoided.

LEDs are available in a wide range of colors, from the near IR (900–1100 nm), where they are used for remote control, infrared LANs, and multimode fiber connections, to the ultraviolet (350 nm). Some LEDs have more than one spectral peak, so if this matters, make sure you measure yours.

Red LEDs are the most efficient, but the other wavelengths are catching up fast. Already the best LEDs are more efficient than tungsten bulbs, and their efficiencies continue to climb. White LEDs are really a species of fluorescent bulb—a UV LED excites a fluor in the package, and it's the fluor that limits the coherence and lifetime.

There are a few good reasons for the historical inefficiency of LEDs. The efficiency of generating photons from carriers is reasonably good, approaching 100% in direct bandgap devices. Unfortunately, these photons are generated deep inside a highly absorbing semiconductor material with a refractive index of about 3.3. This hurts in two ways: the semiconductor absorbs the light, and the critical angle for light exiting the semiconductor into air is $\arcsin(1/3.3)$, or about 18° , so that the solid angle through which photons can exit is only 0.3 steradians, or about 2.3% of a sphere. Fresnel losses on reflection limit the actual extraction efficiency of uniformly distributed light to below 2%. If the substrate is absorbing, the totally internally reflected photons will mostly be absorbed before getting another chance. The absorbing substrate problem is addressed by building heterojunction LED with transparent top layers (and newer ones have transparent substrates as well, with only the junction region being absorbing). The total internal reflection problem is helped by putting the LEDs in a plastic package ($n \approx 1.5$), which more than doubles the extraction efficiency, and by roughening the surfaces, which gives the light lots of chances to escape.

An LED's intermediate source size means that its spatial coherence is not very good compared with a single-mode laser, but is still better than you can do with a tungsten bulb. Again, this is usually a virtue except in interferometry.

Unlike nearly any other source, LEDs can have their intensities varied continuously over a huge range (from 0 to their maximum rated power) with little or no spectral shift, by merely changing their drive current. Spectral shifts remain small until maximum power is approached, at which point most LEDs shift slightly to the red,[†] due to high die temperatures (LED wavelengths tend to drift at about +100 ppm/K) and to high level injection effects, which narrow the effective bandgap. This throttling capability is of great value in extending the dynamic range of really low cost sensors.

Ordinary LEDs are fairly slow compared with laser diodes, being limited by carrier lifetime rather than stimulated emission. You can modulate an LED at a few megahertz, but garden variety ones don't go faster than that. IR LEDs intended for communications can go a bit faster, up to about 100 MHz.

[†]That is, toward longer wavelengths, not necessarily closer to 630 nm. This is really visible with yellow LEDs, which turn orange at very high drive currents.

The main problem with LEDs is their packaging. For efficiency reasons, nearly all LEDs come in molded plastic packages with cylindrical sides and very steep lenses on the top surface. The optical quality is poor and the die centration erratic. Furthermore, the wide angle of illumination of the sides leads to lots of internal reflections in the package. Together, these defects produce far-field patterns that look like a cheap flashlight's. The plastic package also limits the LED's power dissipation, because its thermal resistance is very high, so that the heat has to travel down the leads. For applications where lower spatial coherence is desirable, an integrating sphere will produce a nice Lambertian distribution, and even a ground glass, 3M Magic Invisible Tape, or holographic diffuser will help considerably. Some LEDs have glass powder mixed with the plastic, producing a frosted effect with better uniformity but lower efficiency and wider angular spread.

If higher spatial coherence is needed, consider using a surface-mount LED with a flat surface. Another way is to use a laser diode below threshold as a spatially coherent LED. The spatial characteristics are lithographically defined and the windows are good, so their beam qualities are excellent. Their coherence length is short and depends on how hard you drive them (it gets longer as threshold is approached), but is generally shorter than in multimode lasing operation, and of course very much shorter than in single-mode operation. Laser diodes are more expensive and (below threshold) won't produce as much total power as a bright LED. Compared to the work of removing the encapsulant from an LED or using LED chips and wire bonds, a diode laser below threshold is a bargain, unless you need more than a few hundred microwatts of power.

2.6.2 Superluminescent Diodes

The laser-diode-below-threshold idea can be extended to higher gains by removing the regeneration (e.g., by AR-coating one or both facets) or by gouging the back facet with a scribe. The resulting device has stimulated emission but little or no regeneration, and is called a superluminescent diode (SLD). Its coherence length is short, although not as short as a LED's, and it is a bit less sensitive to back-reflections than a diode laser. SLD linewidths are a few percent of the center wavelength.

SLDs are far from immune to back-reflection, however; residual resonator effects usually leave ugly peaks and valleys in their spectra, so the fringe visibility doesn't usually wash out to 0 as smoothly or as quickly as one would like. If you're using SLDs with fiber pigtails, make sure you measure the spectrum *after* the pigtail is attached, and generally be suspicious.

Commercial SLDs (sometimes confusingly called SLEDs) are available with output powers up to 10 mW or so. They're made with AR coatings and have to be run at very high current density due to the lack of regeneration. This makes them even more sensitive to optical feedback than the homemade ones. A perpendicularly cleaved fiber end is enough to make them lase, which (due to the high level of pumping) will usually blow them up. Their modulation bandwidths are more like LEDs' than laser diodes' (~10 MHz).

2.6.3 Amplified Spontaneous Emission (ASE) Devices

Optical amplifiers all have spontaneous emission. Leaving out the input signal and cranking up the pump power causes the spontaneous emission to be amplified like any other signal, producing an ASE source. Commercial ones produce tens of milliwatts from

ytterbium (1000–1100 nm), neodymium (1060 nm), or erbium (1560 ± 40 nm), but are expensive. Their center wavelengths can vary a lot depending on the fiber material and the emission bands selected, and their linewidths are a few percent, becoming narrower at higher gains. A better behaved ASE device is the *frequency-shifted feedback* laser, with an acousto-optic cell in the cavity. These permit multipass gain while preventing normal lasing, because the N th pass gets shifted by $2N$ times the acoustic frequency.

2.6.4 High Pressure arc Lamps

Thermal sources cannot exceed the brightness of a black body at the same temperature. The evaporation rate of the most refractory metals limit filament temperatures to below 3500 K, so plasmas are the only way to get brighter thermal sources. There are several types of plasma sources, divided into high and low pressure plasmas, either continuous (arc lamps) or pulsed (flashlamps).

High pressure devices are much brighter than low pressure ones. Although they are thermal sources, they are far from thermodynamic equilibrium, so their spectra are usually dominated by strong emission lines corresponding to electronic transitions in the component gases. These lines are very broad (tens of nanometers typically) due to collisional broadening from the high pressure of hot gas. There is also a weaker black body continuum from the plasma, and still weaker, lower temperature thermal radiation from the hot electrodes.

Arc lamps are available with a variety of fill gases, ranging from hydrogen to xenon, including especially mercury/argon and sodium/argon. Mercury arcs are the most efficient, but produce most of their output in the UV (especially near 254 nm), and their visible output is very green due to a strong line at 546 nm. Sodium vapor is nearly as efficient as mercury, but is very orange (590 nm). Xenon arcs appear white to the eye, which makes them popular for viewing applications. Hydrogen and deuterium lamps produce a strong continuum in the UV, with D_2 units going far into the vacuum UV, limited mainly by the windows (for VUV operation, they are used windowless, with differential pumping).

The emission lines get broader and weaker, and the continuum correspondingly stronger, as the gas pressure is increased; extremely high pressure arcs (200 atmospheres) can make even mercury vapor look almost like a 6000 K light bulb.

The plasmas are inherently unstable, so the spectrum and power output fluctuate with time in a $1/f$ fashion, with $\pm 10\%$ power variations being common. There are variations caused by strong temperature gradients in the tube, causing big changes with time and warmup. The pressure increases with temperature, which changes the broadening. The position of the arc is unstable with time as well, at the level of 0.1 to a few millimeters, which will cause the light from your condenser to wander around. In some bulb types, the spot moves erratically at a speed of 50 m/s or faster, leading to lots of noise in the kilohertz range. They are a great deal noisier than tungsten or LEDs, and their noise has a spatially dependent character that makes it hard to correct for accurately.

Commercial arc lamps are divided into short arc and long arc styles. The length of the arc partly governs the source size, and thus the spatial coherence of the light: you can do a better job of collimating the light from a short arc lamp. They come in two main package types: squat ceramic and metal cylinders with quartz windows and integral reflectors, and long quartz tubes with a bulge in the middle. The ceramic type is much easier to use but has a few idiosyncrasies. Most such lamps have an arc formed perpendicular to the window, so that the cathode gets in the way of the beam. The package contains a

parabolic reflector, which roughly collimates the light reflected from it (typical angular spreads are $\pm 1^\circ$ to $\pm 3^\circ$, with the spread increasing as the bulb ages). Unreflected light floods out into a wide angle; the illumination function is thus rather strange looking, with a strong doughnut-shaped central lobe and a weaker diffuse component. On the other hand, the arc and reflector come prealigned, which substantially simplifies changing the bulb.

EG&G sells ceramic arc and flashlamps that circumvent this problem, by arranging the electrodes transversely (so that the current flow in the arc is parallel to the window and the electrodes are not in the way), and using a spherical reflector with the arc near its center of curvature. All the light therefore emanates from the arc, and both the strange pupil functions are avoided. You still have to collimate the light yourself, however. One choice will probably be better than the others in your application.

2.6.5 Flashlamps

A flashlamp or flashtube is a pulsed arc lamp, usually filled with xenon or krypton, and is used where very bright, hot, short-duration illumination is needed. Flashlamps can reach source temperatures of over 10^4 K, making them instantaneously about 100 times brighter than a tungsten source of the same size. When used at high power, their plasmas are *optically thick* (i.e., highly absorbing[†]), so that their radiation is dominated by the continuum component. At lower pulse energies, the plasma is cooler (closer to 5000 K than 10,000 K), and the line spectrum is more pronounced.

Flashlamps are powered by capacitor discharge; you charge up a special capacitor[‡] that is connected to the lamp through a very small inductor (which may be just the leads). When the capacitor is fully charged, you trigger the arc, either with a tickler coil, wound round the cathode end of the tube and driven with a 10 kV pulse, or by momentarily increasing the voltage across the tube with a pulse transformer.

Once the arc gets going, it discharges the capacitor in a time controlled by the impedance of the arc, wiring, and capacitor parasitic inductance and resistance (including dielectric losses), until the voltage drops below that needed to keep the arc going. This variable excitation means that the light intensity is strongly peaked with time; unless you do something special, there's no flat top to the light pulse, regardless of its length. In fact, the current can easily oscillate due to LC resonance in the capacitor and leads. Flashlamp manufacturers publish data allowing you to design a critically damped network, so that the total pulse width is minimized for a given tube, capacitor, and drive energy.

The total duration is usually between 100 μ s and 10 ms, although active controllers (using big power MOSFETs) can make it as short as 1 μ s. Pulse energy is limited by explosion of the bulb and ablation of the electrodes. Getting long life (10^5 – 10^7 flashes) requires operating the bulb a factor of 6–10 below the explosion energy. Tubes with explosion energies between 5 J and 10 kJ are readily available.

Peak power is limited by how fast you can push that much energy into the bulb, and average power is limited by cooling. Flashlamps can work at kilohertz repetition rates if the pulses are weak enough. There is considerable variation in the pulses; a well-designed flashlamp supply with a new bulb may achieve 0.1% rms pulse-to-pulse variation, but

[†]It is perfectly correct from a thermodynamic point of view to call the central plasma core of a 1 kJ flashlamp "black." Try it some time.

[‡]Capacitors intended for rapid discharge have very low internal inductance and resistance, to improve their discharge time constants and avoid dumping lots of heat into the capacitor itself.

typical ones are more like 1%. They have the other noise characteristics of regular arc lamps, for example, strong spatial variations, wandering arcs, and lots of $1/f$ noise. Their wall plug efficiency can be over 50% when integrated over all ν , but is more typically 10%.

Flash initiation depends on the generation of the first few ions in the gas, which is inherently jittery compared with laser pulses, for example. Typical jitters are in the hundreds of nanoseconds and are a strong function of the bulb design. Jitter can be reduced with a “simmer” circuit, which maintains a low current arc between flashes, or a pretrigger that initiates a weak arc slightly before the main pulse.

Flashlamps and arc lamps with concentric reflectors are intrinsically more resistant to source noise from arc wander, because the reflector projects an inverted image back onto the arc. If the arc moves to the left, the image moves to the right. For an optically thick plasma, the effect can even help stability a little; if the left side gets hotter, its image dumps heat into the right side.

2.6.6 Spark and Avalanche Sources

You can get 400 ps rise time from a spark in a mercury switch capsule.[†] If subnanosecond speed is enough, and you don't need narrow linewidth or high brightness, this can be a convenient and cheap alternative to a pulsed laser. The sharp current pulse will couple into everything in the neighborhood, so careful electrical shielding will be needed. Spark sources have all the jitter and pulse-to-pulse variation problems of flashlamps, and since the rapid breakdown is what we're after, they can't easily be fixed with simmer circuits and so on.

2.7 INCOHERENT LINE SOURCES

2.7.1 Low Pressure Discharges

We are all familiar with low pressure gas discharge tubes, such as Geissler tubes (beloved of undergraduate optics labs) and ordinary domestic fluorescent bulbs. Low pressure discharges produce a line spectrum by electrically exciting gas molecules and letting their excited states decay by radiating photons. The positions of the spectral lines are characteristic of the molecular species doing the radiating, while the line strengths and linewidths are more strongly affected by experimental conditions such as pumping strategy, temperature, pressure, and other gases present.

The gas fill is often a mixture of gases to improve arc characteristics and aid energy transfer to the desired excited state of the radiating molecule, for example, argon in mercury lamps and helium in neon bulbs. Argon is a particularly popular choice; besides being cheap and inert, it has a low first ionization potential, so that arcs are easily struck. The striking voltage depends on pressure, with a minimum at about 1 torr, rising at high vacuum and at high pressure.[‡]

Their fairly narrow linewidths (moderate temporal coherence) and low spatial coherence make low pressure lamps useful for qualitative imaging interferometers. The most

[†]Q. A. Kerns, F. A. Kirsten, and G. C. Cox, Generator of nanosecond light pulse for phototube testing. *Rev. Sci. Instrum.* **30**, 31–36 (January 1959).

[‡]M. J. Druyvesteyn and F. M. Penning, *Rev. Mod. Phys.* **12**, 87 (1940).

common example is preliminary testing of optical surfaces via Newton's rings and Fizeau fringes (see Section 12.6.2).

The exceptions to the line spectrum rule are low pressure hydrogen and deuterium lamps, which emit a bright continuum extending into the VUV. Deuterium lamps emit almost no visible light at all, which is helpful in UV experiments because the filters and gratings that would otherwise be needed are often highly absorbing. These arcs are often run at extremely low pressure, especially deuterium, where the interior of the bulb connects immediately with the vacuum system, and deuterium leakage is controlled only by differential pumping.

2.8 USING LOW COHERENCE SOURCES: CONDENSERS

2.8.1 Radiometry of Imaging Systems

We saw in Section 2.4.1 that the product of area and projected solid angle is invariant under magnification and Fourier transforming, so that in the absence of vignetting or loss, the source and image radiance are the same, and the average radiance is the same in the pupil as well. Unfortunately, it is not easy to avoid vignetting when the source radiates into 4π steradians.

It is a common error to assume that a higher power bulb will get you more photons for your measurement. Thermodynamics dictates that no point in the optical system can have a higher radiance than that of the source (i.e., the arc or filament). Since filament and arc temperatures are lifetime-limited, this means that there is an absolute upper limit to how bright you can make your illumination with thermal sources. A higher wattage bulb is merely larger, not brighter. Accordingly, upgrading the bulb will produce a larger illuminated area or aperture, but the same radiance. If that's what you need, a higher wattage bulb can help. If not, you'll have to improve your condenser and collection optics instead, or switch to laser light.

2.8.2 The Condenser Problem

The challenge in illumination systems is to achieve the characteristics needed for good measurements: constant angular distribution, stability in intensity and polarization, and uniformity of brightness and spectrum across the field. Ideally it should also use the available source photons reasonably efficiently. Optical systems using tungsten illumination are almost always extremely inefficient with photons, because being efficient is tough. It isn't usually worth the trouble, either, because the equipment is mains powered and tungsten bulb photons are very cheap,[†] but other illuminators don't have this luxury.

As shown in Figure 2.2, the simplest kind of condenser is an imaging system, basically a lens corralling the source photons and pointing them toward the sample. Trying to get more photons will involve increasing the solid angle of the light collected (i.e., the NA of the condenser on the bulb side).

The envelopes of some kinds of bulbs (especially tungsten halogen) are serious impediments to illuminator design. If the optical quality of the bulb envelope is good enough, it may be possible to use the reflector to put the image of the filament right next to the real filament, which will improve the illuminated area and hence the efficiency.

[†]About $\$10^{-23}$ each.




Figure 2.2. A typical condenser is a low quality imaging system with care taken in controlling thermal loads.

On the other hand, an irregularly shaped bulb envelope, or one with a poorly designed internal reflector, makes it nearly impossible to approach the thermodynamic limit for light coupling. Most fall into this category.

Example 2.1: Fiber Illuminators. Getting the highest radiance from a tungsten bulb may not always involve a condenser. For fiber bundle coupling, just image the filament on the bundle, using an NA (on the bundle side of the lens) high enough to overfill the fiber NA. The rest of the housing can just toss light back into the filament to reduce the electrical power required. If the bundle has a very high NA (some are 0.5), it may not be easy to get lenses like that. Figure 2.3 shows ways of dealing with this problem, using an array of lenses, a Fresnel lens, or a fast spherical mirror. With the bulb just off to one side of the center of curvature and the bundle at the other side, this does a good job of filling the bundle NA and diameter. The mirror doesn't have to be too good; there's no reason the image quality can't be reasonable, since both source and bundle are near the center of curvature. The mirror can have a cold mirror coating (which passes IR and reflects visible light) to reduce the heat loading of the fibers. (Pretty soon we'll just stick a white LED at the business end and get rid of the bundle entirely.)




Figure 2.3. Condensers for fiber bundles: a simple lens or a spherical cold mirror.

Most people are better off buying condensers when they need them, since condenser design is a somewhat specialized art, and combination lamp housings and condensers are widely available commercially. If you design your own, take apart a couple of different commercial ones to see the tricks they use. Be sure to remember the heat absorbing glass, which is a tempered filter that dissipates the heat that would otherwise toast your optical system or your sample.

2.9 LASERS

Lasers have an immense literature, and come in a bewildering variety. We'll concentrate on the properties of the ones most commonly seen,[†] especially diode lasers. See Siegman[‡] for much more theoretical and pedagogical detail.

Lasers rely on an external energy source, the *pump*, to provide a population inversion and hence laser gain. Oscillation requires a minimum pump power, known as the *threshold*. How many times above threshold the pumping is determines how rapidly the oscillation builds up, among many other parameters. For instruments, the distinctions between laser types center on wavelength, tunability, power level, pumping source, pulse width, and linewidth.

2.9.1 Mode Structure

A laser is a resonant cavity with a gain medium inside, and some way of coupling the light in the cavity to the outside, usually a partially transparent cavity mirror, the output coupler. The resonant condition requires that the cavity be carefully aligned, which is a fiddly process if it's user adjustable.

Cavities always have more than one mode, so that the laser's spatial and temporal behavior will be a superposition of cavity modes. Real lasers have complicated mode properties, but we can do a good job by decomposing the fields in terms of the modes of an unperturbed, empty resonator. These break down into the transverse modes, which are well described by Gaussian beams, and longitudinal modes, which are like the resonances of an organ pipe. All will have slightly different frequencies in general, constrained by the limits of the gain curve of the laser medium. For our purposes, transverse modes higher than TEM₀₀ are unwanted, but fortunately they are usually easily suppressed, so that most commercial lasers have a single transverse mode.

The longitudinal modes are defined by the requirement that the round-trip phase in the laser resonator should be a multiple of 2π radians, so that multiple bounces will give contributions that add in phase and so produce a large total field. Accordingly, longitudinal modes are spaced by very nearly $\Delta\nu = c/(2n\ell)$, where 2ℓ is the round-trip distance (twice the cavity length) and n is as usual the refractive index of the material filling the cavity. For a large frame gas laser, $\Delta\nu$ is on the order of 200 MHz, but it ranges up to ~ 150 GHz for a cleaved cavity diode laser and 1 THz for a VCSEL. Diffraction effects and the pulling of the resonant frequency by the gain slope of the laser medium and coatings often prevent the longitudinal modes from being exact harmonics of each other,

[†]The most important neglected types are dye lasers. These were once very important for lab systems but have been eclipsed lately, except for those few applications requiring wide tunability in the visible.

[‡]Anthony E. Siegman, *Lasers*. University Science Books, Mill Valley, CA 1986.

which produces some noise and spurious signals in the detected photocurrent. Coupling between the modes will make them lock to each other if these effects are small enough.

Simplified laser theory suggests that a single longitudinal mode should always dominate, with a linewidth that narrows enormously for pumping rates well above threshold. In practice, multiple longitudinal modes are more commonly found, and pumping harder makes it worse. For interferometric applications, multiple modes are undesirable, because they cause the coherence length to be much shorter than it needs to be and produce noise and sensitivity variations with a periodicity of twice the cavity length. These applications often need single longitudinal mode (single-frequency) lasers.

A laser may even oscillate in more than one spectral line at once; this is suppressed by using a dispersing prism or grating as part of the cavity, so the cavity can only be aligned for one line at a time.

The equal spacing of these modes means that their configuration should be periodic with time, like an ideal piano string.

2.9.2 Relaxation Oscillation

Lasers work by running electrons through a loop consisting of three or four quantum transitions. Each of these transitions has its own time constant, which as any circuits person knows, can lead to instability due to accumulated phase shifts. Good lasers have one dominant time constant, namely, the long lifetime of the upper state of the laser transition. The other time constants lead to excess phase shift, which generally causes ringing in the laser output and level populations whenever a sharp change in pumping occurs, just like a feedback amplifier with too small a phase margin. Due to the historical fact that lasers weren't developed by circuits people, this ringing is misnamed *relaxation oscillation*. It isn't really oscillation, but it does limit the modulation response of the laser, and it causes excess noise near the peak—generally about 1 GHz for cleaved cavity diode lasers and more like 10 GHz for VCSELs. Diodes and solid state lasers show much stronger relaxation oscillations than gas lasers in general.

2.10 GAS LASERS

Gas lasers cover the entire visible, UV, and infrared range in isolated lines with very narrow tuning ranges. Typical examples roughly in decreasing order of usability are helium–neon (HeNe, 632.8 nm), argon ion (488 and 514.5 nm), helium–cadmium (HeCd, 442 and 325 nm), carbon dioxide (CW and pulsed IR at 10.6 μm), nitrogen (pulsed, 337 nm), and excimer (pulsed deep UV). Except for HeNe and CO₂, these are really suitable only for lab use or specialized applications such as laser light shows and medicine. Gas lasers are big, bulky, and fairly noisy (0.5–2% intensity noise).

Diode lasers are very popular just now because they're cheap, small, and mechanically rugged. Don't despise HeNe's, however—they are much less sensitive to optical feedback, are very frequency stable, have long coherence lengths (300 m for single-frequency units), are really well collimated, and work as soon as you plug them in. A HeNe is just the right medicine for a lot of ills. Besides, they come with their own power supplies and are totally invulnerable to electrostatic discharge (ESD) damage—if you don't smash it, it'll always just work. Low gain gas lasers such as HeNe's have no spatial side modes to worry about, so there aren't too many really weird sources of noise, apart from the occasional baseband mode beat (see Section 2.13.7).

HeCd lasers are very noisy, and N_2 lasers cannot be run CW due to bottlenecks in their level transition rates (you get 10 ns pulses no matter what you do—an example of relaxation oscillations that really oscillate). Excimer lasers are used for laser ablation of materials and are also widely used for semiconductor lithography. Their spatial coherence is very low; although their output may be below 200 nm, you can't focus their output more tightly than a millimeter or so.

Some gas lasers can be made single-frequency, either by reducing the gain and carefully tuning the cavity length (HeNe, $P < 1$ mW), or by putting a Fabry–Perot etalon (see Section 1.6.2) in the cavity (ion lasers). This is effective but painful, requiring temperature stabilization of one or both resonators to keep the tuning sufficiently stable.

Gas lasers are usually pumped by running an electric arc through the gas fill in a sealed tube with mirrors or Brewster windows fused to its ends. They are astonishingly inefficient, due to short upper state lifetimes and lots of deexcitation paths. For example, an Ar-ion laser is doing well to have a wall plug efficiency (laser emission/electrical input) of 0.02%, which leads to big AC supplies and water or fan cooling. The main exception is CO_2 lasers, which can hit 25% wall plug efficiency. All those Teslaesque features cause severe low frequency noise in most gas lasers, with the high power ones being worse.

Grating tuned gas lasers such as ion and metal vapour units usually have a diffraction grating or Littrow prism inside their cavities to select a single line, although multi-line units do exist. Getting a laser aligned is sometimes a painful process; it's discussed in Section 12.9.8. Apart from HeNe's and sealed CO_2 s, gas lasers need significant amounts of tender loving care; if you're going to use one, learn how to clean and maintain the mirrors and windows properly, and if there's someone available who can show you how to align it, ask.

2.11 SOLID STATE LASERS

Solid state lasers are based on electronic transitions in impurity sites in solid crystals or glasses. Typical examples are Nd:YAG (1.06 μm), ruby (694 nm, pulsed), Alexandrite (700–820 nm), Ti:sapphire (0.66–1.2 μm , femtosecond pulses possible), and color center (0.8–4 μm , widely tunable).

Solid state lasers have better characteristics than gas lasers in general, and are much more flexible. Their efficiencies are limited by the pumping strategy employed; diode laser pumped Nd:YAGs have very respectable wall plug efficiencies (in the percents). The host crystals change both the center wavelengths and the strengths of the emission bands, so that the laser wavelength of a given ion will move by tens of nanometers depending on the host. Ruby ($Cr^{3+}:Al_2O_3$) was the first laser of all, and ruby units are still sold, though you'll probably never use one. Neodymium YAG (Nd^{3+} ion in yttrium aluminum garnet) lasers are much more common, because they have really good performance over a wide range of powers and pulse widths, from CW to the tens of picoseconds. The very long upper state lifetime of a YAG laser (250 μs) makes it efficient, and also allows it to store lots of energy. Powerful pulses can be achieved using Q -switching, where the cavity resonance is spoilt while the upper state population builds up, and is then rapidly restored (usually with an acousto-optic or electro-optic cell, but sometimes with a bleachable dye in low energy units). The cavity finds itself way above threshold and very rapidly emits a huge pulse, typically 1–10 ns wide.

A somewhat less well-behaved pulse can be obtained by letting the cavity build up a big circulating power, and then using the AO or EO cell to allow this energy to exit the cavity in one round-trip time, a technique called cavity dumping.

Pulsed YAGs are usually pumped with flashlamps, whose pulse width is a good match to the upper state lifetime. Efficiency suffers because the pump bands (the spectral regions in which pump light can be converted to laser light) are not well matched to the flashlamp spectrum. CW YAGs are nowadays usually pumped with 808 nm diode lasers. Single-longitudinal-mode diode-pumped YAGs are highly efficient, can be pretty quiet, too—the best ones are 10–30 dB above the shot noise at moderate modulation frequencies, and even closer above 20 MHz or so.

YAG lasers are often frequency doubled, yielding green light at 532 nm. These are familiar as low quality green laser pointers, but with a bit more work, A diode-pumped, frequency doubled, single-frequency YAG laser is a very attractive source: intense, efficient, quiet, and visible. High peak power units can be tripled or quadrupled into the UV. Doubling has to be done inside the cavity to get decent efficiency, and there are a number of sources of potential instability in the process.

Other neodymium-doped crystals are Nd:YV₀₃ (yttrium vanadate) and Nd:YLF (yttrium lithium fluoride). Both have better thermal behavior than YAG—they're birefringent, which reduces losses due to thermal stress-induced birefringence (just as in PM fiber), and both have low dn/dT , which reduces thermal lensing and so improves beam quality. Vanadate lasers are also easier to run on the weaker 914 and 1340 nm lines. Neodymium-glass lasers can produce higher pulse energy than YAG, but the very low thermal conductivity of glass compared with YAG or vanadate severely limits the repetition rate, so that the average power is lower for the same size device. The spectrum is shifted slightly and shows a lot of inhomogeneous broadening. The upper state lifetime is much shorter and hence efficiency is lower.

Titanium-sapphire lasers are now the predominant choice for femtosecond pulses, having taken over the honor from colliding-pulse modelocked (CPM) dye lasers.

Diode-pumped YAGs, usually called diode-pumped solid state (DPSS) lasers, are potentially extremely quiet; their long upper state lifetime has the effect of filtering out the wideband AM noise of the pump, and the pump laser can be power-controlled over a wide bandwidth, so in the absence of mode jumps, mode-partition noise, thermal lensing, and photorefractive instability, a single-frequency DPSS laser is a very well-behaved device.

Decent solid state lasers are not simple or cheap devices and must be pumped with another light source, usually a flashlamp or diode laser. They typically run \$5000 to \$50,000, and lots more for the really fancy ones. Sometimes there is no substitute, but you'd better have a decent grant or be selling your system for a mint of money.

2.11.1 Modelocked Lasers, Optical Parametric Oscillators, and Other Exotica

Like vacuum systems, laser design is an arcane and lore-rich subject on its own—see Koehner's *Solid State Laser Engineering* and Siegman's *Lasers* if you need to know more details. All there's space for here is that the longitudinal modes of an ideal laser form a complete basis set for representing any pulse shape you like—including a δ -function. By forcing all the modes to oscillate in phase at time t_0 , $t_0 + \Delta$, $t_0 + 2\Delta$. . . , many cavity modes combine to produce a train of narrow pulses, with the period of cavity

round-trip time. The phase and amplitude shaping is done by a modulated attenuator or by a saturable absorber, which attenuates bright peaks less than their dimmer wings, thus narrowing the pulse on every bounce. Until the limit set by loss and dispersion is reached.

Pulsed lasers are much harder to maintain than CW units in general, with the high pulse power, low rep rate, lamp-pumped modelocked units being generally the worst, on account of the thermal transients, vibration, coating damage, and general beating up that the crystals take when they're clobbered that hard. If you're exclusively a CW person, count your blessings—you've led a sheltered life.

Aside: Beam Quality. A complicated system such as a picosecond optical parametric generator pumped by the third harmonic of a lamp-pumped YAG laser is not going to have the beam quality of a TEM₀₀ HeNe, no matter what. The good ones have decent spots in the near field that turn into mildly swirly messes in the far field; bad ones produce beams that resemble speckle patterns regardless of where you look. Small changes in conditions can produce severe beam degradation, so make sure that whatever source you have, you can measure its beam profile accurately and easily. Especially in the IR, it is amazing how many people trying to make complicated measurements (e.g., sum-frequency generation spectroscopy) don't have any idea of their beam quality. Many kinds of nonlinear sources, especially OPOs, have beam profiles that are pulse-height dependent, so you can't normalize them with an energy meter no matter how hard you try.

2.12 DIODE LASERS

The champs in efficiency and cost effectiveness are diode lasers. Diode laser photons are more expensive than tungsten photons, but cheap for lasers, and their output quality is good. Linewidths of 10–100 MHz or so are typical for single frequency units. Wall plug efficiencies can reach 40%, with slope efficiencies ($\partial P_{out}/\partial P_{in}$) of 80% or even higher.

Most diode lasers are of the cleaved cavity (Fabry–Perot) type: the die is cleaved to separate the lasers, and the cleavage planes form the cavity mirrors. They are usually coated at the rear facet with a moderately high reflector. The front facet can be left uncoated or may have a passivation layer. The Fresnel reflection on leaving the die is sufficiently strong (40% or so) to sustain oscillation with the extremely high gains available in the active region.

Current is confined to a very small active region by giving the surrounding material a slightly higher bandgap, so that the forward voltage of the diode in the active region is smaller, and it hogs all the current. The resulting spatial restriction of laser gain helps guide the beam, and lasers relying on this are said to be *gain guided*. A better method is *index guiding*, where the refractive index profile makes the active region a stable waveguide as well as confining the current.

Due to this monolithic construction, diode lasers are mechanically extremely rugged, although their associated collimators generally are not. Their packages are similar to old-style metal transistor cans, for example, TO-3, 9 mm (TO-8), and 5.6 mm (TO-72) for the smallest ones. Most have a monitor photodiode in the package, to sense the light emitted from the rear facet of the diode for intensity control. The monitor is big enough that it is not completely covered by the laser die, so that it will pick up light scattered back into the diode package, which makes it somewhat unreliable in practice.

Diode lasers are available in narrow wavelength ranges centered on 380, 405, 635, 650–690, 750–790, 808, 830, 850, 915, 940, 980, 1310, 1480, and 1550 nm (plus a few at oddball wavelengths like 1.06 and 1.95 μm), which are defined by their intended use. The shortest-wavelength diodes widely available are 405–410 nm high power multimode units (from 20 mW up to 100–200 mW, linewidth 1–2 THz) for Blu-Ray discs. At this writing (late 2008) the cheapest way to get these is to cannibalize them from the drives, because they cost \$1000 apiece otherwise. There's a big hole between violet (405 nm) and red (633 nm), where there are currently no good laser diodes, so diode-pumped solid state lasers are the usual choice. Lasers used in optical drives have lots of power—roughly 150–200 mW for 658 nm DVD burner diodes. Longer wavelengths, 2–3 μm , can be obtained with quantum cascade lasers based on superlattices.

The market for diode lasers is dominated by optical storage and telecommunications, so diode types come and go, and they're expensive in ones and twos. CD/DVD player lasers are the cheapest: 650 nm, 5–7 mW or less. These cost around \$10 unless you're a DVD player manufacturer and get them for \$0.50 each in 100,000 piece quantities. Fancier ones, such as the SDL-5400 series of Figure 2.4, can give you >100 mW of single frequency power. Multiple longitudinal mode diodes reach about 1 W, and multiple transverse mode units (based usually on many apertures side by side) are approaching

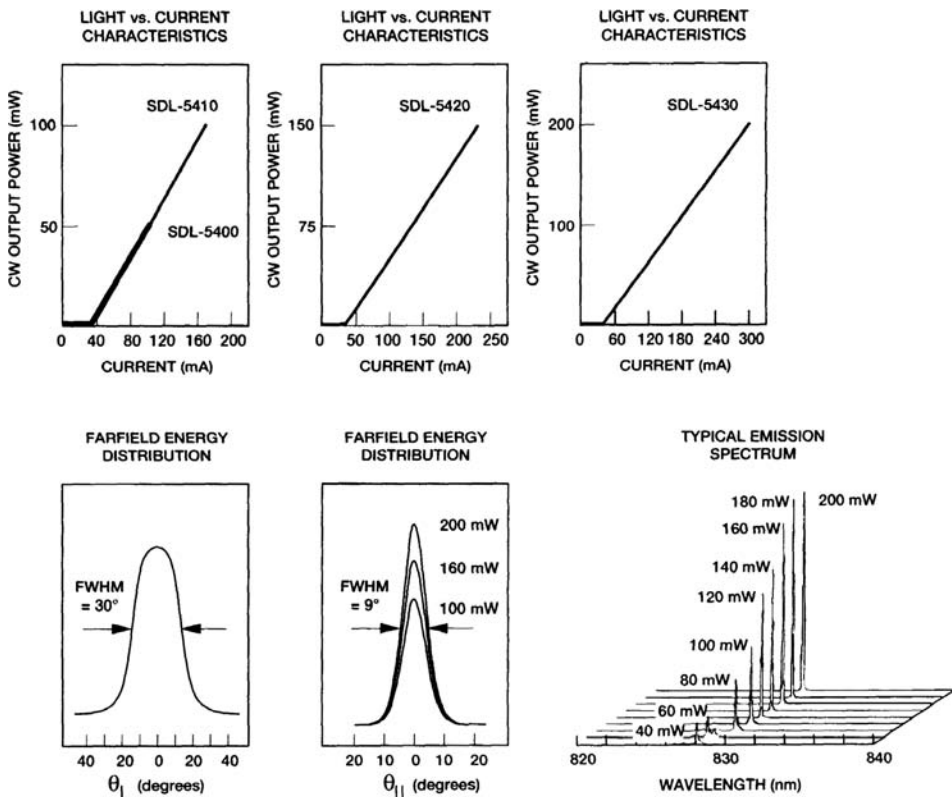


Figure 2.4. Beam parameters of Spectra Diode Labs (now JDSU) 5420 diode laser (From SDL, Inc., and reprinted with permission © 1994, 1997 SDL, Inc.)

1 kW CW. These bar-type diodes are best used as light bulbs, because their coherence properties are poor.

The bad things about diode lasers are that they are extremely delicate electrically, that they are inconvenient to drive and to collimate, that their tuning depends on everything under the sun, and that their mode hops and instabilities will drive you nuts if you're unprepared or need better-than-ordinary performance.

The other major drawback with all these low cost devices is that they are tightly tied to consumer applications in printers, DVD players, and so on. When the consumer technology moves on, the lasers go away; if you're using them in your instrument, and you don't have a big stock of them, you're out of luck.

2.12.1 Visible Laser Diodes

Working in the infrared is more difficult than in the visible, and the available spatial resolution is less for a given NA. Commercial visible laser diodes (VLDs) work near 670, 650, 630, and 405 nm, which is convenient for human eyes. High power VLDs are used in CDR and DVDR drives, so they've become quite common. Unfortunately, VLDs behave significantly worse than their IR brethren, with poorer tunability and frequency stability and more mode hopping. If you need good noise performance from your diode laser, look elsewhere.

Aside: Wavelength-Division Multiplexing (WDM) and Chirp. The intrinsic bandwidth of optical fiber is extremely wide—much wider than any foreseeable electronic switching frequency. Thus the most convenient way to get more capacity from a fiber-optic data link is by sending many independent signals down the same fiber, using a grating device or a sequence of narrowband couplers to combine the signals at one end and separate them at the far end. The wavelengths of the optical carriers conform to a standard spacing, the so-called *ITU grid*. The grid frequencies are multiples of 100.0 GHz, from 184.8 to 201.1 THz, with 50 GHz offsets also available. Achieving low crosstalk between channels requires that the laser tuning be very stable, both with time and temperature and (which is more of a challenge) with fast modulation. Normally the frequency of a semiconductor laser changes quite a bit—as much as a couple of nanometers (~ 250 GHz at $1.5 \mu\text{m}$)[†] during a fast turn-on, which would scribble all over the grid if it weren't controlled. If the chirp is linear in time, it can be used with a grating to compress the pulse, but most of the time chirp is just a nuisance. VCSELs, even multimode ones, have much wider spectra but only a few modes; they have much lower chirp than FP lasers, which have many modes for energy to slosh about in (see Section 2.12.7).

2.12.2 Distributed Feedback and Distributed Bragg Reflector

Diode lasers need not use the cleaved ends of the die as mirrors. Distributed feedback (DFB) lasers use an active waveguide with a grating superimposed on it, resulting in very high selectivity. The similar distributed Bragg reflector (DBR) scheme uses separate gratings in passive waveguide regions, which can be tuned separately with a second bias

[†]Paul Melman and W. John Carlsen, Interferometric measurement of time-varying longitudinal cavity modes in GaAs diode lasers. *Appl. Opt.* **20**(15), 2694–2697 (1981).

current (some have more than one tunable grating segment). At one time, DFB lasers had better tunability since the same mechanisms that tune the frequency also tune the gratings to match, resulting in wide temperature and current tuning ranges with few or no mode hops. DFB lasers are expensive and specialized, so they're only available in the telecom bands. Chirp is very damaging there, as we saw, so lots of work has gone into reducing it; although you can tune modern DFB lasers with temperature, they hardly current-tune at all. For wider current tuning, DBR lasers are superior, if a bit less convenient to drive.

2.12.3 Tuning Properties

Everything in the world tunes diode lasers: temperature, current, pressure, and any sort of optical feedback (see Section 2.13.6). All of these effects can be controlled, but you have to be aware of the need. Highly stable diode laser instruments need clever control systems for temperature, current, and mode hop control, which we discuss in Chapter 20 (<http://electrooptical.net/www/beos2e/thermal2.pdf>) and Section 15.9.1.

Among Fabry–Perot lasers, the 800 nm ones are the best behaved. They tune at rates of around -0.015 to -0.08 cm^{-1}/mA and -0.1 cm^{-1}/K in a single mode (for a 5 mW unit), and around -4 cm^{-1}/K on average due to mode jumps. They can be tuned through 1 – 2 cm^{-1} by current in between mode jumps, and much further via temperature.

Tuning red VLDs is a different story. A small bandgap difference between the active area and its surroundings hurts the current confinement and makes VLDs very temperature sensitive and generally unstable. It is not trivial to find a nice single-mode operating region with a VLD, although it can be done in the lab if you don't need more than 0.5 cm^{-1} of current tuning range and are prepared to hunt. Violet and UV lasers are generally multimode. Stick with the 750–850 nm ones if you need to do anything fancy.

2.12.4 Mode Jumps


The tuning of single-mode diode lasers is discontinuous in both temperature and current, as shown in Figure 2.5, and the discontinuities unfortunately move around slowly with laser aging; even with perfect temperature and current control, your continuous tuning range won't stay the same. Another odd thing is that the tuning curve is multivalued: in some regions, the same temperature and current can support oscillation at two different frequencies. These two modes typically do not oscillate at the same time; it's one or the other, depending on history, a phenomenon called *hysteresis*.

The mode jumps move slowly downwards through the operating current range as you increase the temperature; they move much more slowly than the tuning, unfortunately, so you can't always get to your desired wavelength without external cavity stabilization.


If you're building instruments relying on tuning F-P diodes, you normally have to ensure that you're working at a point where ν is a single valued function of T and I , and that takes some work. On the other hand, you only get DFB lasers at 1.3 and 1.55 μm , and they are two orders of magnitude more expensive, so fixing up F-P diodes to tune reliably is a worthwhile thing to do.

2.12.5 Regions of Stability

Besides the intrinsic Fabry–Perot mode structure, diode lasers in real optical systems always have some feedback, which is liable to lead to instability, and certainly modifies



(a)



(b)

Figure 2.5. Tuning properties of an FP visible diode laser (Toshiba TOLD9211): (a) versus current and (b) versus temperature.

the stable regions. If you vary the temperature and bias current to a diode laser, and look at its noise, you find two-dimensional islands of stability surrounded by a sea of mode hopping. There is some bistability at the edges, since the edges of the islands sometimes overlap, as we saw (they remain separate if you consider wavelength as a third dimension).

These islands can be mapped out fairly easily, and if you are using diode lasers in instruments it is a good idea to do that. Use some simple software for the job, because otherwise you'll be at it quite awhile; since the noise goes up so much during mode

hopping, it is straightforward to design self-testing software for the instrument. The islands change slowly with time. VLDs have much smaller islands than 800 nm diodes.

If you want the islands to be as stable as possible, temperature control the collimating lens and the lens-to-laser spacing too. A collimator design that makes this easy is shown in Example 20.7 at <http://electrooptical.net/www/beos2e/thermal2.pdf>.

2.12.6 Checking the Mode Structure

Use a fine pitch grating (e.g., 2400 lines/mm at 670 nm), aligned so that the first-order beam exits near grazing incidence. If the laser is well collimated, the mode structure will show up clearly. You get a bunch of dim dots at the far-off Fabry–Perot transmission maxima, which are just filtered spontaneous emission, and a few bright ones that are actually lasing. You can spot mode jumps right away with this trick, and sticking a photodiode in just one of the bright spots will give you a visceral feel for just how bad mode partition noise can be.

2.12.7 Vertical Cavity Surface-Emitting Lasers

Almost all lasers have cavities that are long and skinny. This means that they have a great many longitudinal modes to jump around in, which causes noise problems. In addition, they emit out the ends, which is okay for built-up lasers but is inconvenient for diodes. Diode lasers are hard to test before dicing and can't easily be made in two-dimensional arrays. They also have that nasty beam asymmetry and astigmatism to deal with.

A partial solution to some of these problems is the *vertical cavity surface-emitting laser* (VCSEL). A VCSEL is rather like a DFB laser built on end. The cavity mirrors and active region are made by multilayer deposition, and the bias current flows vertically right through the whole thing. The high reflectance of the mirrors allows the active region to be very thin, so that the longitudinal mode spacing is very wide (1 THz). That and the limited mirror bandwidth mean that a typical VCSEL has only one longitudinal mode. The NA of a typical VCSEL is 0.3–0.4, and due to its approximate rotational symmetry, its beam is roughly circular. That same symmetry means that its polarization usually wanders all over the place.

VCSELs also have lots of transverse modes, and they're nearly all multimode. It is quite common for a VCSEL to jump between two or three spatial modes as the current is increased, and wind up in some $N = 6$ mode that looks like a chrysanthemum (See Figure 2.6.) They can be really quick (20 GHz modulation bandwidth), so they work well for things like high speed communication via multimode fiber. There has been some progress made in improving the polarization purity and mode structure, usually by breaking the circular symmetry somehow, but VCSELs are far from a panacea.

Aside: VCSEL Pathologies. Even in datacom sorts of jobs, some VCSELs exhibit weird turn-on artifacts at low repetition rates. The pump current has to flow through all those layers, so there are lots of opportunities for trap states to form. You have to fill them all before the laser will turn on, and at low rep rates, they'll all have recombined before the next pulse arrives. Repopulating them each time slows down the leading edge badly.

VCSELs also tend to have really horrible $1/f$ mode partition noise, which is worse than in ordinary multimode lasers since the modes are randomly polarized. Here's another instance where a good quality polarizer at the laser can help a lot in a differential measurement.

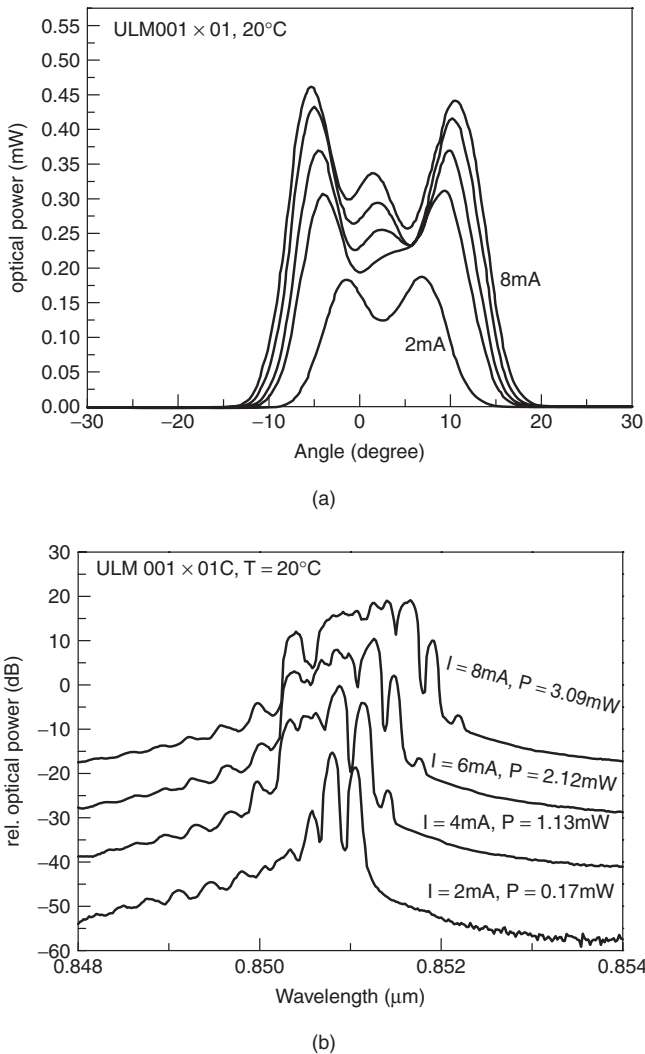


Figure 2.6. Angular beam width (a) and spectrum (b) of an 850 nm multimode VCSEL versus bias current. (Reproduced by courtesy of U-L-M Photonics GmbH.)

2.12.8 Modulation Behavior

You can modulate a laser diode fast enough for almost anything. Generic diodes can be modulated strongly up to 1 GHz, and special ones such as 850 nm communications VCSELs can go much faster, up to 20 GHz. The limiting factors are parasitic inductance and capacitance, and the intrinsic speed of the laser mechanism due to transition rates and the relaxation oscillation peak.

The modulation sensitivity increases toward the relaxation peak and drops off very rapidly beyond it, but you don't want to be working there anyway. The relaxation frequency (and hence the maximum modulation frequency) generally increases as the operating current goes up, up to somewhere near maximum output power, and then starts to slow down slightly.

Like an RF power transistor, the modulating impedance of the diode is normally low, which leads to vulnerability to parasitic inductance; in fact, wiring an RF power transistor right across the diode is a good way to modulate a high power device, but watch out for oscillations at 300 MHz to 3 GHz when you do this—it's possible to have too much of a good thing, and an anti-snivet resistor can come in very handy (see Section 19.7.4).

The main problem in modulating diode lasers is that the intensity and frequency modulate together, so that it isn't easy to get pure AM or pure FM from current tuning. If you're doing a spectroscopic measurement with a diode laser, you need some way of suppressing the giant intensity slope superimposed on your data. Typical methods are FM spectroscopy (Section 10.6) and noise canceler spectroscopy (see Section 10.8.6).

2.12.9 ESD Sensitivity

Diode lasers are so extremely sensitive to electrostatic discharge that ESD is the major cause of reliability problems. ESD in reverse bias causes hot carrier damage to the junction region, increasing the threshold, but forward bias ESD is much more spectacular: because the laser radiation grows so fast, a carpet shock with a 1 ns rise time can generate such a high peak laser power that the output facet of the laser gets blown right off. This is usually easy to spot; it puts stripes in the laser output pattern and dramatically reduces the output power.

Always use ground straps, and design protection into your diode laser mounts: use a DIP reed relay to keep the anode and cathode shorted together when the laser is not in use. A big bypass capacitor (1 μF) will absorb carpet shocks without letting the laser blow up, although this severely constrains the AC modulation possibilities. One way round this is to use a transformer with its secondary wired in series with the the diode to get fast modulation without ESD risk.

2.12.10 Difficulty in Collimating

Edge-emitting diode lasers emit radiation from a stripe-shaped aperture about 1 μm wide by a few microns long. The aperture is small, so that light comes out into a large solid angle, which is very asymmetrical: typical NAs are 0.09–0.15 in one plane and 0.3–0.5 in the other, roughly an elliptical Gaussian beam. This large NA means that the laser–lens distance is very critical; the depth of focus is only a few microns, so that extreme stability is required. Very thin layers of UV epoxy are a good way of affixing the lens and laser together.[†] Stability against mode hopping requires significantly tighter control than focus stability does.

The light is strongly polarized with **E** along the minor axis of the ellipse. The beam from a gain guided laser is a complete mess. It has such large amounts of astigmatism that getting decent collimation is extremely difficult (they're getting rare now anyway). Index guided lasers have some astigmatism, usually about 0.6 wave. This amount is big enough that it needs correction, which is a pain, but not so large that correcting it is unduly difficult if you have a measuring interferometer.

Astigmatism can be reduced with a weak cylindrical lens, but the commercially available ones are normally a poor match to your laser, unless you're very lucky. Fortunately, good collimation can be reliably obtained without the use of astigmatism correction optics

[†]See Example 20.7 at <http://electroptical.net/www/beos2e/thermal2.pdf>.

(0.95 Strehl ratio is often attainable).[†] The trick is to compensate for the astigmatism with a small amount of defocus, so that instead of the wavefront being a cylinder of 0.6λ p-p, it is a saddle shape with a p-p error of 0.3λ . You can also use the collimator slightly off axis, to get a bit of coma and astigmatism to knock the laser's astigmatism flatter still. This approach is unfortunately a bit fiddly; in a production system, it may be better to use a more powerful laser and just chop off all but the center 20% with a circular aperture.

The elliptical beam can be circularized with one or two prisms, where the beam enters near grazing (producing a patch elongated into a circle) and leaves near normal, so that the circular patch defines the refracted beam profile. Gratings can be used in the same sort of way, with lower efficiency. If you're using gratings, make sure you use the deviation-canceling configuration in Figure 7.8, or your beam will wander around with tuning.

An anamorphic lens system,[‡] in this case a telescope made of cylindrical lenses, can circularize the beam as well as cancel the astigmatism, but these are limited by the difficulty of fabricating good quality cylindrical lenses at low cost.

Blue Sky produces lasers with diffraction-limited circular output beams by using a cylindrical microlens right near the laser die, inside the package, which in principle makes a lot of things easier. These are not particularly cheap, but they're good for lab use. (You still have to collimate them yourself.)

2.12.11 Other Diode Laser Foibles

Diode lasers have a number of other funny properties. One is that dust and crud are photophoretically attracted to regions of high light intensity, which unfortunately means the diode laser facet or the center of the laser window. In a high power system, the beam path will get dirty long before the rest of the system, so make sure to control outgassing and dust inside your diode laser-based instruments. Interestingly, dry air turns out to be much better than inert gas—having oxygen around gives the organic crud a chance to oxidize before it builds up enough to cause damage.[§]

2.13 LASER NOISE

Lasers exhibit noise in both intensity and frequency. Apart from ultrahigh resolution spectroscopy, most measurements suffer more from intensity noise than frequency noise. As we'll see in Chapter 10, a lot of ingenuity has been expended on getting rid of the effects of laser noise.

2.13.1 Intensity Noise

Intensity noise pollutes laser-based measurements in two ways. Most measurements have a nonzero baseline, so that a fluctuation in the laser power causes a fluctuation in the

[†]The Strehl ratio is the best single measure of beam quality for instrument purposes, where the beams aren't too ugly. It cannot be greater than 1, and 0.8 corresponds roughly to Rayleigh's $\lambda/4$ criterion for diffraction-limited image quality. See Section 9.5.4 for more discussion of the aberration terms used here.

[‡]That is, one with different magnifications in x and y .

[§]R. Jollay et al., *Proc. SPIE* **2714**, 679–682 (1996).

background signal, which shows up as additive noise in the measurement (additive means that this noise doesn't depend on the signal strength). This additive noise reduces the sensitivity of the measurement and, since laser noise tends to be non-Gaussian in character, may not average out well.

The other effect of laser intensity noise is to cause the signal itself to fluctuate in strength. In nearly all laser-based measurements, the detected signal is proportional to the laser power, so an intensity fluctuation produces a signal strength fluctuation. This is called *multiplicative* noise, or noise intermodulation. For example, consider a single-beam tunable diode laser spectrometer looking at a gas cell. The spectrum consists of absorption lines on a smooth baseline. Intensity noise causes the baseline to fluctuate (additive noise), but also causes the absorption peaks to fluctuate in magnitude (noise intermodulation). (See Section 13.6.11.)

There are a variety of differential and ratiometric detection schemes to help with this problem, of which laser noise cancellation is by far the most effective; it can provide as much as 70 dB reduction in both additive laser noise and noise intermodulation, and gets you down to the shot noise reliably, even with noisy lasers. It needs an extra beam plus about \$10 worth of simple electronics. If your measurement suffers from laser intensity noise, have a look in Section 10.8.6.

2.13.2 Frequency Noise

An oscillator is composed of an amplifier plus a frequency determining device, usually a resonator. The resonator attenuates signals outside its narrow passband but, more to the point, exhibits a phase shift that varies rapidly with frequency. Since oscillation requires that the round-trip phase be an exact multiple of 2π , changes of the resonator length force the frequency to move. The frequency noise of the oscillator is determined by the combination of the amplifier's phase fluctuations and the phase slope of the resonator. Lasers are a bit more complicated in that the resonator may exhibit phase fluctuations too, as when fan vibrations or cooling water turbulence jiggles the mirrors of an argon ion laser. Resonator noise is confined to low frequencies, so it can be dealt with separately by mechanical means.

This picture of frequency noise suggests that lasers can be stabilized by using longer cavities with higher Q , which is true; external cavity diode lasers have linewidths a factor of 10^3 – 10^4 narrower than when the diode's F-P resonator is used. There are also various active locking techniques such as Pound–Drever stabilization, which are beyond our scope but are covered well in Ohtsu.

2.13.3 Mode Hopping

Most lasers have a large number of possible oscillation modes within the gain curve of the medium. Ideally, one of these should dominate all others, but this is often not the case, due to spatial hole burning.[†] Even when it is, the difference between adjacent modes is often small enough that it takes very little perturbation to switch the laser from one mode to another. This will happen during warmup, for example. Sometimes, though,

[†]Hole burning is the picturesque name given to the local reduction of laser gain near peaks of the standing wave pattern in the gain medium. This reduces the gain of the strongest mode, without reducing that of the others equally. See Siegman for more details.

there is no one stable mode of oscillation (usually due to spurious feedback of one kind or another coupling the modes), leading to mode hopping.

Diode lasers are especially sensitive to mode hopping, because their cavities are strongly coupled to the outside (reflectivities of about 40%, vs. 99.5% for a HeNe). A spurious reflection on the order of 1 part in 10^6 can set a diode laser into mode hopping; this leads to a lot of flaky failures that come and go, making testing difficult. The actual mechanism of mode hopping is a complicated interaction between the thermal, plasma-optical, and current confinement behaviors of the diode; a change in the laser tuning changes the power output, which changes the dissipation in the channel, which changes the temperature, which changes the cavity length and index, which changes the tuning, and so forth. Since the active region can be cooled very rapidly by the laser output itself, there is a strong coupling between die temperature, tuning, and power output that leads to instability. Visible diode lasers are the worst for this. VCSELs are designed with cavities so short that their free spectral range is wider than the natural linewidth, so they don't hop between longitudinal modes, but frequently they do between transverse modes.

Mode hopping is most obviously a frequency noise phenomenon, but it results in strong (0.1–1%) intensity noise as well, because the gains of the laser system in adjacent modes are not the same. Mode hopping causes irregular jumps and spikes in the laser power, at rates of 100 kHz or so. Even in a pure intensity measurement, diode laser mode hopping due to incidental feedback is very obnoxious and can render your measurement extremely difficult. Mode hopping makes all your etalon fringes dance around, so that the frequency noise gets converted to intensity noise as well.

2.13.4 Mode-Partition Noise

The total power output of a laser is limited by the pump power among other things, and various saturation effects couple the intensities of the modes, so that the instantaneous power of the laser varies less than that of the individual modes. The switching back and forth of the laser power is called *mode-partition noise*. It is insidious, because it doesn't show up on a power meter but is silently at work screwing up your measurement and producing seemingly irrational results; with a gas laser, it is easily possible for a spatial filter, a knife edge, or even an iris diaphragm to cause the intensity noise to go up by 20 dB; a stray etalon fringe can do the same, since the different modes will see different phase shifts and hence will be demodulated differently. It's pretty puzzling if you don't know the secret.

The quietest lasers are single longitudinal mode units, followed by highly multimode ones, and the worst usually have only a few (2–10) modes. The same is true of optical fiber devices (see Section 8.4.3). It's worth trying a quick experiment with a few-mode diode laser and a grating—if you catch one mode with a photodiode, you'll usually find it much noisier than the entire beam, even in absolute terms.

2.13.5 Gotcha: Surface Near a Focus

Unless you use really effective—60 dB or more, optical—Faraday isolators to protect the diode laser from feedback, make sure that there is no surface in your optical system that coincides or even nearly coincides with a focus. Even if the specular reflection goes off at a steep angle, and so misses the laser, there will be enough scatter to make the

laser mode hop. If you have a diode laser system that mode hops, and you're sure you've eliminated all the near-normal surfaces, look for this one. If it isn't a surface near a focus, it's probably feedback right inside the collimator.

An insidious possibility you may need to watch for is that the focus in question may not be a focus of the main beam, but of a converging stray reflection; a concave coated surface will give rise to a converging beam 1% as strong as the main beam, and even after two more bounces to get back into the laser, it can easily get to the 10^{-6} mode hopping danger level. The ISICL sensor of Example 1.12 has this difficulty if it isn't mounted properly.

2.13.6 Pulling

The oscillation frequency of a laser has to be at a frequency where the round-trip phase delay is an integral multiple of 2π radians. Normally, as the cavity length ℓ changes slightly, or small amounts of contamination collect on the mirrors, the oscillation frequency changes so as to preserve the round-trip phase. However, if one of the mirrors suddenly were to develop a time-dependent phase shift of its own, the oscillation frequency would have to respond to it by shifting, even if the cavity length were perfectly stable.

This is more or less what happens with pulling; some external influence, for example a Fabry–Perot resonator such as an optical spectrum analyzer or merely an incidental reflection, sends delayed light back to the output coupler of the laser. Interference between the light inside the laser cavity and this spurious reflection causes the phase of the light returned from the output coupler to change.

The resulting frequency shift is

$$\Delta\nu \approx \frac{-\Delta\phi}{\partial\phi/\partial\nu + 2\pi\ell/c}, \quad (2.12)$$

where $\partial\phi/\partial\nu$ is the phase slope of the total reflection. Note that even if the spurious reflection is weak, so that the total phase excursion is small (see Section 13.6.9), $\partial\phi/\partial\nu$ can be made very large by using a sufficiently long delay. As a result, the stray reflection can in principle take over the frequency determining role from the cavity. This happens especially in diode lasers, where the cavity is short and the mirrors leaky; sometimes their tuning behavior is determined more by the back-reflection from their collimator than by the cavity mirror.

2.13.7 Mode Beats

If you shine a HeNe laser onto a photodiode and examine the results on a spectrum analyzer set to DC–10 MHz, you'll probably see the occasional strong, narrow spur[†] appear, sweep toward 0, and disappear when it gets to 1 MHz or so.[‡] It may be as strong as 0.1% of the total photocurrent. These odd objects are baseband mode beats. Lasers, like all oscillators without automatic level control, operate in a strongly nonlinear region, which means that all their (supposedly orthogonal) modes are in fact coupled together

[†]That is, spurious signal, see Section 13.5.

[‡]If it doesn't do it right away, put your hand on one end of the housing to cool it down a bit, and then look. A bit of bending caused by thermal gradients will unlock the modes.

with each other. Those mode beats are caused by one mode mixing with a third-order intermodulation product of two others (Section 13.5.3 for more details). Since optics people always have fancier names for things, this intermodulation is called *four-wave mixing*. Small anharmonicities in the laser medium or cavity cause the two products to differ in frequency by about 1 part in 10^9 , causing the few-hundred-kilohertz mode beats we see. The disappearance at low frequency is caused by the modes jumping into lock with each other.

2.13.8 Power Supply Ripple and Pump Noise

Like other oscillators, lasers are sensitive to power supply ripple. The laser output depends on how much gain is available, which depends on how hard it's pumped, which depends on the supply voltage. This is usually worst in big gas lasers, whose large power requirements make quiet supplies harder to build. The advent of good quality switching supplies, whose high operating frequencies make filtering much easier, have improved this problem, but it still persists. You'll definitely see your power supply's operating frequency come through loud and clear. If it sits still, you can usually avoid it, but some supplies change their frequency with load, and that makes it much harder to avoid. Make sure you know this about your particular laser.

Diode lasers have very high electrical to optical conversion efficiencies, so the current noise on the supply translates more or less directly into photon noise. Most diode laser supplies use lousy voltage references, inadequately filtered, to define their output current, which makes the lasers themselves noisier. For bright-field measurements, it is really worthwhile to make sure that your diode laser's bias supply is quieter than the shot noise. This isn't too hard to do—see Section 14.6.7.

Lasers that are pumped optically will also be modulated by the noise of the pump source. Flashlamps are usually the prime contributors, but ion laser pumped dye lasers also suffer from instability due to mode noise in the pump laser.[†] As noted in Section 2.11, single longitudinal mode DPY lasers are very quiet.

2.13.9 Microphonics

The narrow linewidth of a laser comes not from the narrowness of the spectral line doing the lasing, but rather from its being put in a Fabry–Perot interferometer a meter long, and then subjected to regeneration, which narrows it further. Since the narrowness comes from the cavity selectivity, anything that causes the cavity length to fluctuate will produce frequency noise in the laser. This includes vibrations from ambient sound (microphonics), cooling water turbulence, fans, and conducted vibrations from the table. Lasers whose mirrors are firmly positioned (e.g., sealed HeNe units and especially diode lasers) are much less prone to this problem.

Some types of lasers (e.g., medium power diodes and diode pumped YAGs) come in styles with and without fans in the laser head. Avoid fans wherever possible.

[†]Martin C. Nuss, Ursula H. Keller, George T. Harvey, Michael S. Heutmaker, and Peter R. Smith, Amplitude noise reduction of 50 dB in colliding-pulse mode-locking dye lasers. *Opt. Lett.* **15**(18), 1026–1028 (1990).

2.13.10 Frequency Noise

Frequency noise in an oscillator comes from the noise of the active element, restrained by the selectivity of the resonator. The low frequency noise of the amplifier (and any noise or instability in the resonator) gets translated up to become low modulation frequency sidebands on the laser oscillation (see Section 15.9.4), and high frequency noise becomes a more or less white phase noise background. A resonator run at its 10^6 th overtone, such as a laser cavity, makes this a bit more complicated by introducing competition between modes, and the complexities of the laser gain mechanism contribute intrinsic noise and instability, but this picture is still basically right.

Laser frequency noise can be reduced by reducing the noise forcing: quieting down the pumping, using a stable gain medium (e.g., Nd:YAG rather than N_2) when possible, or using a highly mechanically stable cavity. It can also be improved by using a more selective resonator (a longer or lower loss one). An external cavity stabilized diode laser uses both approaches.

There is a region of optical feedback in which temporal coherence collapses completely, and the laser output becomes chaotic. You'll recognize it if it happens to you.

2.13.11 Spatial and Polarization Dependence of Noise, Wiggle Noise

Laser noise is not merely a well-behaved wandering around of the intensity or frequency of the entire beam at once. There are important variations in the noise with position and polarization; for example, vignetting the beam of a single frequency argon ion laser has been known to increase its residual intensity noise (RIN) by an order of magnitude; a weak side mode that was previously orthogonal to the main beam was made to interfere by the vignetting, producing a huge noise increase. Diode lasers are especially prone to this pathology for some reason, but all lasers have spatially dependent noise.

Interference with small amounts of laser light and spontaneous emission whose spatial pattern is different causes the laser beam to wiggle back and forth ever so slightly, a phenomenon called wiggle noise.[†] A gas laser with a second spatial mode that is close to oscillating, or a "single-mode" fiber run at too short a wavelength (so there are really two or three or five modes) are especially bad for this; you don't know what pointing instability is like until you've used a system like that. On the other hand, a truly single-mode fiber with all the extraneous light stripped out is unsurpassed for pointing stability (see Section 8.2.2). The angular size of the wiggle goes as the ratio of the amplitudes of the laser mode and the spurious signal, that is, the square root of their intensity ratio, so this isn't as small an effect as you might think.

Polarization dependence is easier to understand; for example, lasers produce spontaneous emission, which is more or less unpolarized. Because it has a different dependence on pump power, the modulation of the spontaneous emission by pump noise will be different. Because the laser radiation is usually highly polarized, the detected noise will differ in character depending on the rotation of an analyzer.[‡]

[†]M. D. Levenson, W. H. Richardson, and S. H. Perlmuter, Stochastic noise in TEM₀₀ laser beam position. *Opt. Lett.* **14**(15), 779–781 (1989).

[‡]An analyzer is just the same as a polarizer, but the name specifies that it is being used to select one polarization for detection, rather than to prepare a single polarization state for subsequent use.

2.14 DIODE LASER COHERENCE CONTROL

2.14.1 External Cavity Diode Lasers

The tuning and frequency stability of a Fabry–Perot diode laser are limited mainly by the poor selectivity of its cavity. It is possible to use a diode as the gain medium in a conventional laser cavity; this is done by antireflection coating the output facet to get rid of etalon resonances in the diode, and using an external reflector and a grating for selectivity. In these external cavity diode lasers (ECDLs), the external cavity takes over the frequency determining function; because of its length and the good optical properties of air, the linewidth is narrower and the tuning much more stable. Because the diode has such high gain, the cavity needn't be all that efficient, and typically fixed-tuned ECDLs use a grating in Littrow as the cavity mirror, with the specular reflection being the laser output. This makes the output beam alignment tuning-sensitive, so it is usually restricted to fixed-tuned applications. Tunable ECDLs use two bounces off a fixed grating, with a rotatable mirror as the tuning element; that makes the output beam pointing very stable. Very fast-tuning ECDLs use polygon mirrors but these have amplitude spurs due to Doppler shifts. ECDLs can also be made from uncoated diodes, but the two competing cavities make it much harder to find a stable operating region, and the tuning is no longer continuous as it is with the grating-tuned, AR-coated version. You'd think that the grating feedback would help to prevent mode hops by sharply distinguishing the allowed modes, but in practice it doesn't really. Even fancy commercial ECDLs are highly sensitive to feedback, so budget for a two-stage free-space Faraday isolator in addition to the laser itself.

2.14.2 Injection Locking and MOPA

Small diode lasers have better mode characteristics than large ones, just as small signal transistors make better oscillators than power units. Larger lasers can have their facets AR coated to eliminate regeneration, turning them into amplifiers instead. They can be used as power amplifiers for the radiation from smaller ones, the so-called master oscillator–power amplifier (MOPA) approach. The main drawback of MOPA, besides its cost and complexity, is the large amount of spontaneous emission from the amplifier.

Alternatively, the power stage can be left to oscillate, but seeded by shining the master oscillator's output into it. Under the right conditions, the power stage will *injection lock* to the seed. (This is an old microwave trick and also helps a lot with OPOs.) Injection locking requires less pump power than MOPA, and the Fabry–Perot resonance of the power stage filters out most of the spontaneous emission, but it's just flakier. Its bad behavior arises from the coupling between the two resonators, which are sensitive to temperature, current, and the phases of the moon (microwave versions have amplifiers with good input–output isolation, a luxury we'll have more than one occasion to envy before this book is finished). MOPA seems to be a much more reliable approach.

2.14.3 Strong UHF Modulation

Rather than increase the temporal coherence of a laser, sometimes it's more helpful to destroy it. Mode hops can be eliminated by the use of *quenching*, by analogy with the superregenerative detectors of early radio. When gain was very expensive, positive (regenerative) feedback offered an appealing if unstable method of getting more gain

for less money. Superregens work by coupling the input into an oscillator circuit that is turned on and off at ultrasonic frequency by a second (*quench*) oscillator. The exponential buildup of the oscillations produces a waveform proportional to the size of the input, but many times larger, with the amplification and linearity controlled by the quench frequency (see Terman, it's beautiful). Lower quench rates give a logarithmic response.

Laser quenching isn't as pretty, but it's still useful. In a mode hopping laser, the situation is a bit more complicated, since the laser oscillation itself builds up very rapidly (1 ns or faster) and it's the mode hops we want to quench. Using large-signal UHF modulation to essentially turn the laser on and off at 300–500 MHz suppresses mode hopping completely, at the cost of enormously increased linewidth. Commercial ICs are available, or you can use an RF transistor in parallel with your diode laser. (Gallium nitride RF FETs are especially good for this.) This trick was widely used in magneto-optical storage applications and DVD players—there were even self-pulsating diode lasers that turned themselves on and off at UHF rates. (Note that the linewidth is still fairly small compared with a light bulb, so that interferometers built with these will have phase noise problems unless their path differences are extremely small.)

Optical Detection

You can't hit what you can't see.

—Walter Johnson (American baseball player)

3.1 INTRODUCTION

Electro-optical systems normally consist of an optical front end and an electronic and software back end, with an optical detector occupying the uncomfortable place of honor in between. The detector subsystem consists not only of the detector itself, but includes any baffles, coolers, windows, or optical filters associated with the detector, as well as amplifiers, electrical filters, and analog signal processing taking place in its vicinity. The optical front end is normally the most expensive part of the system, but is the only place where its photon throughput (and hence the maximum SNR) can be improved; the electronic back end is where the filtering, digitizing, and postprocessing occur, and is what produces the valuable output.

The guiding principle of detector system design is this: a photon, once lost, cannot be recovered. This includes those lost due to poor coatings, inefficient detectors, or poor matching of the optical characteristics of the incoming light to those of the detector, as well as those that are needlessly swamped in technical noise due to a poor choice of amplifier, load impedance, or circuit topology.

Once the measurement principle has been chosen, the choice of detector and the design of the detector subsystem are usually the most critical tasks in engineering a high performance electro-optical instrument; it is easy to get these badly wrong, and even serious errors are not always immediately obvious. Vigilant attention to every decibel there will be repaid with a sensitive, stable, repeatable measurement system. Decibel-chasing should not be limited to sensitivity alone, but should include stability as well; if efficiency is purchased at the price of instability, a measurement that was once merely slow may become impossible. Careful attention must be paid to many second-order sources of artifacts, such as ground loops, spectral response changes with temperature, etalon fringes, parametric effects such as memory in photoconductors, and nonlinear effects such as overload in photomultipliers or Debye shielding and lateral voltage drops in photodiodes. Achieving high stability is a somewhat subtle task and requires the cultivation of a certain healthy paranoia about unconsidered effects. Nevertheless, it is quite possible

for a measurement system to achieve stabilities of 10^{-5} to 10^{-6} in 1 hour, even at DC, if good design practices are followed.

Linearity is often the most important parameter after efficiency and stability. Many types of photodetector are extremely linear, as well as time invariant; this good performance can only be degraded by the succeeding circuitry. Thus another useful maxim is: if there's one operation in your signal processing strategy that has to be really accurate, do it right at the detector. Examples of this sort of operation are subtracting two signals, where very small phase shifts and amplitude imbalances matter a great deal, or amplifying very small signals in a noisy environment.

Detectors differ also in their resistance to adverse conditions. Photodiodes tend to be vulnerable to ultraviolet damage, photomultipliers to shock, APDs to overvoltage. (PDs are bulletproof by comparison, of course.)

3.2 PHOTODETECTION IN SEMICONDUCTORS

A semiconductor *PN junction* is the interface between regions of n-doping (excess electrons) and p-doping (excess holes). Electrical neutrality (zero E field) would require the extra electrons to stay in the N region and the extra holes in the P region. However, the electrons are in rapid motion, and their thermal diffusion flattens out the density gradient of the free carriers. The mismatch between the bound charge (ions) and free charge (carriers) causes an E field to form in the junction region, even at zero bias. The magnitude of E is just enough to cause a drift current equal and opposite to the diffusion currents. Absorption of light causes the formation of electron–hole pairs, which are pulled apart by the E field, yielding a current at the device terminals. Away from the junction, the E field is shielded out by the free charge. Thus an electron–hole pair generated there will usually recombine before it can be collected, which reduces the quantum efficiency of the device. Applying a reverse bias pulls the free carriers back from the junction, forming a *depletion region* with a large E field throughout. If the doping level in the depletion region is low, applying a reverse bias can cause a very large change in the depletion layer width (Figure 3.1a,b). This reduces device capacitance typically $7\times$ (by effectively separating the capacitor plates), and the extra depletion layer thickness improves quantum efficiency at long wavelengths, where many photons would otherwise be absorbed outside the depletion region. In an avalanche photodiode (Figure 3.1c), the doping profile is changed to make a separate high field region deep inside the device, in which electron multiplication takes place. (See Sze for more.)

3.3 SIGNAL-TO-NOISE RATIOS

There is a significant amount of confusion in the electro-optics field about how to calculate and quote signal-to-noise ratios, and power ratios in general. This arises from the fact that detectors are square-law devices; the electrical power emerging from the detector subsystem is proportional to the square of the optical power incident.

3.3.1 Square-Law Detectors

One way to look at this is that a photon of a given frequency ν carries an energy equal to $h\nu$, so that the optical power is proportional to the number of photons incident on the




Figure 3.1. Photodetection in semiconductor diode: (a) PIN diode, zero bias; (b) PIN diode, high bias (fully depleted); (c) avalanche photodiode; v and π are very low-doped n and p regions, respectively.

detector. In a quantum detector, each photon gives rise to an electron–hole pair, so that the electrical current is proportional to the photon flux. Electrical power is proportional to i^2 , however, so the electrical power goes as the square of the optical power. Since signal-to-noise ratio (SNR) is defined in terms of power, rather than amplitude, the signal-to-noise ratio on the electronic side is the square of that on the optical side. The same is true of thermal detectors, since $\Delta T \propto P_{\text{opt}}$ and $\Delta V \propto \Delta T$.

An optical signal of 10^6 photons per second has an RMS statistical noise of 10^3 photons in a 1 second DC measurement (0.5 Hz bandwidth); since the optical power is proportional to the photon flux, the signal-to-noise ratio on the optical side is 10^3 in 1 second. On the electrical side, however, the signal power delivered to a load resistor R is $(10^6 e)^2 R$, while the noise power is $(10^3 e)^2 R$, so that the signal-to-noise ratio on this side is 10^6 in the same 1 second measurement. The two ratios behave differently with bandwidth, as well. In a 0.01 second measurement ($B = 50$ Hz), the optical signal-to-noise ratio is $\sqrt{(0.01 \times 10^6)} = 100$; it goes as $B^{-1/2}$. The electrical SNR is $100^2 = 10^4$ in 0.01 second; it goes as B^{-1} .

The author uses the electrical signal-to-noise ratio exclusively, because he finds he makes many fewer blunders that way. It is much easier to convert the signal and shot noise into electrical terms than to convert load resistor Johnson noise, multiplication noise, recombination noise, and so forth into optical terms, because the one conversion is physical and the other is not. Test equipment such as spectrum analyzers, lock-in amplifiers, and A/D boards all work in volts and amps. By using the electrical signal-to-noise ratio, it is unnecessary to mentally extract square roots all the time in order to have good comparisons between different measurements. One drawback to this approach is that if your colleagues are used to quoting SNR in optical terms, they may feel that you're grandstanding by using the electrical SNR. The only defense against this charge is to state clearly and often which ratio is being quoted.

Aside: Where Does the Power Go? Square-law detection is a surprisingly deep subject, leading as it does to coherent detection and large dynamic ranges, but there is a

simpler aspect that people puzzle over too. It is this: If I send an optical signal of power P_O into a photodiode, the electrical power $P_E = R_L(R \cdot P_O)^2$, which is less than the optical power when $R \cdot P_O < h\nu$ and greater than P_O when $R \cdot P_O > h\nu$. What's going on? Am I really wasting almost all my optical power when I'm stuck with small signals?

We need to distinguish between *energy* and *information*. A solar cell, whose job is to turn optical into electrical energy, wastes a lot of the incident energy: even if its quantum efficiency is 1, its power efficiency is $eV_F/(h\nu)$. However, if shot noise dominates, the counting statistics and hence the information content are identical before and after detection, so no information is lost or gained, even though the SNR is squared. The root of the confusion here is that we get so accustomed to identifying electrical SNR with information carrying capacity (correctly, as we'll see in Section 17.11.1) that we tend to forget that it's really the detection statistics that are fundamental for measurements. This is another reason to stick with electrical SNR throughout.

3.3.2 Photons

The photon concept is central to the theory of quantum electrodynamics, the relativistic quantum theory of electromagnetic fields. The interaction of light with matter is quantized; a very weak beam of light, which spreads out over a large area, nevertheless gets absorbed one photon at a time, each one in a very small volume. This is easy to see with an image intensifier. The mathematics of quantum field theory are the most accurate way of predicting the interactions of light with matter that are known at present, and that accuracy is very impressive in many cases. It is clear that this represents one of the towering achievements of 20th century physics.

Now let's come down to Earth and talk about instrument design. Photons are a very important bookkeeping mechanism in calculating photocurrent shot noise, and they are a useful aid in keeping straight the operations of acousto-optic cells. Beyond that, *the photon is the most misleading concept in electro-optics*. The moment people start thinking of photons flying around from place to place, they start finding mares' nests and hens' teeth. For our purposes, it's only the interaction of light and detectors that's quantized—in every other way, light behaves as a wave, and obeys Maxwell's equations to absurdly high accuracy.[†]

3.4 DETECTOR FIGURES OF MERIT

Comparing detectors based on widely differing technologies, manufactured in different sizes, and requiring different circuitry is often difficult. Various figures of merit have been developed to make it easier; unfortunately, these frequently have only oblique connections with the issues facing designers. A healthy scepticism should be maintained about the value of a small difference in a figure of merit, and practical considerations kept firmly in mind.

[†]Nobody knows what a photon *is*, for one thing: see Willis E. Lamb, Jr., "Anti-photon," *Appl. Phys. B* **60**, 77–84 (1995), and the supplementary issue of *Optics & Photonics News* from October 2003, which was devoted to the subject. Each of the eminent contributors had a strong opinion about what a photon was, and no two of them agreed.

3.4.1 Quantum Efficiency

A *quantum detector* is one in which absorbed photons directly create free carriers: a photodiode, photoconductor, or photocathode. The quantum efficiency (QE) η of such a detector is the ratio of the number of photodetection events to the number of photons incident, before any amplification takes place. (We usually get one carrier pair per detection event.) It is the most basic measure of how good a particular quantum detector is—the detector analogue of the transmittance of the optical system. It determines our signal to shot noise ratio, which sets an upper bound on the SNR of the measurement.[†]

Detectors with gain can reduce the effects of circuit noise but are powerless to improve the shot noise (and indeed contribute significant excess noise of their own). Close attention to quantum efficiency is an excellent defense against being led astray in the design of a detector subsystem.

3.4.2 Responsivity

The responsivity of a detector is the ratio of the output current to the optical power input, and is quoted in amps per watt or cute substitutes such as milliamps per milliwatt. It is the most appropriate figure of merit to use when considering the biasing of the detector and the gains and dynamic ranges of subsequent amplifier stages, and also in determining how much optical power will be needed to overcome Johnson noise in the electronics. It is easily seen to be given by

$$\mathcal{R} = \frac{M\eta e}{h\nu}, \quad (3.1)$$

where η is the quantum efficiency and M is the multiplication gain (unity for photodiodes). For a detector of unit quantum efficiency and unit gain, the responsivity is the reciprocal of the photon energy in electron volts; it ranges from 1A/W at 1.24 μm to 0.32 A/W at 400 nm in the violet. The responsivity of real detectors is typically a strong function of wavelength; not only are there fewer photons per joule at short wavelengths, but the detectors themselves are selective.

Responsivity is a reasonable basis for comparison of detectors without gain, such as different types of photodiodes, but should be used with caution otherwise; for example, a PMT with lower photocathode quantum efficiency but higher dynode gain may exhibit a higher responsivity than a high QE device—perhaps yielding a better measurement in very low light, where Johnson noise in the load resistor may dominate, but certainly a worse one in brighter light, in which the amplified shot noise from the photocathode is the major noise source.

The term *responsivity* is reused for a rather different parameter in photoconductors: the change of terminal voltage produced by an incident optical signal, in V/W, with some combination of bias current and bias resistor.

3.4.3 Noise-Equivalent Power (NEP)

The most important single parameter of a detected signal is its signal-to-noise ratio. Because most commonly used optical detectors are very linear, their intrinsic noise is

[†]There are very special situations (“squeezed states”) in which this has to be qualified, but they aren’t of much practical use.

additive in nature: it remains constant regardless of the incident optical power (PMTs and APDs are exceptions).

Detectors differ in their intrinsic gains and readout mechanisms, so that it is inappropriate to compare the noise performance of different detectors solely on the basis of their output noise currents. It is common to convert the noise into a noise-equivalent (optical) power, NEP. The NEP is defined as the optical signal power required to achieve a SNR of 1 (0 dB) in the given bandwidth. The output noise is broadband, but not in general flat, so that the SNR of your measurement depends on how fast your light is modulated. Accordingly, the NEP is quoted in (optical) watts, at a certain wavelength λ , modulation frequency f (see Section 13.3), and bandwidth B (usually 1 Hz)—NEP(λ, f, B). NEP is somewhat awkward to use, because a good detector has a low NEP; thus its reciprocal, the *detectivity*, is often quoted instead.

In order to be able to compare the performance of detectors of different sizes, the NEP is sometimes normalized to a detector area of 1 cm², and then is quoted in units of W·cm⁻¹·Hz^{-1/2} (the somewhat peculiar unit is another example of optical vs. electrical power units—this one is optical). (Thermal detectors do not in general exhibit this area dependence, so their noise should not be normalized this way—see Section 3.10.7.)

Example 3.1: Silicon Photodiode NEP. Designing with visible and NIR photodiodes is easy, at least near DC, because there are really only three noise sources to worry about: Johnson noise from the load resistor, shot noise from leakage, and the shot noise of signal and background. As a concrete example of a noise-equivalent power calculation, consider a 3 mm diameter silicon PIN photodiode, Hamamatsu type S-1722, used with 820 nm diode laser illumination. This device has a quantum efficiency of about 60% at 820 nm, a room-temperature dark current of about 100 pA at 10 V of reverse bias, and a shunt impedance of 10¹⁰ ohms. Its responsivity is

$$\mathcal{R} = \frac{\eta e}{h\nu}, \quad (3.2)$$

which is 0.39 A/W at 800 nm. The two contributors to the dark NEP (of the diode alone) are the shot noise of the dark current and the Johnson noise of the shunt resistance R_p , so that the total current noise is

$$i_N^2 = \frac{4k_B T}{R_p} + 2ei_{\text{DC}}. \quad (3.3)$$

With these parameters, the second term dominates the first by a factor of nearly 20, and the result is $i_N = 5.7 \times 10^{-15}$ A/Hz^{1/2}. The shunt resistance is not in general linear, so it is not safe to assume that $R_p = V_{\text{bias}}/i_{\text{leak}}$, although it'll be the right order of magnitude. To get from here to the NEP, we just multiply by the energy per photon of $hc/\lambda = 1.51$ eV and divide by the quantum efficiency, since a notional signal photon has only 0.6 probability of being detected, so that the optical power required to equal the noise goes up by 1/0.6. The final result for this case is that the NEP is

$$\text{NEP} = \sqrt{2ei_{\text{DC}}} \frac{hc}{\lambda\eta} \approx 1.6 \times 10^{-14} \text{ W}\cdot\text{Hz}^{-1/2}. \quad (3.4)$$

With a load resistance of 500 k Ω , this device will be shot noise limited with a current of $2kT/(eR_L)$, or 100 nA, corresponding to an optical power of 260 nW. The shot noise

of the dark current is so small that it would take a 500 M Ω load resistor to make it dominate. This is very typical.

3.4.4 D^*

The most used figure of merit for infrared detectors is D^* , the *specific detectivity*, that is, detectivity normalized to unit area. More specifically,

$$D^* = \frac{\sqrt{A}}{\text{NEP}(f, \text{BW})}. \quad (3.5)$$

This choice of normalization takes account of the area dependence of the noise, the wavelength dependence of the responsivity, and the frequency dependence of the noise and the sensitivity. (Consider a photoconductive detector with significant $1/f$ noise below 1 kHz, and a 3 dB cutoff at 1 MHz due to carrier lifetime.) This allows meaningful comparisons of detector types that may not have exactly the same area, and accounts to some degree for the behavior of detectors with gain, such as APDs and PMTs, whose huge responsivities may mask the fact that their noise is multiplied at least as much as their signals.

To choose a detector technology, we first prepare a photon budget, which estimates how much light is available with how much noise, and see if we can meet our electrical SNR target. Lots of common sense and sanity checking are required here, to make sure all relevant noise sources are considered, including signal-dependent terms such as shot and multiplication noise. The maximum NEP in the measurement bandwidth is equal to the optical power divided by the square root of the target electrical SNR (optical vs. electrical again). Once the detector area A has been chosen (normally the minimum size required to collect all the signal photons), the minimum D^* required is given by

$$D_{\min}^* = \frac{\sqrt{A}}{\text{NEP}_{\max}}. \quad (3.6)$$

It must be emphasized that normalizing the response this way is intended to aid comparisons of different technologies and does not necessarily help the designer choose the right detector unit. In a background or Johnson noise limited system, if the signal beam can be focused down onto a 100 μm detector, then choosing a 5 mm one will result in a NEP needlessly multiplied by 50 for a given bandwidth, but D^* won't change. Thermal detectors' NEP is controlled by their thermal mass and thermal conductivity, so they often don't scale with area this way, making D^* useless.

D^* is not very useful in the visible, because it is so high that the detector's intrinsic noise is seldom a limitation. Visible light measurements are typically limited by the background and its noise, the shot noise of the signal, or the Johnson noise of the detector load resistor. In a poorly designed front end, amplifier noise may dominate all other sources, but this need not be. Chapter 18 discusses in detail how to optimize the signal-to-noise ratio of the detector subsystem. In brief, because visible-light photodiodes are such good current sources, the effects of Johnson noise can be reduced by increasing the value of the load resistance R_L . The signal power is proportional to R_L , whereas the Johnson noise power is constant, so the (electrical) SNR goes as R_L (there are circuit tricks to keep the bandwidth from vanishing in the process—see all of Chapter 18).

In the infrared, D^* appropriately includes the shot noise of the leakage current and of the 300 K background, shunt resistance Johnson noise, and lattice G-R noise. D^* is sometimes quoted in such a way as to include the shot noise of the nonthermal background (e.g., room lights), and even nonlinear effects such as multiplication noise in APDs and PMTs. While this is a legitimate way of quoting the noise performance of a measurement, it is unhelpful. Since it lumps all noise contributions into a single number, it totally obscures the questions most important during design: Which is largest, by how much, and what can be done about it? Furthermore, as discussed in the previous section, nonoptical noise sources are important, and translating all noise into optical terms is confusing and unphysical. Accordingly, all the D^* numbers quoted in this book are for detectors operating in the dark.

Another problem with D^* is that an inefficient detector with very low dark current may have a huge D^* but still be a very poor choice for actual use; for example, a 1 cm² detector with a quantum efficiency of 10^{-6} , but with a dark current of one electron per week, would have a D^* of about 1.6×10^{15} cm·Hz^{1/2}/W in the red, an apparently stunning performance, but actually it's useless. This illustration is somewhat whimsical, but nonetheless illustrates an important point: NEP and D^* are not the whole story.

Example 3.2: Indium Arsenide. Consider the case of a 2 mm diameter InAs photodiode (EG&G Judson type J12-5AP-R02M), operating over the wavelength range of 1.3–2.7 μm , at a temperature of 240 K and zero bias. This temperature was chosen to reduce the otherwise large response nonuniformity and to improve R_{sh} . The device has a long wavelength cutoff of about 3.5 μm . Its shunt resistance R_{sh} is 100 Ω , and the Johnson noise of the shunt resistance dominates the noise. Its quantum efficiency η is about 0.6 across this wavelength range, set mainly by the front surface reflection from the detector, which lacks an AR coating. The mean square noise current in a bandwidth B is given by

$$\langle i_N^2 \rangle = \frac{4kTB}{R_{\text{sh}}}, \quad (3.7)$$

where the thermal background has been neglected. The Johnson noise current is 11.5 pA/Hz^{1/2}, equal to the shot noise of a photocurrent of 400 μA (about 0.5 mW of optical power, far more than this detector would ever see). D^* is given by

$$D^* = \frac{\eta e}{h\nu} \sqrt{\frac{A}{\langle i_N^2 \rangle}}, \quad (3.8)$$

which at $\lambda = 2 \mu\text{m}$ is about 1.5×10^{10} cm·Hz^{1/2}/W.

Infrared detection is where the concept of D^* is really crucial, because here the intrinsic noise of the detector is far from negligible. To verify that the detector is Johnson noise limited, we will estimate the shot noise of the background.

We assume that the planar detector is surrounded by a hemispherical black body at a temperature T , and that background photons whose wavelength is shorter than 3.5 μm are detected with $\eta = 0.6$, while all others are lost. Since $h\nu \gg kT$, the thermal photons are uncorrelated, and the thermal photocurrent exhibits full shot noise (see Section 3.10.2).

The photon flux per unit area, between ν and $\nu + d\nu$, from this hemisphere is $M_{q\nu} d\nu$, where

$$M_{q\nu}(T) = \frac{2\pi\nu^2}{c^2(e^{h\nu/kT} - 1)}. \quad (3.9)$$

Since the exponential in the denominator exceeds 10^6 over the wavelength region of interest, the -1 can be neglected, and the result integrated analytically, giving the approximate photocurrent due to the thermal background,

$$\begin{aligned}
 I_{BG} &= \frac{2\pi\eta e}{c^2} \left[\frac{kT}{h} \right]^3 \int_{\frac{hv_0}{kT}}^{\infty} dx' x'^2 e^{-x'} \\
 &= \frac{2\pi\eta e}{c^2} \left[\frac{kT}{h} \right]^3 (x^2 + 2x + 2)e^{-x} \Big|, \tag{3.10}
 \end{aligned}$$

where $x = hv_0/kT$ and $v_0 = hc/3.5 \mu\text{m}$. After converting to current noise spectral density (see Section 3.10.2), this corresponds to a noise current i_N in 1 Hz of 0.6 pA, well below the Johnson noise. The limiting D^* if only this noise contributed would be $2.7 \times 10^{11} \text{ cm}\cdot\text{Hz}^{1/2}/\text{W}$.

Neglecting the 1 in the denominator causes less than a 1% error in the integrand whenever $hv/kT > \ln(100)$, which applies for $\lambda < 10.4 \mu\text{m}$ at 300 K, or $\lambda < 40.6 \mu\text{m}$ at 77 K.

3.4.5 Capacitance

The product of the load resistance and the detector capacitance usually dominates the high frequency performance of the detector system (transit time and carrier lifetime effects may also be important in some instances, especially in photoconductors). Bandwidth is often the limiting factor in how large the load resistor can be, so capacitance partly determines the sensitivity of the detector subsystem. The capacitance per unit area of the detector is thus another important figure of merit.

Reverse Bias. If the detector type allows use of reverse bias, this can reduce the capacitance by as much as 7–10 times in some devices. High speed applications, which require small load resistors to achieve small RC products, may benefit from detectors with intrinsic gain; their gain allows them to overcome the high Johnson noise current of the load. See Sections 18.4.4 and 18.5 for other ways to sidestep the problem.

In CCD and CID detectors, which are inherently integrating, capacitance per pixel should be large, to increase the full well capacity. The maximum signal level depends on how many electrons can be stored, so increasing the capacitance makes the statistical fluctuations and the readout noise a smaller fraction of the maximum signal level.

3.4.6 Spectral Response

A detector with a high peak quantum efficiency isn't much use if it is insensitive to some wavelength you care about. On the other hand, low quantum efficiency can be very helpful sometimes, as in solar blind UV photomultipliers and in GaP or other direct bandgap detectors, whose quantum efficiency drops extremely rapidly at wavelengths longer than cutoff. Spectral flatness means different things in different situations; a bolometer tends to be flat in terms of resistance change per watt, irrespective of the photon energy, whereas a photodiode is flat in terms of electrons per photon. Thermal band broadening gives detectors large values of $\partial\eta/\partial T$ near their long- λ cutoff.

3.4.7 Spatial Uniformity

Detectors are not perfectly uniform in sensitivity. Silicon photodiodes typically have 1–10% variation in η across their active areas, due to coating nonuniformity and the sheet resistance of the very thin top layer. Large-area diodes are usually worse, as one might expect. Fringes due to windows and nonuniform passivation layers can make this even worse. Beyond about 1 μm , you can even get etalon fringes from the back surface of the silicon (especially in CCDs, which have no spatial averaging). Nonuniformity can be a problem in situations such as interferometric detectors, where it can degrade the angular selectivity of the measurement by preventing spatial interference fringes from averaging to zero as they should, in position sensing applications, and in any measurement calling for absolute radiometric accuracy. Some detectors are much better than others; specifications may be available from the manufacturer, but there's no substitute for your own measurements. For the highest accuracy applications, homogenizing devices such as integrating spheres provide a good solution; they are of course better for white light sources than for lasers, due to speckle (see Section 5.7.11). With lasers, it's usually best to put the photodiode at the pupil, because small angular shifts generally cause much less sensitivity variation.

3.5 QUANTUM DETECTORS

3.5.1 Photodiodes and Their Relatives

Photodiodes are the most popular detectors for optical instruments. This is no accident; they come in a huge variety of types, sizes, and characteristics, their performance is excellent, and their cost is usually low.

A shallow PN junction is formed in a semiconductor. Light entering the chip is absorbed, creating electron–hole pairs. Provided the absorption occurs in or near the depletion region of the junction, the large electric field there (produced by the junction itself and by any externally applied bias) separates the carriers rapidly. Thus they do not recombine appreciably until the charge accumulation is sufficient to overcome the field in the junction, either through forward conduction (as in an open-circuited solar cell) or through Debye shielding in areas of high light intensity.

Ordinary PN junction devices are useful in low speed applications but tend to have very large capacitances, especially at zero bias. PIN diodes have a thick layer of very low-doped (*intrinsic*) semiconductor between the electrodes, which increases the depletion layer thickness, and so reduces the capacitance enormously, to the point where carrier transit time can become the speed limitation. Moderate reverse bias can cause the whole thickness (500 μm or so) of the device to be depleted.

Photodiodes tend to have constant, high quantum efficiency over broad bands; AR-coated silicon photodiodes easily reach 90% η over the visible, and even with no coating, often reach 60%. They roll off at long wavelengths as the semiconductor becomes transparent, so that many of the photons pass completely through the junction region, and at short wavelengths as light is absorbed too shallowly for the carriers to even reach the junction before recombining. So-called blue- and UV-enhanced diodes use very thin top layers to minimize this, at some cost in linearity. Another approach is to use a Schottky barrier instead of a PN junction; a nearly transparent top electrode of thin metal or transparent conductor such as indium tin oxide (ITO) can substitute for the top semiconductor layer.

The quantized nature of the light-to-electricity conversion ensures that photodiodes are extremely linear, which is one of their most important characteristics. Furthermore, photodiodes made of silicon are extremely good current sources when operated at zero or reverse bias; as we saw in Example 3.1, their intrinsic noise arises almost entirely from the shot noise of their leakage current, which is itself low. The low intrinsic noise and leakage are largely traceable to the fact that silicon has a relatively wide bandgap, nearly 50 times the thermal voltage kT/e at room temperature, so that thermally generated carriers and thermal radiation are not significant contributors to the noise in most cases. Leakage can be reduced even further by cooling, although that is rarely necessary in practice.

Nonlinearity in reverse-biased photodiodes comes mainly from excessive current density. High photocurrent densities (> 1 mA for a uniformly illuminated 3 mm diameter unit, for example) cause lateral voltage drops and intensity-dependent local increases in the carrier density. Both effects reduce the electric field in the junction, leading to enhanced recombination and slow drift. In other words, if you hit it too hard, it'll go nonlinear and slow down on you. A very small illuminated spot will exhibit this effect at surprisingly low photocurrents—microamps, even CW. See Section 3.5.4 for how much worse this gets with pulsed lasers. *Don't focus the beam down on the diode.*

Photodiodes used at zero bias are considerably less linear than reverse-biased ones, because lateral voltage drops in the thin epitaxial layer cause the diode to be locally forward biased, so that small forward currents flow and partly short out the photocurrent. Radiometrists resort to 20 mm diameter cells for photocurrents of 100 or 200 μA , whereas in reverse bias, the same detector is probably good to 10 mA. It takes a pretty big lateral voltage drop to overcome 5–20 V of reverse bias.

Silicon is the most common material used for photodiodes, both because it is a very good material and because silicon processing is a mature field. In general, silicon is best from 0.2 to 1 μm , and InGaAs from 1 to 1.7 μm (new InGaAs devices reach 2.6 μm). Germanium is also widely used for detectors out to 1.8 μm , but it is being eclipsed by InGaAs, which has been greatly developed for fiber optic communications at 1.3 and 1.55 μm . Beyond there, infrared devices become more exotic and difficult to use. High bandgap semiconductors such as GaP and GaN photodiodes ought to be the best kind in the ultraviolet, but don't seem to be as good as Si in practice. The one exception is silicon carbide, which has the advantage of being nearly totally solar blind; SiC devices have negligible leakage, work from 200 to 400 nm, and can reach $\eta = 0.72$ at 270 nm, which is remarkably good.

3.5.2 Shunt Resistance

Low light applications, where you want to use big diodes and huge feedback resistors (100 M Ω to 10 G Ω) are one place where cooling silicon photodiodes helps, for an interesting reason. To reduce leakage, you run the photodiode at exactly zero bias. Here's the subtlety: the leakage goes to 0, all right, but the shunt resistance deteriorates very badly at zero bias, because of the diode equation. In Section 14.6.1, we'll see that a diode's zero-bias shunt resistance r_0 decreases dramatically at high temperatures and low bandgaps, which is why IR detectors are worse than visible ones, and why cooling them helps. That resistance is a real resistance with real Johnson noise current, and furthermore if it gets too small, it forces the transimpedance amplifier to run at a huge noise gain

(see Section 13.1), which multiplies the amplifier's noise voltage, offset, and drift—bad news for measurement accuracy. Fortunately, we're very rarely in this situation.

Aside: The Zero-Bias Heresy. It is a sad fact that almost all the photodetector circuits published in the history of the world have been designed by circuits people who weren't also optics people. The author has no statistics to back this up, but he feels that it must be so, because almost all show the photodiode being operated at zero bias, often with great care being exerted to make the bias exactly zero. This will reduce the dark current through the photodiode, all right, but that isn't the problem we need to solve (see Section 18.2.1). Photodiode dark current is almost never the limiting factor in a visible or near-IR measurement. Fixing this nonproblem costs you a factor of 5–7× in bandwidth (or the same factor in high frequency SNR), as well as destroying the large-signal linearity, which makes it an expensive blunder. *Don't do it.*

3.5.3 Speed

For the highest speed applications, such as 40 Gb/s optical Ethernet, the speed of photodiodes becomes a serious issue. There are two effects that limit speed: transit time and RC delays. The transit time is how long it takes a carrier to reach the terminals. It can be reduced by making the diode thin, and maximizing the field in the junction by using as high a bias as possible and making the diffusion-dominated p and n layers extremely thin. This of course tends to increase the capacitance, so fast photodiodes tend to be very small. Another problem is that the QE suffers, because at the 850 nm wavelength used in fiber LANs, the absorption depth in the silicon is several microns; thus narrow-gap III–V materials such as InP are commonly used in fast NIR applications, which costs extra. One approach to fixing this is to use transverse devices, where a very small detector is coupled to a guided wave—the light travels perpendicular to the current flow, so the light path can be long and the current path short.

Aside: Plastic Packages. It appears that a few types of plastic-packaged photodiodes are susceptible to QE changes due to triboelectric charging of their plastic packages. The effect usually goes away with a bit of surface leakage (e.g., by breathing on the package) or slight heating. If your accuracy requirements are high, you may want to test for this, or stick with metal-can devices.

3.5.4 Photodiodes and Pulses

Now that femtosecond lasers are so popular, it's worth giving a bit of thought to the plight of the poor photodiode used to detect the pulses. Of course, the photodiode doesn't have the glamour job—or the easy one either. Consider a 100 fs laser with a 100 kHz pulse repetition rate, and an average detected power of 2.5 mW. A 100 fs pulse of 250 nJ (assuming it's a Ti:sapphire around 800 nm) produces a peak photocurrent of 160,000 A, and even though the diode can't possibly respond that rapidly, one wouldn't expect it to be terribly linear under such treatment, as indeed it isn't. One very good solution is to use a small (25 mm) integrating sphere on the diode (see Section 5.7.7). If it has a reflectivity of 97%, then incoming photons will typically bounce around over about a meter's distance before being absorbed, which will broaden the pulse to 2.5 nanoseconds or so. This is a better match to the capabilities of photodiodes. If your detector's diameter

is larger than the input aperture's, and larger than about 1/8 of the sphere's diameter, you'll collect most of the photons. Make sure the first bounce of the beam in the sphere isn't in the field of view of the detector, or you may still have nonlinearity problems. (Don't ablate the paint.)

3.5.5 Phototransistors

A phototransistor looks like a silicon bipolar transistor with a photodiode connected between base and collector. They have gain and are widely available at low cost, which about exhausts their virtues. Although their gain makes smallish photocurrents conveniently large, they are slow, leaky, nonlinear, and very noisy. They come only in very small sizes. Photodarlingtons have an additional gain stage and are even slower. These devices are classical examples of why amplification is not always useful; except for the very lowest performance applications, avoid them at all costs.

3.5.6 Prepackaged Combinations of Photodiodes with Amplifiers and Digitizers

Several detector and op amp manufacturers build packaged combinations of photodiodes with transimpedance amplifiers, current to frequency converters, "current amplifiers" (Hamamatsu), or even $\Delta-\Sigma$ type A/D converters. These devices are intended for people who are not comfortable designing detector circuits, and you pay in cost and performance for the convenience of using them despite frequent claims to the contrary by their manufacturers. (Their data sheets very seldom mention noise performance, which is a very bad sign.) The performance problems mostly arise from their use of fixed value internal feedback resistors (often 1 M Ω), and failure to put in a cascode transistor to reduce junction capacitance effects or a tee network to decrease the second-stage noise contribution (see Section 18.4.12 for more details). They can be useful where there is lots of light available and speed is not a limitation, or where skilled engineering time is very limited; people commit much worse blunders than these all the time, and using these devices at least limits the damage. These devices may improve in quality and cost effectiveness in the future; at present, however, designers who are willing to do a bit more work can usually achieve much better performance at a considerably lower cost. The exception to this rule is some packaged APD/GaAs FET amplifier combinations, which provide a level of speed and noise performance that is not trivial to reproduce.

3.5.7 Split Detectors

It is often handy to be able to measure the position of a beam on a detector. If image sensing is not absolutely required, the best way to do this is by using split detectors (bi-cells and quadrant cells) or lateral effect devices. A split detector is just that: two or more detectors spaced very closely, often arranged as sectors of a circular disc. Each detector has at least one lead of its own (often all the anodes or all the cathodes are common, to save leads and ease fabrication). By combining the separate photocurrents in various ways, such as subtracting them or computing their ratio, small changes in the beam position can be measured very accurately (often to 0.1 nm or better near null, at least at AC). Because their geometry is defined lithographically, they are

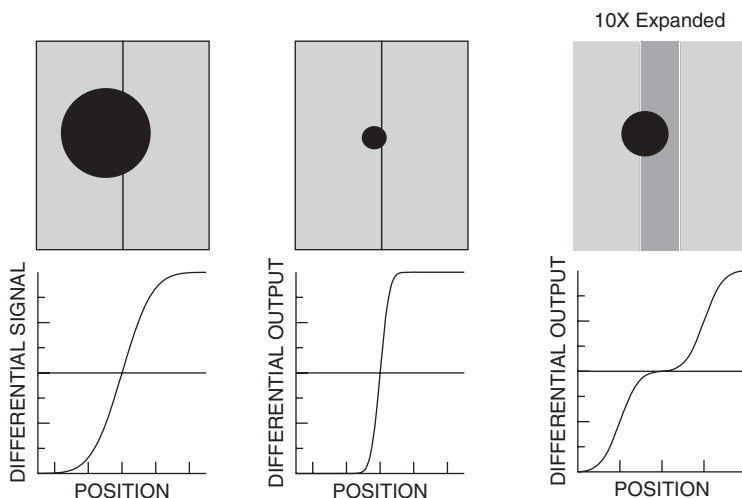


Figure 3.2. Split photodiode, showing the dependence of output signal on beam diameter and position.

very stable, and because each segment is operated as an independent photodiode, split detectors have the same virtues of linearity, low noise, and high efficiency we expect from single photodiodes. Split detectors are currently available in Si, Ge, InSb, and HgCdTe.

The main problems with split detectors are that some of the light is lost by falling into the kerf between the cells, and that the position sensitivity depends reciprocally on the beam diameter (See Figure 3.2). This is physically obvious, since a full-scale change results as the beam moves from being entirely on one side to entirely on the other. If the position information is obtained by subtraction, the sensitivity will depend on the optical power as well. Using analog division instead of subtraction helps.

Another approach is to use the cells in the open-circuit photovoltaic mode, where the photocurrent is allowed to forward bias the cell (as in a solar cell), and subtract the open-circuit voltages. These voltages depend logarithmically on the photocurrent, so when they are subtracted, a ratiometric measurement results. The circuit responds slowly, but this trick is good for alignment sensors (see Section 12.9.11).

3.5.8 Lateral Effect Cells

Another position sensing device, which avoids the problems of split detectors, is the lateral effect cell. Both 1 D and 2 D devices are available, in Si, Ge, InSb, and InAs. They are single large photodiodes that use a thin, highly resistive layer for the top electrode of the cell, each end of which has its own lead. The output leads are connected to low impedance points (such as op amp summing junctions). The light beam appears as a current source located somewhere on the surface, so that the photocurrent divides itself between the output pins in proportion to the conductance of each path. Because the conductance depends on the distance from the light beam to the output pin, the ratio of the currents in each pin gives the location of the light source. Because the cell surface is uniform, and the current division linear (i.e., two or more light sources shining on the

same cell produce the sum of the outputs each would produce by itself), the position signal is much less sensitive to beam diameter, until the edges of the beam approach the edges of the cell or the current density or lateral voltage drops get so high that response nonlinearity sets in. Two-dimensional lateral effect cells come in two varieties: *pincushion*, where the readouts are in the corners, and *linear*, where the anode is split in one axis and the cathode in the other, so the anode sheet gives x and the cathode y . The linear ones give position information that is accurate to 1% or so, with the best ones achieving 0.1% over 80% of their aperture. Lateral effect cells tend to be much slower than split detectors, since they have a fairly large series resistance (1–200 k Ω), and also more difficult to use. The difficulty is that, because of the series resistance, even moderate photocurrents can cause lateral voltage drops large enough to locally forward bias the junction, leading to serious nonlinearity. The simplest way to avoid this problem is to reverse bias the junction. There are also some inferior devices where the top electrode has four long contacts, along the sides of a square, and the bottom has only one. These superficially resemble linear lateral effect devices but differ in one crucial respect: the two axes compete for current. If the illuminated spot is near the $-x$ electrode, that electrode will suck up almost all the photocurrent, leaving little to drive the y axis, which therefore is rendered nearly indeterminate by noise and drifts.

The low-light performance of lateral effect cells is relatively poor, because the resistance of the silicon layer appears in parallel with the outputs. It thus contributes a large amount of Johnson noise current, with a noise source connected between the two outputs of each axis. This is an obnoxious problem, because although these noise currents of course sum to zero, they must be subtracted or divided instead in order to make a position measurement; this causes them to augment instead of canceling. Unless the photocurrent is sufficient to drop $2kT/e$ (50 mV at room temperature) across the cell resistance, the measurement will be Johnson noise limited. It does not help much to increase the load resistance, since the differential signal will always be loaded by the cell resistance, so that the effective load for differential measurements will be the parallel combination of the cell resistance and the external load. In general the lateral effect cell is a fairly noisy, kilohertz-range device for sensing beam positions, but it is relatively immune to beam-size effects that can plague split-cell measurements.

3.5.9 Position Sensing Detector Pathologies

Split detectors are more sensitive to etalon fringes than single element ones. A smooth beam profile on a single element detector does a much better job of preserving the orthogonality of misaligned beams, whose fringe patterns average to zero over the surface. In a split cell, there's a great big cliff in the middle of the fringe pattern, so a small phase shift with temperature can cause an overwhelming amount of drift—all the more so since the desired signal is a small difference between two large currents. Position-sensitive detectors are also vulnerable to nonorthogonalities of the measurement directions caused by sensitivity variations across the detector.

3.5.10 Other Position Sensing Detectors

For position sensing applications, lower cost solutions such as a few discrete diodes plus a shadow mask should not be overlooked (Figure 3.3). For many purposes, they can be very useful and are dramatically cheaper than the combination of a lens plus a position

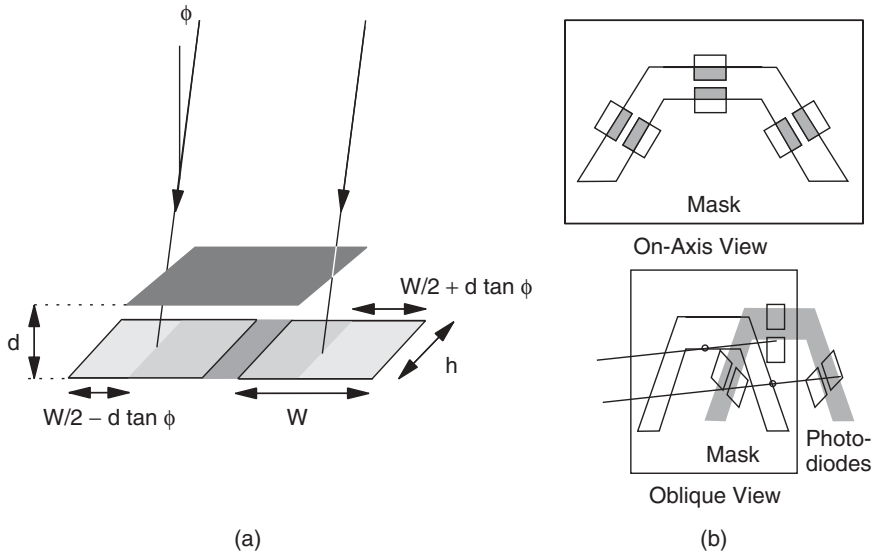


Figure 3.3. Three-dimensional shadow mask detector: (a) schematic and (b) drawing of actual device.

sensing diode. When using mid- and far-infrared detectors, where split detectors and lateral effect cells are harder to get, this may be the only choice feasible.

Example 3.3: Two Solar Cells and a Mask Versus a Bi-cell. As an example of a cheap position sensing detector, consider two rectangular solar cells of width W and height h , plus a thin mask at a height d , covering the right half of the left cell and the left half of the right cell, as shown in Figure 3.3. Uniform illumination of intensity I comes in at an incidence angle ϕ , illuminating $W/2 + d \tan \phi$ of cell 1, and $W/2 - d \tan \phi$ of cell 2. Obviously, if the angle is too large, one diode will be completely covered, and the other one completely illuminated. For angles smaller than this, if the two photocurrents are subtracted, the result is

$$i_- = 2I\mathcal{R} dh \sin \phi, \quad (3.11)$$

where I is the power per unit area, measured in a plane normal to the incoming beam. If the difference is normalized by dividing by the sum of the two currents, the result is

$$i_{\text{norm}} = \frac{2d \tan \phi}{Wh}, \quad (3.12)$$

which is plotted in Figure 3.4, together with the angular uncertainty due to shot noise alone. You can also divide instead of subtracting, which gives

$$\frac{i_2}{i_1} = \frac{W + 2d \tan \phi}{W - 2d \tan \phi}. \quad (3.13)$$

It is apparent that the sensitivity of the shadow mask/multiple element detector can be adjusted over a wide range by changing d , a desirable property. If the light comes




Figure 3.4. Output and RMS angular error of a shadow mask sensor, with 5 mm square detectors and a 3 mm mask spacing. $P_{opt} = 10 \text{ mW/cm}^2$.

from a wide spectrum of angles, it is sensible to integrate Eq. (3.11) over angle. If the differential characteristic is not required, a single photodiode with a shadow mask can be used instead, at an even greater cost saving.

If two axes are required, two such single-ended detectors can be used, with the edges of their shadow masks mutually perpendicular. If the differential character is important, then three or four detectors can be combined. If four are used, the 1D expressions are approximately correct, except for the additional effect of obliquity in the other axis. For the three-detector case, as shown in Figure 3.4, the X and Y directional signals are given by

$$X = \frac{i_3 - i_1}{i_1 + i_2 + i_3} \tag{3.14}$$

and

$$Y = \frac{i_3 + i_1 - 2i_2}{i_1 + i_2 + i_3}, \tag{3.15}$$

respectively. This saves one photodiode, while preserving the differential property: light coming in exactly on axis gives $X = Y = 0$; the third dimension comes from intensity or (for a nearby point source) parallax, in which case the simple expressions given have to be modified. The noise performance of shadow mask detectors, like that of quadrant cells, is excellent. In a 1 Hz bandwidth, the shot noise limited RMS angular uncertainty is

$$\left\langle \frac{\delta\theta}{\Delta\theta_{pp}} \right\rangle \approx \sqrt{\frac{2e}{i_0}}, \tag{3.16}$$

where i_0 is the response of one unobscured detector to the incident optical intensity. A HeNe laser beam of 1 mW per detector will give $i_0 \approx 300 \mu\text{A}$. Assuming a mask spacing

of half the detector width, which results in a full scale range of $\pm\pi/4$ radians, the 1 Hz RMS angular uncertainty is around 50 nanoradians, or 0.01 arc second.

The accuracy attainable by this technique is limited mainly by nonideal beam characteristics, such as amplitude nonuniformity, multiple scattering between mask and detectors, and diffraction of the beam at the edge of the mask. †

3.5.11 Infrared Photodiodes

Photoelectric detectors for the mid- and far-infrared region also exist, based on compound semiconductors such as indium antimonide (InSb), indium arsenide (InAs), platinum silicide (PtSi), and mercury cadmium telluride (HgCdTe, sometimes written MCT, and pronounced “mercadtell”). Their characteristics are not as good as those of silicon and InGaAs photodiodes, and they are much more expensive. Compound semiconductor materials are more difficult to process, because small errors of stoichiometry appear as huge dopant densities, and because their small markets dictate that their processing is much less developed than that for silicon (GaAs and InGaAs are the exceptions). In the mid-IR, the standard detector at present is HgCdTe.

Of the near-IR devices, Ge was king for a long time. It has the advantage of being an element, so the stoichiometry is not a problem, but its poor leakage performance means that it requires cooling in situations where compound semiconductor devices such as InGaAs do not. In the past, InGaAs detectors were too expensive for widespread use, but now that demand for detectors in fiber optics applications has funded their further development, costs have fallen to a few dollars for a usable InGaAs detector at 1.5 μm .

Mid- and long-wavelength detectors, such as HgCdTe, PbS, PbSe, PtSi, and InSb, frequently require cryogenic cooling, which is inconvenient and expensive (\$2500 for a single cooled detector element). InAs, with response out to 3.5 μm , is often useful with only thermoelectric (-20 to -60°C) cooling.

Unlike silicon detectors, the shunt resistance of an infrared diode can be very low, as in Example 3.2; the Johnson noise of this low resistance is the dominant additive noise source in IR devices that are not cryogenically cooled. Because of the low shunt resistance, carriers generated far from the electrodes are sometimes lost, leading to very large spatial nonuniformities of response (Figure 3.5 shows a 4 \times changes in η with position).

3.5.12 Quantum Well Infrared Photodiodes

A promising new mid- and far-IR detector technology is based on quantum wells. These are basically the famous square box potential well, whose energy levels can be tailored by adjusting the depth and width of the well. Accordingly, the bandgap of the detector can be as narrow as desired, without having to deal with the poor properties of narrow-gap semiconductors. Thermionic emission is still a problem with these devices, so they must be cooled, but their D^* can be as high as 10^{12} at 9 μm , a remarkable performance. They can also be made in arrays. We’ll see more of these in the future; the technology is driven by space-borne and military sensors.

†Do the diffraction ripples at the leading and trailing edges of the shadow cause ripples in the XYZ outputs versus angular position? Why or why not?

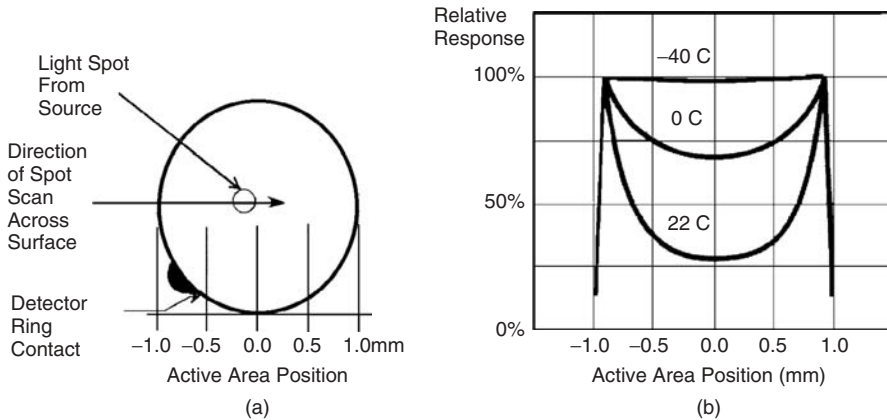


Figure 3.5. Response nonuniformity and shunt resistance of a commercial InAs photodiode (EG&G Judson J12-5AP-R02M) versus temperature. High lateral and low shunt resistance leads to poor uniformity at 300 K. (Courtesy of EG&G Judson Inc.)

3.6 QUANTUM DETECTORS WITH GAIN

3.6.1 Photomultipliers

At wavelengths shorter than $2 \mu\text{m}$, thermally generated photons are very rare, so in principle a detector should be limited only by signal photon statistics (shot noise). However, Johnson noise in load resistors, amplifier input noise, and other circuit noise typically dominates shot noise at low photocurrents. That's where photomultiplier tubes (PMTs) come in. PMTs use *electron multiplication*, which we will see more of later, to amplify the generated photoelectrons before they become mixed with the noise currents of the electronics. Electron multiplication occurs when an electron hitting a surface causes it to emit more than one *secondary electron*. The way this works in a PMT is as follows: a photocathode is exposed to incoming light. By the photoelectric effect, some fraction of the incoming photons cause the photocathode to emit photoelectrons from its surface. These photoelectrons are electrostatically accelerated and focused onto another electrode, the *first dynode*, which is coated with a material selected for high secondary electron yield. The secondaries from the first dynode are accelerated and focused onto the second dynode, and so on for 5 to 14 stages, before finally being collected by the anode. By the end, each photoelectron has become a pulse 300 ps to several nanoseconds long, containing perhaps 10^5 to 10^7 electrons, which is easily detectable above the Johnson noise.

Aside: Electron Affinity. When a photoelectron is generated inside a photocathode or dynode, it must escape into the vacuum before it's any use to us, and to do so it has to overcome a potential barrier. The height of this barrier is a material property called the work function W . The photoelectron is continually losing energy through collisions, and the requirement that it arrive at the surface with at least W limits the photoelectron yield. The work function is made up of two parts, the classical image potential, which is the work required to separate the electron from its image charge in the metal surface,

and the electron affinity.[†] The image potential is always positive, but negative electron affinity (NEA) materials exist, and their lower work function leads to improved electron yield, at the expense of slower response.

3.6.2 PMT Circuit Considerations

This elegant scheme requires a certain amount of circuit support: a power supply of from -500 to -2000 volts, with taps for the photocathode and all the dynodes. This is usually provided by a powerful (1–2 W) high voltage supply and a multitap voltage divider string made of high value (100–300 k Ω) resistors. The high supply power is almost all dissipated in the resistors, which sets a practical lower limit on their values. The exact bias voltages are not terribly critical (unlike APDs, for example).

Because of the high electron gain, the last few dynodes need a fair amount of bias current, which is supplied only poorly by the resistor string. Bypass capacitors or Zener diodes on the last couple of stages are a help, but nevertheless the nonlinearity of most photomultiplier systems at high light intensity is dominated by voltage drops in the dynode bias string. It is often convenient to change the supply voltage to vary the gain, and having zeners on in the string makes this hard, because the voltage distribution on the dynodes will change as the supply voltage is altered. Since the linearity depends significantly on this distribution, zeners reduce the flexibility of the system.

The Cockroft–Walton (C-W) generator is a many-section voltage multiplier based on a diode–capacitor ladder structure. An N -stage C-W generator produces N nearly equally spaced voltage taps, making it a natural fit for biasing PMT dynodes. Some photomultiplier modules now include C-W multipliers rather than voltage dividers.[‡] C-Ws have the useful property that the lower taps have much lower impedance than the high voltage end, a good match for the needs of PMTs. Besides compactness and low power, this leads to enormously improved linearity in Cockroft–Walton devices, so that the resistor scheme is obsolescent for most purposes. Watch out for the modules with preamps built in—they all cut off at 20 kHz to avoid the power supply ripple.

The most appropriate uses for resistor biasing nowadays are in applications where linearity is not vital but power supply ripple is extremely objectionable, or in devices where the photocathode cannot be run at a large negative bias, and DC coupling is not needed. Running the photocathode near ground means that the anode and last dynodes must be run at a high positive voltage, which significantly reduces the advantages of the Cockroft–Walton technique, since the high current electrodes are at the high impedance end of the supply.

Applications requiring a grounded photocathode include scintillation detectors, where the scintillator is an ionic crystal such as NaI, which must be operated at ground for safety reasons. If the photocathode is not grounded, the large potential gradient across the glass envelope of the PMT in an end-on tube can lead to electrophoretic motion of ions toward the photocathode, which destroys the tube by photocathode corrosion. Side-looking PMTs with opaque photocathodes (electrons are emitted from the same side the light hits) are immune to this since the photocathode doesn't touch the envelope.

[†]For insulators the electron affinity may be less than this, because an added electron goes into the conduction band, whereas a photoelectron comes from the valence band.

[‡]This idea took awhile to catch on—it was first published in 1960 (R. P. Rufer, Battery powered converter runs multiplier phototube. *Electronics* **33**(28), 51 (1960)).

The noise of PMTs is dominated by thermionic emission from the photocathode, leading to dark current spikes, and by variations in the dynode gain (especially at the first dynode), which leads to multiplicative noise. The average number of secondary electrons is only 5 or so, although the NEA material GaP(Cs) can achieve 20–50. As we'd expect from a low yield random process, the size of the pulse from a single photon event varies within a 1σ range of $\pm\sqrt{5}/5 \approx \pm 45\%$, although with larger photocurrents these error bounds are greatly reduced through averaging. PMTs require very high insulation resistances between their internal elements, but use volatile metals such as cesium and antimony, which are prone to migrate at high temperatures; thus PMTs cannot be baked out as thoroughly as most vacuum systems. There is always some residual gas (perhaps 10^{-6} torr, mainly water), which leads to artifacts known as *ion events*. An ion event takes place when a positive ion is generated in the residual gas inside the PMT near the photocathode. Because of its positive charge, it accelerates and hits the photocathode hard enough to knock loose many electrons. This large pulse is amplified through the dynode chain, producing a very large current pulse at the anode. Ion events are even more common in old or poor quality tubes, or those operated near radioactive sources.

Related to ion events are afterpulses, which are secondary pulses that sometimes occur, usually 20–100 ns after a photon is detected. These arise from photoemission inside the PMT due to electron impact on a surface, or a generated ion that hits a dynode instead of the photocathode. Afterpulses are a problem with fast photon counting setups; putting in 100 ns of dead time after each photocount will more or less cure the problem. This dead time reduces the integration period at higher light levels, so it has to be corrected for (see Section 7.3.1).

PMTs get old and wear out, at a rate largely controlled by the anode current. They exhibit strong (several percent) effects due to warmup, intensity and voltage hysteresis, and other historical causes. Under good conditions, over its life a PMT can produce an integrated anode charge of a few hundred coulombs per square centimeter of photocathode. In long-term dark storage, or in low current applications, the lifetime approaches a constant value of a few years, limited by helium diffusion through the tube envelope and by surface changes inside. Tubes with soda lime glass envelopes are the most vulnerable; high partial pressures of helium can kill one of those in an hour.

PMTs can be quite fast; ordinary ones have rise and fall times of a few tens of nanoseconds and the fastest are around 250 ps, with timing uncertainties of 40 ps or thereabouts. Fall times are rather slower than rise times. In a high gain PMT, a single photon can produce 10^7 output electrons, which in an 8 ns wide pulse amounts to 200 μA . Such a large current can produce 10 mV signals even across a 50 Ω load, making it feasible to count individual photons at high speed. The pulses are repeatable enough in height that simple thresholding can easily distinguish a pulse corresponding to a single detected photon from noise, and from multiphoton or ion events. Such *photon counting* is an important application of PMTs, and photon rates of up to 30 MHz can be accommodated with good accuracy with commercial gear (200 MHz with special hardware). The combination of high gain and low dark count rate makes PMTs uniquely suited to photon counting.

In photon counting, the photoelectrons arrive more or less one at a time. PMTs can also be used in analog mode, where the average anode current is measured in much the same way as with a photodiode. Of the two, photon counting is better behaved. This is mainly because the gain of a PMT depends on everything you can think of. In the analog

mode, this directly affects the signal level, whereas in photon counting it merely changes the pulse height, without making a pulse significantly more or less likely to be detected.

A photon counting PMT has a sensitivity similar to that of a cooled CCD. It has no spatial resolution, but on the other hand you don't have to wait for the end of the integration time to see the data.

The amplification mechanism of PMTs is very clever and effective, and their optical performance has recently improved by a factor of nearly 2. Conventional alkali PMTs tend to be narrowband, peaking strongly in the violet. PMTs are available in large sizes, up to about 300 mm in stock devices, and up to 600 mm in custom devices. The combination of huge area and low dark count rates is unique to PMTs.

Choosing a Photocathode Material. The choice of photocathode material depends on the application; they are typically made of a mixture of antimony with alkali metals or of a direct bandgap compound semiconductor with traces of cesium. Infrared units are available too. The classic Ag-O-Cu S-1 material reaches $1.1 \mu\text{m}$, but has extremely low quantum efficiency (0.01–1%) and has a lot of dark current, often requiring cryogenic cooling.

NEA photocathodes work further into the IR than alkali metal ones, but have a much slower response (1 ns rather than < 50 ps), making them less suitable for exotic applications such as streak cameras. The best InGaAs NEA photocathodes reach $1.7 \mu\text{m}$, with quantum efficiencies of around 20%, and the best enhanced alkali and GaAsP ones achieve about 45–50% peak quantum efficiency in the blue and near-UV.

How to Kill a PMT. Photomultipliers can be destroyed by exposure to daylight when powered, and their dark current can be dramatically increased by such exposure even when unpowered; several days of powered operation may be required to bring it back to normal.

Making Accurate Measurements. The photon counting mode is pretty trouble-free, but it is not trivial to make accurate analog measurements of absolute light intensity. For instance, the total gain typically varies $\pm 10\%$ with position on the photocathode. Accuracies of around 1% result from ordinary care, but by taking sufficient pains to control or calibrate small effects, accuracies and repeatabilities of 0.1% can be achieved.

The gain of a PMT is affected by many internal and external effects, including shock, age, dynode fatigue due to high current operation, stray fields, and even changes in geometry caused by gravity or acceleration. The efficiency and speed with which the first dynode collects the primary photoelectrons depends on position, leading to variations of ~ 10 –20% in sensitivity and a few percent in delay. Due to Fresnel reflection at the photocathode surface, the quantum efficiency also depends on incidence angle and (off-normal incidence) on polarization. Besides these parabolic-looking variations, the sensitivity of some PC types shows ripples with position, which get much worse at longer wavelengths, so it's often helpful to put diffusers in front of PMTs.

The gain of some tubes can be reduced 40% by a 1 gauss magnetic field (the Earth's field is about 0.5 gauss), although others can work in 1 kilogauss. Mu-metal shields improve this greatly, but mu-metal's shielding properties are easily destroyed by shock or bending. The lower energy electrons between the photocathode and first dynode are most susceptible to magnetic steering, so mount the tube so the shield sticks out about one diameter in front of the photocathode, to allow the fringing fields space to die off.

Static charges on the envelope or nearby grounded objects can cause electrons to strike the tube envelope, generating spurious light. Photoemission from the gas and from dynode surfaces, as well as Čerenkov light from cosmic rays and radioactive decay, cause spurious counts. Light from these sources can be guided by the glass tube envelope directly to the photocathode, where it will cause spurious counts. Graphite paint (DAG) applied to the envelope and kept at cathode potential (via a 10 M Ω high voltage resistor for safety) eliminates the light guiding and provides an electrostatic shield. High electric fields inside the envelope can cause scintillation as well, so use DAG and really good insulation (e.g., 4 mm of high quality silicone rubber, not 10 layers of PVC tape). Taking care here can reduce the dark counts by a factor of 10. High humidity is a disaster for PMT performance due to external leakage currents.

Being vacuum tubes, PMTs are easily destroyed by shock and may be microphonic in high vibration environments; also, of course, they are expensive. PMT manufacturers provide 200-odd page manuals on how to apply PMTs in measurements, and these are full of lore. If you need PMTs, get a few of these application manuals.

3.6.3 Avalanche Photodiodes (APDs)

When a high electric field is applied to a semiconductor, free carriers can acquire enough energy to excite other carriers through *impact ionization*. These newly generated carriers can themselves create others, so that a chain reaction or *avalanche* results. At higher fields still, carriers can be generated spontaneously by the field, and breakdown occurs. When this mechanism is used to multiply the photocurrent in a photodiode, the result is an APD. Reasonably stable multiplication gains (M) of up to 100 or so are possible by controlling the bias voltage carefully at a value of about 90% of the breakdown voltage, and compensating its strong temperature dependence.

Holes and electrons have different values of the ionization coefficient α (a normalized cross section). One might think that the performance would be best when both contribute equally, but in fact that's the worst case. All the carriers have the same speed, so in a pure electron avalanche, all the secondary electrons arrive at the same time as the primary photoelectron, and the holes are spread out by the transit time τ . In a bipolar avalanche, the holes cause ionizations, so that the avalanche spreads out in both directions and bounces back and forth until it dies away due to statistical fluctuations in the rates. This makes it become very slow and very noisy as the gain increases. The figure of merit for this is $k = \alpha_h/\alpha_e$, the ratio of the ionization cross sections of the less ionizing species (holes) to the more ionizing (electrons). The situation to avoid is $k \approx 1$. In silicon, k is very small, because holes cause almost no impact ionization, but in low bandgap materials like InGaAs, k is around 0.3 at low voltage, rising to nearly 1 at high voltage. Heterostructure APDs exist, in which the detection and multiplication are done in different semiconductors; this gives the best of both worlds at the price of complexity and expense. The noise is always worse at high gain; the noise power tends to increase as M^{2+m} , where the noise exponent m is 0.3–1.0, so that the SNR goes down as $M^{-0.3}$ to M^{-1} . The exact exponent is device dependent, and manufacturer's specifications should be consulted—if you're lucky enough to find a data sheet that has that level of detail. (Optical detector manufacturers should be ashamed of themselves for the uniformly poor quality of their data sheets.) The excess noise performance of APDs has been improving, due to innovations such as separating the photodetector region from the multiplication region, so this may become less of a problem in future.

The other major excess noise contributions come from uncertainties in the position of the first ionization event of the avalanche. Simple designs where the multiplication occurs in the absorption region have a 6 dB noise penalty, because some photons are absorbed deep into the multiplication region, so that the available gain is much less (they also have horrible variations of gain with wavelength, for the same reason). Newer designs that separate the absorption and multiplication regions via control of the doping density are much quieter and flatter with λ .

APDs exhibit a gain–bandwidth trade-off almost like an op amp’s. Due to the finite value of k , the avalanche takes more time to build up and die away as M increases, so the bandwidth tends to go as $1/M$. Even at low gain, the time it takes carriers to transit the (thick) multiplication zone limits the ultimate bandwidth. That’s why the quickest photoreceivers (40 Gb/s or ~ 30 GHz) use Schottky photodiodes and accept the reduced sensitivity.

The inherent SNR of the photocurrent generated by an APD is monotonically decreasing with gain; the signal quality gets worse and worse—so why on earth use them? The great virtue of APDs shows itself in wide bandwidth systems. For speed, these systems have to use low load impedances, and so are limited by the large additive Johnson noise. When the additive noise dominates all other noise sources, the (electrical) signal-to-noise ratio improves by M^2 until the excess noise from the APD becomes comparable to the Johnson noise, at which point the SNR peaks and begins to deteriorate. Thus the operating M should be chosen so that the excess noise roughly equals the Johnson noise of the amplifier (actually a bit higher since the falloff in SNR is very slow in the direction of increasing M , so the peak is not exactly where the two noises are equal). Alternatively, compared to a PIN diode, you can reduce the load resistance by a factor M^2 , which can help the bandwidth a lot. The price you pay is slightly lower SNR due to multiplication noise, narrower optical bandwidth, extra cost, and uncertainty in the exact operating gain.

The gain of an APD is a very strongly increasing function of bias voltage near breakdown, as shown in Figure 3.6 for a Hamamatsu S5343; a $\pm 20^\circ\text{C}$ change will make a nominal gain of 70 vary between 30 and 200. While this can be calibrated, and the bias




Figure 3.6. Gain of a Hamamatsu S5343 Si avalanche photodiode versus bias voltage, for various temperatures.

voltage or temperature controlled to keep it within closer bounds, we're never going to get even 1% accuracy over temperature with wild swings like that. This makes APDs poorly suited to accurately calibrated jobs in analog mode.

If you have matched APDs, either monolithically or by sorting, it is possible to stabilize the multiplied dark current (and hence the gain) of an illuminated APD by putting a constant current into the dark diode and applying the resulting voltage to the illuminated device. (You have to use a lowpass filter and a buffer, plus some thoughtfully chosen safety limits to avoid blowing up expensive APDs.) This is twice as expensive, somewhat noisier, and vulnerable to gradients, but saves the trouble of running temperature calibrations, which tend to be slow and hence expensive.

3.6.4 Photon Counting with APDs

APDs operated in the breakdown region (*Geiger mode*) are often used in medium-performance photon counting applications. The bias must be reduced after each event to stop the avalanche, typically by putting a huge resistor ($\approx 100\text{ k}\Omega$) between the device and the bias supply, with the output taken across the diode. The resulting RC time constant makes the recovery slow ($1\ \mu\text{s}$), but this does not affect their rise time, which can easily be below $1\ \text{ns}$ — $20\ \text{ps}$ has been reported.

Since all light pulses produce an output of the same size, pulse height discrimination is impossible. Compared with compact photomultipliers, these devices are more rugged, have higher quantum efficiency, and are available in larger diameters, but because of their slowness they have no compelling advantage in signal detection performance. Circuit improvements[†] can get them down to around $50\ \text{ns}$, which is somewhat better.

APDs emit a small amount of light when they break down, so if you have a multi-APD setup, you can get optical crosstalk. It's a good idea to reset all the APDs whenever any of them fires. You can also get segmented APDs intended for counting bursts of multiple photons in Geiger mode. In these devices, the active region is pixellated, but all the segments (as many as 14,000) are wired in parallel, via integrated quench resistors. These have higher dynamic range (since many segments can break down at once) but no better timing characteristics. Pulses from different segments are reasonably uniform in size, so that bursts of a few dozen photons can be counted with good pulse height discrimination. Besides the ordinary dead time correction (see Section 7.3.1), there is a small nonlinearity due to segments with multiple events. However, the main drawback of APDs for photon counting is their very high dark count rate—something like $50\ \text{MHz}/\text{cm}^2$ for an APD (Hamamatsu S10362-33-100C $3 \times 3\ \text{mm}$ segmented APD) versus $30\ \text{Hz}/\text{cm}^2$ (Hamamatsu H9319-01/11 $25\ \text{mm}$ PMT module), more than six orders of magnitude worse. This restricts photon counting APDs to very small areas.

APDs should be avoided if possible, since a PIN diode operating anywhere near the shot noise limit will always be superior, as well as cheaper and much easier to use. Unfortunately, we can't always have all the light we'd like, and when we can't use coherent detection, APDs can be a big help in the nasty $10\ \text{pA}$ to $1\ \mu\text{A}$ region.

[†]A. Spinelli, L. M. Davis, and H. Dautet, Actively quenched single photon avalanche diode for high repetition rate time gated single photon counting. *Rev. Sci. Instrum.* **67**, 1 (January 1996); A. Lacaïta et al., Performance optimization of active quenching circuits for picosecond timing with single photon avalanche diodes. *Rev. Sci. Instrum.* **66**, 4289–4295 (1995).

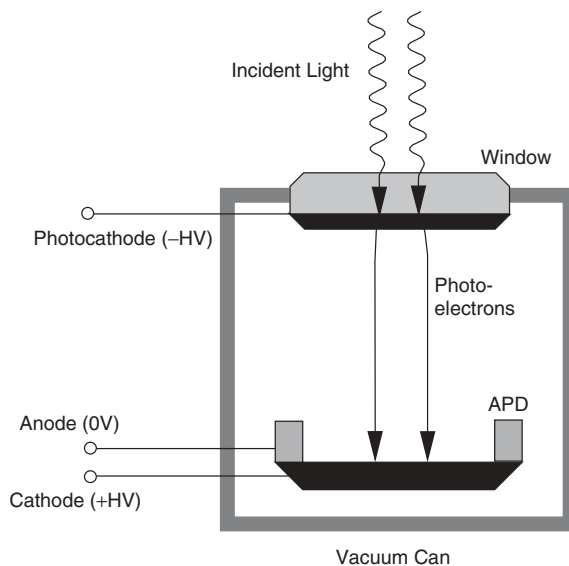


Figure 3.7. The vacuum APD (or hybrid PMT), an imaginative cross between APD and PMT; the high energy of the photoelectron hitting the APD produces a very well-defined pulse.

3.6.5 Vacuum APDs

An imaginative hybrid between the PMT and APD has been developed, the vacuum APD or hybrid photomultiplier (Figure 3.7). It consists of a photocathode and a silicon APD sealed into opposite ends of a vacuum tube, with a high potential (5–10 kV) applied between them. A photoelectron emitted by the photocathode slams into the APD surface with enough energy to generate as many as 2000 carrier pairs. Combined with avalanche multiplication, the overall gain is 10^5 – 10^6 , easily enough for photon counting. Because of the huge electron gain due to impact (equivalent to the first dynode gain in a PMT), the pulse height is much better controlled than in either a PMT or a regular APD, and is much larger than APD dark pulses, so the two are easily distinguished by thresholding.

You can also get these with regular PIN diodes as the anode, and those have the usual PIN advantages of stability and low noise, at the expense of gain. Either type gives a tight enough pulse height distribution that you can accurately compute the number of photons in a burst, at least up to about 10 photons or so. This is important, because it extends the usefulness of photon counting measurements to higher flux levels. The disadvantages of VAPDs include high cost, photocathode inefficiency, and the requirement of two very high voltage (+2 kV and –10 kV) power supplies. Still, these devices appear to be the state of the art for nonimaging detectors with gain.

3.6.6 Photoconductors

True photoconductive detectors must be distinguished from photodiodes operated at reverse bias, in the misnamed “photoconductive mode.” The two bear a superficial resemblance, in that the same circuit is used for both, but their operating principles and performance are quite distinct. A true photoconductor is a resistive chip or film made of a substance whose conductivity changes when it is illuminated, due to the generation of

electron–hole pairs when photons are absorbed. If a bias current is applied, the change in conductivity gives rise to a voltage change between the terminals.

Except in special circumstances, such as far-IR detection, where no good photodiodes exist, photoconductors are a poor choice of detector.

An ideal photoconductor is 3 dB noisier than an ideal photodiode; in a photoconductor, recombination becomes a Poisson process with the same variance as generation, so that the noise power doubles. More than that, though, photoconductors exhibit a fundamental trade-off between sensitivity and speed that photodiodes are free of. Once a carrier pair is generated in a photoconductor, it contributes to the conductivity until it recombines. Recombination can occur at the surface or in the bulk; surface recombination is usually faster, so that at short wavelengths, where the carriers are generated near the surface, the responsivity may drop.

A photoconductor exhibits gain equal to the ratio of the carrier lifetime τ to the transit time τ_{tr} , since it basically gets to reuse the carriers M times, where

$$M = \frac{\tau}{\tau_{tr}}. \quad (3.17)$$

This can be increased by increasing the lifetime or shortening the transit time. This relation appears very favorable, since it allows the possibility of useful gain with noise lower than that of PMTs or APDs (since the multiplication contributes no additional noise except for recombination). Unfortunately, the carrier lifetime limits the speed of the device, and so it must be made short for practical devices.

The only common photoconductors used in the visible are cadmium sulfide (CdS), cadmium selenide (CdSe), and their offspring, CdSSe. Besides being very slow, these devices have a sharply peaked spectral response that depends strongly on processing, a memory effect that can lead to factor-of-5 changes in cell resistance due to previous history, and a response time that depends on the illumination level; on the other hand, they exhibit a truly gigantic photoconductive response, which is helpful in applications such as night lights, where the amount of circuitry must be minimized. They are not terribly useful in high performance applications.

Infrared photoconductors are a bit better, but their main attraction is that they exist in a field with limited alternatives. They include HgCdTe, whose cutoff wavelength can be tuned from 5 to 22 μm by adjusting the Hg/Cd ratio; InSb, useful to 5.5 μm ; and lead salts PbS and PbSe, which are very slow and should be avoided if possible. Most of these require cryogenic cooling for good performance, which tends to cost thousands of dollars per unit. Room temperature photoconductors are available out to 15 μm , but are so insensitive that they are good only for detecting laser beams.

Far-infrared ($>20 \mu\text{m}$) photoconductors are also available; these are made from doped silicon and germanium. Due to dopant solubility limitations, they are nearly transparent to the radiation they are designed to detect, which seriously limits their quantum efficiencies; they also require cooling to liquid helium temperatures, and cost \$10,000 or more.

Designers who find themselves thinking that there must be an easier way should consider pyroelectric and thermal detectors. Photoconductor performance varies widely; consult the manufacturers' data sheets.

3.7 THERMAL DETECTORS

Not all photodetectors are based on quantum effects; the other major class is thermal devices and includes bolometers, thermocouples and thermopiles, pyroelectric detectors,

and mixed technology devices such as Golay cells. These devices are characterized by room temperature operation, low sensitivity, broad wavelength response, and low speed (except for some pyroelectrics).

A bolometer is a device whose resistance changes with temperature, such as a carbon thermistor or platinum film, and which has a very black surface. It is limited by the Johnson noise of its resistance and (if care is not taken) by the noise of its bias current and by thermal gradients. Increasing the bias current improves the sensitivity, but self-heating effects set a limit. Bolometers are used in a bridge configuration with a second, similar element shielded from incident radiation, to provide compensation for ambient temperature variations.

Thermocouples and thermopiles were the first IR detectors, dating from the 1830s. These devices do not require bias currents, as they generate their own output voltage from the temperature difference between two junctions of dissimilar metals or semiconductors. A thermopile is a series array of thermocouples, which gives a higher output voltage.

Bolometers and thermocouples are frequently used in spectrometers, where their room temperature operation and wide wavelength range are important advantages. Superconducting bolometers exploit the very large temperature coefficient of a superconductor near its transition temperature; they are extremely sensitive but obviously cumbersome.

Infrared bolometer arrays are becoming popular, in spite of poor sensitivity, because room temperature operation makes them much easier to use than cooled quantum detectors; they are somewhat cheaper as well.

Pyroelectric detectors are basically capacitors made of lithium tantalate (LiTaO_3), molecular crystals such as triglycine sulfate (TGS) and its derivatives,[†] or ferroelectric plastics such as polyvinylidene difluoride (PVDF), which is basically fluorinated Saran Wrap. By *poling* the material (causing a net dielectric polarization to be “frozen in”), the dependence of the polarization on temperature can be converted to a surface voltage change. High impedance AC amplifiers are used to detect these changes. Pyroelectric (PE) detectors do not work at low frequency, so PE detectors require their inputs to be modulated (e.g., with a chopper). Pyroelectrics work at room temperature but are not as sensitive as cooled semiconductor detectors. Their response increases as their Curie temperature (where they depole spontaneously) is approached. These devices have a very wide range of speeds and sensitivities, from submicrosecond devices used with pulsed lasers to pyroelectric vidicons that can see small ($0.1\text{--}0.5^\circ\text{C}$) variations around room temperature. Circuit hacks can get very competitive sensitivity ($\sim 0.13\text{ K NE}\Delta T$) in a low resolution image sensor very cheaply (see Section 3.11.16 and Example 17.1).

3.8 IMAGE INTENSIFIERS

3.8.1 Image Tubes

These are two kinds of spatially resolved image amplifiers. Image intensifier tubes use a photocathode and a scintillator in a vacuum tube with an electrostatic lens in between, to make a single photoelectron emitted from the photocathode produce many photons from the scintillator (up to 2×10^3 in one stage). These can be detected directly or can be run into another image intensifier tube. Stacks of three tubes have been popular, requiring

[†]Philips used to use *deuterated triglycine fluoroberylate* in their PE vidicons. To the author’s knowledge this stuff has the most jaw cracking name of any material in practical use.

power supplies of 15–50 kV at low current, and yielding photon/photon gains of 5×10^4 or so.

Image tubes are somewhat inconvenient to use, since electrostatic lenses have curved focal surfaces, so that the tube faces must be convex. Controlling field curvature requires negative lenses. Laplace's equation forbids this, because the field would have to be a maximum away from the boundary. This is a particularly inconvenient shape to match to the focal surface of a lens, which also tends to be convex (in the other direction of course). Even when a fiber optic face plate is used to flatten the focal surface, image tubes have relatively poor resolution. The image intensifier tube is not used as widely as it was 20 years ago, because of the advantages of microchannel plate (MCP) image intensifiers.

3.8.2 Microchannel Plates

A microchannel plate is a slab of glass containing a close-packed hexagonal array of small holes or channels, 5–15 μm in diameter, running from one face right through to the other. It is typically made by fusing together a bundle of specially made optical fibers, whose cores are of an easily etched glass. After fusing and slicing, the cores are etched away, leaving channels. The channels are lined with a thin layer of an electron multiplication material like that used in photomultiplier dynodes. A small current flows along this layer, allowing the layer to replace both the dynodes and the bias string of a PMT. A few kilovolts' bias is applied across the thickness of the plate, so that due to the voltage drop across a strong potential gradient exists along the length of the channel. An electron hitting the wall of the channel gives rise to a few secondary electrons, which cascade down the length of the channel, hitting the walls and resulting in a large multiplication factor from face to face. The length of the channels is typically 50 times their diameter.

A MCP image intensifier has a photocathode near one face and a scintillator near the other; due to the electron multiplication, a much higher photon–photon gain ($>10^4$) can be realized in a single stage with MCP, so multiple stages are usually unnecessary. The spatial resolution of a MCP is much better than that of an image tube, since it is defined by the geometry of the channels, rather than that of a rather fuzzy electrostatic lens. Microchannel plates are vulnerable to ion events; a straight MCP channel is an effective ion accelerator. Modern MCPs have their channels at an angle to the slab faces or arranged in a chevron or curved (“J”) arrangement. This increases the gain and reduces its spread, since it tends to force all electrons to hit the wall near the top of the channel. In addition, tortuous channels help force any ions to hit the sides instead of being launched into the photocathode.

MCPs are extremely fast; typical rise times for MCP electron multipliers are around 100 ps, and fall times 100–500 ps. The transit time spread is usually a few times less than the rise time. They are also very linear for channel currents up to perhaps 10% of the bias current per channel, at which point signal-dependent voltage drops along the channel begin to modulate the gain. Their multiplication noise is rather worse than a conventional dynode-chain multiplier's, because the secondary electron yield depends more on the trajectory of the electrons.

You can get MCPs with up to 64-segment anodes for time- and space-resolved photon counting applications. Their time resolution is excellent (~ 40 ps RMS) but they take a lot of circuitry.

Due to their huge internal surface area and constricted channels, MCPs cannot be baked out very well. Their interiors are thus much dirtier than those of PMTs, which limits their lifetime to about 0.2–0.5 coulomb of anode charge per square centimeter, *about 1000 times shorter than that of a regular PMT.*[†]

3.8.3 Streak Tubes

The high gain and fast shuttering of MCPs has one major disadvantage; all the photons from all times are superimposed on one another, so that there is no way of discovering the time evolution of the signal afterwards. One can use several MCPs, or make the signal periodic and use stroboscopic sampling, but these are not always possible; fortunately, there is a better way: the streak tube.

A streak tube is nothing more than an image tube with deflector plates in it, so that the image can be steered about on the output phosphor. A one-dimensional (1D) line image can be scanned across the output phosphor to produce a two-dimensional (2D) grey scale picture of the time evolution of every point in the line image—much like a 2D optical oscilloscope, but with a time resolution down to 1 ps. The moderate photon-to-photon gain of a streak tube, 10–500, is useful, because a signal has to be pretty intense to produce many photons in a picosecond. Fast-responding photocathodes such as S-20 are *de rigueur*—an NEA photocathode will smear the time response out to 1 ns or even slower. Electron storage materials can be used to make IR streak cameras. Streak cameras tend to cost \$100k or so, but they're 30 times faster than the fastest digitizing scopes.

3.9 SILICON ARRAY SENSORS

3.9.1 Charge-Coupled Devices

Silicon photodiodes are very good detectors in the visible and are made on the same substrates as ICs; thus they are natural candidates for making imaging array detectors. In operation, each array element operates as an isolated photosensor, accumulating charge in a potential well for some integration time, after which it is read out. The classical silicon array detector is the charge-coupled device (CCD). In CCDs, the readout mechanism is destructive but quiet; the charge from each element is shifted out by a clever analog shift register technique analogous to a worm screw feed, until it reaches a device lead or an on-chip sense amplifier. CCDs usually have on-chip amplifiers, which have enough gain that the subsequent stages are easy to design (one to a few microvolts per electron is typical).

The limits on CCD performance are set by photon efficiency, dark current, readout noise, and charge transfer losses. Provided the CCD is not shifted too rapidly, transfer efficiency will be 99.99% per stage at least. (This sounds great, but remember that this is raised to the 512th or even 4096th power—in a good device, you can get 99.9999% by slowing down a bit, which is dramatically better.) Nearly all the charge lost in transfer winds up in the following pixel, so that a bit of postprocessing can help.

All but the most expensive CCDs have some bad pixels, which have anomalously low sensitivity (*dead pixels*) or high dark current (*hot pixels*). Really cheap CCDs have whole dead columns. Your processing strategy has to take these into account.

[†]Philips Photonics, *Photomultiplier Tubes: Principles & Applications*, 1994, p.1–21.

3.9.2 Types of CCD

The two major classes of area-array CCDs are interline transfer (ILT) and frame transfer (FT). In an ILT CCD, the columns of sensor pixels are interdigitated with readout columns. When the integration time is finished, the entire image is transferred into the readout columns and is transferred out to be read. The areal efficiency of ILT devices is poor, 25–30%, since the readout columns take up space in the focal plane. On the other hand, the very fast lateral transfer makes these devices good shutters. The sensitive area in ILT CCD pixels is often weirdly shaped, which make its Fourier-domain properties less desirable, and the small size and odd shapes lead to serious moiré effects when imaging strong geometric patterns. Newer ILT CCDs having microlenses deposited on the top can have fill factors nearly the same as FT devices, with the attendant FOV reduction.

A frame transfer CCD has no interdigitated columns. When it is read out, the columns are shifted down into a shielded area array and read out from there. This keeps the fill factor high but requires hundreds of times longer to get the last of the pixels shifted into the dark array, so the shuttering performance is correspondingly worse. On the other hand, the fill factor can be 100%, and the square, full-pitch pixels reduce the moiré effects.

3.9.3 Efficiency and Spectral Response

Some CCDs are cheap, being produced in huge volumes for video cameras and other applications (usually with many dead pixels). These have poor long wavelength response, some cutting off at 600 nm or even shorter, although Sony makes some that work at some level out to 1000 nm. Front-illuminated CCDs are also insensitive in the blue and UV due to absorption of the films on the top of the silicon. Antireflection coating helps somewhat.

Illuminating the CCD from the back helps a lot, because of the greater absorption depth available and the absence of top surface obstacles. This requires thinning the die drastically, which is extremely expensive and hurts yield; nevertheless if you have the money, you can get CCDs whose fill factor is 100% and whose QE exceeds 85% after AR coating.

Typical commodity black-and-white CCDs have $\eta \approx 60\%$ and fill factors of 25%, so that their overall efficiency is only 15% or thereabouts. These specifications must be taken into account when designing a CCD detector subsystem. Some devices are becoming available with microlenses on the chip, to gather more of the incoming light into the active area, which helps a lot but causes spatial pattern problems, especially etalon fringes with lasers (see Section 3.9.6).

3.9.4 Noise and Dark Current

The readout noise of a CCD is additive and can be very low: as low as 2 electrons per pixel in a cooled device with a full well capacity of 10^6 electrons, read out at 20 kHz or so. Floating-gate readout devices have demonstrated sub-electron noise levels but are not commonly available. Over the past 20 years or so, commercial CCD readout noise has been stuck at the few-electron level.

There is a sharp trade-off between speed and readout noise; going faster requires more bandwidth, and the noise power scales linearly with speed. It's possible to achieve sub-electron noise by sampling slowly enough, but then it takes forever to read out the

camera. More commonly, the readout noise of a decent camera is around 30 electrons RMS. Achieving this noise level requires eliminating the kTC noise (see Example 13.3) contributed by the reset operation, which is done by sampling the output level before and after each charge readout, and subtracting the two voltages, a procedure called *correlated double sampling*. (Interestingly, since no charge is dissipated in the shifting operations, no kTC noise is introduced by shifting the charge packet into the readout cell.)

Increasing the integration time increases the signal linearly, so the electrical SNR of low light images goes as t^2 , until the dark current noise $(i_{\text{dark}}t)^{1/2}$ equals the readout noise, after which the SNR is shot noise limited and so increases only linearly with time. The attainable (electrical) SNR is limited by photon statistics to a value equal to the full well capacity in electrons, around 47 dB for the camcorder CCD to perhaps 60 dB for a scientific device.

A garden-variety camcorder CCD operated at room temperature has a dark current of about 100 electrons per pixel in a $\frac{1}{30}$ second integration time, and a well capacity on the order of 5×10^4 electrons. A cooled CCD, as used in astronomy, can achieve dark currents far below 1 electron/s per pixel. These very low dark currents are achieved by multiphase pinning (MPP), which eliminates the effects of surface states (MPP is also called *inverted mode*). Since the dark current is thermally activated, it drops by a factor of 2 every 10°C . Cooling cannot be taken too far, however, due to the “freezing out” of trap sites, which leads to serious charge transfer losses at low temperatures (-30°C for some devices, to -100°C for others). Freeze-out is not subtle; you get severe streaks along the transfer direction.

3.9.5 Bloom, Bleed, and Fringing

Prominent CCD pathologies are bloom, bleed, and fringing. Bloom is the picturesque name given to the artifacts that result when some elements become full and spill over into the readout circuitry or the substrate, causing spurious lines and a general leakage of charge into adjoining elements. A badly bloomed image is full of bright lines and blobs. Bloom is controlled by additional electrodes that extract the charge before it can migrate far; antiblooming compromises linearity, fill factor, and full well capacity, so use it only if you really need it. Note that bloom is not necessarily restricted to illuminated areas; by dumping charge into the substrate, saturated pixels can send charge into storage areas as well, a particular problem in scientific applications, when the readout time is sometimes long compared to the integration time.

Fringing and bleed cause loss of resolution in CCDs operated at wavelengths near the $1.1 \mu\text{m}$ absorption edge of silicon. The silicon becomes increasingly transparent, and light can bounce back and forth between the front and back surfaces of the wafer, causing nasty irregular fringes and scatter. This is not a trivial effect: at short wavelengths the fringes tend to be 1% or below, but in the IR they can get up to 5% or even more, which is very objectionable.

If we try to make the silicon thicker, so that more absorption occurs, lateral diffusion of carriers in the field-free region (away from the depletion zone of the junction) causes bleed, where carriers from one pixel diffuse into neighboring ones. The best solution to this is to use a thick piece of high resistivity silicon, which can be depleted throughout the volume (like a back-biased PIN diode). High resistivity CCDs have high quantum efficiency even at 1000 nm. If you haven't got one of these special devices, there's not a lot you can do about it.

3.9.6 Spatial Pattern

CCDs also have stable, accurate spatial patterns; subpixel position measurements of images a few pixels in size can be done by interpolating. Carefully designed, back-illuminated CCDs have high fill factors, which makes their Fourier domain properties simple, but front-surface devices are much iffier.

Fancier Fourier processing techniques should not assume that the spatial sensitivity pattern of a PMT pixel looks like a nice rectangular box, even if the pixel is rectangular: most of the time, it has sloping sides and a bit of dishing in the middle, so that it resembles a fedora more closely than it does a top hat. Some types have serious asymmetry between two halves of the pixel (e.g., when the sensitive area is L-shaped). These effects make the optical transfer function of your CCD do things you might not expect; for example, the zero in the optical transfer function of the CCD is not at $2\pi/\text{pixel width}$ but is somewhat further out; irregular pixel shapes make it even worse because of their high spatial harmonic content. The only defense against this in serious Fourier processing is to oversample by a big factor (say, $4\times$ to $8\times$ the Nyquist limit), so that the big irregularities in the OTF happen out where you don't care about them, and you can correct for the small ones that remain. Devices that are not efficiently AR coated will do unintuitive things, because of reflectance changes with incidence angle. For an air–silicon interface, the normal incidence reflectance is 0.3, whereas at 30° it's 0.36 for s and 0.26 for p . Measuring pixel sensitivities directly with a high NA flying spot will thus produce some funny results due to this asymmetric pupil apodization.

At long wavelengths, bleed spreads incident light into adjacent pixels and causes the OTF to be seriously wavelength dependent; how much this happens depends on the absorption of the silicon.

CMOS image sensors and interline transfer CCDs with microlenses exhibit horrible etalon fringes when used with temporally coherent sources.

3.9.7 Linearity

Unlike imaging tubes such as vidicons, CCDs are normally extremely linear, although the antiblooming provisions of some devices can cause serious nonlinearity above about half-scale. Infrared focal plane arrays are much less linear and usually require multiple calibrations at different light intensities.

At low temperatures and long wavelengths, front-illuminated CCDs exhibit QE hysteresis due to trapping of charge at the interface between the bulk and epitaxial layers (this is polished away in back-illuminated CCDs). The CCDs used on the Hubble, Galileo, SXT, and Cassini space missions had QE variations as much as 10% depending on signal level and operating temperature.[†]

Aside: CCD Data Sheets. From an instrument designer's viewpoint, the worst thing about CCDs is their data sheets. Op amps, microprocessors, diode lasers—all these have reasonably standard spec sheets, but not CCDs. Designing a CCD system that will be replicated more than a few times depends on the designer having a deep knowledge of the details of each kind of CCD considered. The data sheets are also nearly all hopelessly out of date. Consult the manufacturer of your devices for the latest specs—and don't be

[†]J. Janesick, posted to the CCD-world mailing list <http://www.cfht.hawaii.edu/~tmca/CCD-world/>, March 18, 1999.

too surprised if they pay more attention to their camcorder-building customers than they do to you.

3.9.8 Driving CCDs

CCDs take a lot of circuit support. This is somewhat specialized, so if you're rolling your own CCD drivers, you probably want a copy of Janesick. In general, CCDs are forgiving of mildly ugly clock signals as long as they're highly repeatable. People often use transmission gates connected to well-filtered analog voltages to produce the funny clock levels required.

3.9.9 Time Delay Integration (TDI) CCDs

Linear CCD sensors are used in line-scan cameras and in pushbroom-style remote sensing satellites, where the spacecraft motion supplies the frame scan. The SNR can be improved by *time-delay integration* (TDI), in which the linear array is replaced by a narrow area array (perhaps 4096×64 pixels) and clocked in the narrow direction at the same rate that the image crosses the detector, so that the same object point illuminates the same bucket of electrons throughout. This requires accurate alignment and places severe constraints on the geometric distortion of the imaging system, but a 64-deep TDI sensor can get you a 36 dB signal increase.

Another TDI advantage is uniformity. Because each object point is imaged by (say) 64 pixels in turn, the fixed pattern noise tends to average out. There's a bit more in Section 10.5.3.

3.9.10 Charge-Multiplying CCDs

Because the charge transfer efficiency of a CCD is so high, and its dark current can be made very low, its major noise source is its output amplifier. The noise can be reduced to subelectron levels by slowing the readout clock and filtering or averaging the output. The noise/speed trade-off prevents the use of CCDs for real-time imaging at low light levels, as we've seen. Hynecek[†] and more recently Mackay et al.[‡] have made a clever electron-multiplying CCD that overcomes this trade-off almost completely, the *low-light-level CCD* or LLLCCD (or L^3 CCD). This device can be used either as a normal CCD of quantum efficiency η , or as the equivalent of a noiselessly intensified CCD with a QE of $\eta/2$ and a readout speed of 20 MHz. The way it works is to take an ordinary single-output CCD and add a few hundred extra transfer stages before the readout amplifier. In the extension, one of the three readout phases is run at a much higher voltage, enough that there is a small amount ($\approx 1-2\%$) of electron multiplication in each stage. Because there are many stages, reasonably well-controlled gains from 1 to several thousand are easily realized by changing the voltage in the one phase, so the same device can go from sunlight to photon counting by changing one voltage. Since the multiplication occurs inside the silicon, there is no dirty vacuum to limit its lifetime. The complexity increase is much smaller than that of an MCP. Furthermore, blooming

[†]Jaroslav Hynecek, CCM—a new low-noise charge carrier multiplier suitable for detection of charge in small pixel CCD image sensors. *IEEE Trans. Electron Devices* **30**, 694–699 (1992).

[‡]Craig D. Mackay, Robert N. Tubbs, Ray Bell, David Burt, Paul Jerram, and Ian Moody, Sub-electron read noise at MHz pixel rates. *SPIE Proc.* January 2001.

of the multiplication stages automatically protects against damage from bright lights that would reduce an MCP to lava.

The LLLCCD needs cooling to get its dark current low enough to count photons. In addition, its dynamic range is limited by the full well capacity of the multiplication stages, and multiplication causes the shot noise to go up by 3 dB, which is actually amazingly good. In a PMT, secondary emission is a Poisson process, so 512 dynode stages, each with a secondary electron yield of 1.01, would produce a pulse height histogram that looked like a pancake (its variance would be about 100 times the mean; see the chapter problems in the Supplementary Material). Electron multiplication in the CCD involves a Poisson process too, with one key difference: with 512 stages of 1.01 gain, the 1 is deterministic and hence noiseless—only the 0.01 is Poissonian, so the variance is only twice the mean, not 100 times. (Why?)

There is a slight linearity trade-off involved at very high gains—bright pixels see a somewhat *higher* gain than dim ones, the opposite of the usual situation. The error can be a factor of 2 in really bad situations, but as it's monotonic, with care it can be calibrated out. A significant advantage of L^3 CCDs is the ability to achieve high frame rates, because the usual trade-off of readout noise versus frame rate has been greatly improved by the amplification ahead of the readout amplifier. Because of their sensitivity, flexibility, potentially low cost, and long life, electron-multiplying CCDs will probably replace image intensifiers in applications where cooling is feasible and fast shuttering is not required.

For applications such as astronomical spectroscopy, where the 3 dB SNR loss is too great, L^3 CCDs can be operated in photon counting mode, provided the frame rates are high enough that you don't lose too many counts due to two photons hitting a given pixel in the frame time. Because the device operational parameters don't need to change between analog multiplication and photon counting modes, you could in principle change the voltage on the multiplying phase during readout, which would give different pixels different gains, increase the dynamic range.

3.9.11 Charge Injection Devices (CIDs)

CIDs are like CCDs only backwards: the well starts out full, and light removes charge instead of adding it. This adds noise at low light levels but makes CIDs intrinsically resistant to bloom, and so suitable for high contrast applications. Their quantum efficiencies are typically fairly low, around 30%. CIDs use multiplexers instead of shift registers and can be read nondestructively, since the charge need not be removed from the element during readout. CMOS imagers also use big multiplexers (see below). These multiplexer-based technologies offer random access, so we can use different integration times for different pixels on the array. This means that an image with extremely high contrast can be read adaptively, yielding the best of all worlds: rapid acquisition of bright objects and long integration times to enhance detection of faint ones. Since CMOS imagers don't exhibit bloom either, CIDs are no longer so unique in that way.

3.9.12 Photodiode Arrays

Photodiode arrays (commonly known as “Reticons” for historical reasons) look like CCDs but are actually read out via a big multiplexer instead of a bucket brigade, which makes them much easier to control. They are competitive only in 1D arrays, as in OMA

spectrometers, but they work very well for that. Compared with CCDs, they are not limited by charge transfer inefficiency, which suits them well to high contrast applications, where the data will come under close scrutiny, as in spectroscopy. They are a bit less sensitive than CCDs of the same pixel size, however, because of multiplexer noise and charge injection. More insidiously, they are dramatically less linear than CCDs.

Photodiode arrays, unlike CCDs, generally have no bias applied to the diodes during integration; thus they are somewhat nonlinear at large signals, because of the forward conduction of the photodiodes. The forward voltage drops and the dark current increases with temperature, so the linearity, full well capacity, and available integration time all degrade more quickly with temperature than in a CCD. A 10°C increase doubles the dark current and reduces the charge capacity by about 25%.

3.9.13 CMOS Imagers

The CMOS imager is similar in performance to an interline transfer CCD but is manufactured on an ordinary CMOS process, as used in logic and memory ICs. It consists of an array of pixels, each with its own buffer amplifier. The amplifiers are connected to a gigantic crossbar multiplexer, as in a dynamic memory IC. Often there is a single stage of CCD-style charge transfer between the pixel and the amplifier, to provide the fast shuttering capability of an ILT CCD. The attraction of a CMOS imager is that lots of the ancillary circuitry required for a CCD, such as clock generation, amplification, A/D conversion, panning, zooming, and even image processing (e.g., halftoning), can be done on the imager chip. The optical characteristics are not generally as good as CCDs; the dark current is generally higher, and the linearity and accuracy less. The large number of amplifiers all have slightly different offset voltages, so that there is a lot of fixed-pattern noise in CMOS imagers that is not present in CCDs, which have many fewer amplifiers (often only one). This makes their dim-light performance poor, but work is underway to bring these devices to the capabilities of scientific CCDs. Being able to use imaging sensors without dragging along six tons of ancillary hardware is a powerful and useful capability, but on the other hand the flexibility of custom controllers is sometimes invaluable.

CMOS imagers often exhibit severe fringing due to reflections between the fairly dense metal wiring layers and the silicon surface, and microlenses make it worse. Effective optical path lengths are in the 10 μm range. This isn't too terrible in wideband applications, but it can make narrowband measurements and spectroscopy exciting. On the other hand, since there are real FETs between each CMOS pixel and its neighbors, CMOS imagers generally don't bloom.

3.9.14 Video Cameras

Television is a vast wasteland.

—Newton N. Minow (then Chairman of the US Federal Communications Commission)

Video cameras use CCDs or CMOS imagers, but most of them are poorly suited to precise measurements. They have black level and gain adjustments (usually automatic) that often foul up measurements by changing when we want them to keep still. In order to mimic the behavior of photographic film and vidicons, their response is deliberately

made nonlinear, sometimes with an adjustment for γ , the contrast exponent. There are cameras made for instrumentation use, and they work well enough but tend to cost a lot compared to the entertainment-grade ones. Since gain and black level adjustments are done on the full frame, putting an intensity reference in the field of view will sometimes allow you to fix it afterwards, but the γ problem is still there.

Cameras are useful in full field interferometric techniques (phase shifting, moiré, and holographic) and in structured light measurements (see Section 10.5.9). On the other hand, cameras are often overused by people whose primary strength is in software, and who want to get the data into digital form as early as possible. The importance of getting the detector subsystem right cannot be overemphasized, so this by itself is an inadequate reason for using video.

Sometimes video is necessary, for example, in machine vision systems or where the requirement for using commercially available hardware is more compelling than that for optimizing performance. In the long run, however, it's such a headache that it is vitally important to make sure you really understand why you're using video rather than one or a few silicon photodiodes. Its cost and complexity rise very quickly once you get past webcam quality, and even so the measurements are normally poor, due to the low SNR of image sensors, the 8 bit limit of the A/D converters, and the generally poor fidelity of the electronics used.

Color video is even worse than black and white. Color sensors are made by putting arrays of different-colored filters on top of a monochrome sensor.[†] These filters are generally arranged in groups of four, one pixel each of red and blue, and two of green. In order to prevent bad moiré effects, color sensors have “defuzzing filters,” which are thin walkoff plates (see Section 6.3.5) mounted over the CCD to smear out the image over a four-pixel area. Of course, this reduces the resolution by a factor of 2 in each direction, and equally of course (marketing departments being what they are), the quoted resolution is that of the underlying sensor. When comparing color to monochrome, remember that it takes four *Marketing Megapixels*[™] to make one real monochrome megapixel. Avoid color sensors for instrument use whenever possible.

3.9.15 Extending the Wavelength Range: CCDs + Fluors

The UV performance of a front-illuminated CCD detector can be dramatically improved by applying a thin coating of a fluorescent material to convert incident UV photons into visible light before they are absorbed. The quantum efficiency of this approach is around 10% for many fluors, which is not as high as we'd like but a lot better than nothing, which is what we'd get otherwise.

3.9.16 Electron Storage Materials

You can get near-IR sensor material for converting 0.8–1.6 μm light to visible. The best is Q-42 phosphor from Lumitek Corp (available as IR sensor cards from Lumitek and Edmund Optics, since Kodak went out of the business). It uses rare-earth oxides or sulfides in a calcium sulfide matrix and works by electron trapping: visible or UV excites the Ce^{2+} , Er^{2+} , or Sm^{3+} ions to a long-lived metastable state; when an IR photon

[†]There are honorable exceptions, in which dichroic prisms are used to form RGB images on three separate CCD chips, but you'll probably never see one.

comes along, it is absorbed, and a visible one emitted. The fluorescence dies away very rapidly (picoseconds to nanoseconds) when the IR stops, and the quantum efficiency is near unity. This material could probably be used to extend silicon CCD detectors out to 1.6 μm , in much the way as the UV fluors. It would obviously have to be pumped with a wavelength to which the CCD is insensitive, or refreshed during readout like dynamic memory. For near-IR applications at high power density, you can also get second-harmonic generating material, such as nitrofurazone (5-nitro 2-furaldehyde semicarbazone)[†] (5-nitro 2-furaldehyde semicarbazone, Aldrich Catalog #73340), which can be mixed with paint. Water and polar solvents deactivate it, but it survives well in anisole and NMP (n-methyl pyrrolidone).

3.9.17 Infrared Array Detectors

Platinum silicide arrays are made on silicon substrates, but other infrared arrays are odd hybrid devices, generally consisting of detector chips made of InGaAs (to 1.7 or 2.2 μm), InSb (to 5 μm), or HgCdTe (2.5–14 μm) bump-bonded to Si readout chip. Since the chip metal goes on the top, the chips are bonded face-to-face. Thus the light comes in through the substrate of the photodetector chip, which is the origin of the short-wavelength cutoff in most IR imaging arrays (e.g., the absorption edge at 850 nm due to the CdZnTe substrates used for many HgCdTe detector arrays). Nowadays more of these devices are being back-thinned, so it is possible to get InGaAs detectors with response as far down as the ultraviolet (<400 nm). The long-wavelength edge of HgCdTe is tunable from 2.5 μm to 14 μm by changing the alloy ratios, with the leakage becoming worse for longer wavelength cutoff, as the bandgap decreases. Near-IR sensors such as InGaAs as well as lower performance devices such as pyroelectrics and microbolometer arrays can work at room temperature, but all the others require cryogenic cooling.

IR arrays are much less linear than silicon CCDs, and at the current state of the art, their dark current and response nonuniformities are much worse. Thus calibrating an IR array isn't the simple matter it is with silicon CCDs (see Section 3.9.19). Even with pixel-by-pixel correction for gain and offset, the detection performance of IR arrays is usually limited by their residual fixed-pattern noise. Using calibrations taken at several radiance levels, and fitted with low-order polynomials or sums-of-exponentials, can greatly (20–30 dB) improve matters, at least within the range measured. Platinum silicide arrays have enormously better uniformity than InSb and HgCdTe arrays, so they often offer effectively lower noise in the mid-IR even though their QE is very poor. Pixels with bad $1/f$ noise are usually the limiting factor in the stability of the calibration; PtSi calibrations are good for days, but InSb and especially HgCdTe arrays require recalibration on timescales of minutes to hours, if the spatial noise is going to be below the temporal noise.[‡]

3.9.18 Intensified Cameras

Electron multiplication is often helpful in low light situations, to overcome circuit noise. It is natural to try applying it to imaging detectors such as vidicons and CCDs, to overcome

[†]Used as a topical antibiotic (Furacin), but also good for making 0.5–0.7 μm from 1.0–1.4 μm .

[‡]Werner Gross, Thomas Hierl, and Max Schulz, Correctability and long-term stability of infrared focal plane arrays. *Opt. Eng.* **38**(5), 862–869 (May 1999).

their dark current and readout noise in the same way. Just now, the dominant type of intensified camera is the combination of an MCP image intensifier with a CCD sensor. Older types, such as the image orthicon, the silicon intensifier target (SIT) vidicon, and cameras based on image converter tubes and electron optics are now seldom used (see Section 3.9.10).

An MCP camera consists of a microchannel plate image intensifier whose output is coupled to a CCD imaging detector. These are often proximity focused; the phosphor screen is close to the CCD, making an extremely compact, simple, and robust system. MCPs have spatial resolutions of $10\ \mu\text{m}$ or so, a good match to a CCD.

Such intensified cameras normally produce noisy output, because generally they can't improve the photon statistics of the incident light. You can't just use long exposures, because each primary may produce 2000 electrons in the CCD well, so you get only a few dozen counts before saturating. Such noisy signals often need frame averaging before they become useful, which limits the utility of intensified cameras. For instrument use, they're usually at a disadvantage compared with cooled CCDs using long integration times. On the other hand, image intensifier cameras are very suitable when the images are primarily intended to be viewed by eye; real-time imaging becomes more important then, and spatial averaging in the human visual system can take the place of frame averaging.

Another reason for using MCP intensified cameras is their very fast time-gating capability; an MCP has a rise time of 100 ps or so, and its gain can be turned on and off in a few nanoseconds, making it an excellent shutter as well as an amplifier. The two methods for doing this are to use an avalanche transistor circuit to gate the whole MCP bias voltage, or to use an MCP with a grid between photocathode and MCP array, which requires only a few tens of volts. Cameras exist that will take a few frames at 10^7 fps.

Aside: Night Vision Goggles. Direct-view image intensifiers, (e.g., night vision goggles) are helpful for a few other reasons. Photocathodes have about four times the QE of rod cells, plus a much wider wavelength band, especially toward the IR where the sky glow is brighter; they bring the cone cells into play, which have much higher resolution and higher response speed; and the collection system can have a much larger étendue than the eye, both in area and in solid angle.

3.9.19 Calibrating Image Sensors

Achieving imaging performance limited by counting statistics is a bit more of a challenge than this, because image sensors have characteristics that vary from pixel to pixel. Each output amplifier will have its own characteristic bias voltage; each pixel will have slightly different sensitivity and dark current. Separating out and correcting these effects is nontrivial.

What we initially measure is a raw data frame R , but what we want is the true image intensity I , corrected for gain and offset. Providing that the sensor is very linear (e.g., a properly operated CCD), the best way is to take long- and short-exposure dark frames, $D_l(t_l)$ and $D_s(t_s)$, and a flat field frame, $F(t_f)$, using the real optical system aimed at a featureless surface such as the dawn sky or an integrating sphere. A bias frame (just the bias voltage) B can be constructed as

$$B = \frac{t_s \cdot D_l - t_l \cdot D_s}{t_s - t_l}, \quad (3.18)$$

a normalized thermal frame (just the dark current in 1 s) as[†]

$$T = \frac{D_l - B}{t_l}, \quad (3.19)$$

and a sensitivity frame as

$$S = F - B - t_F \cdot T, \quad (3.20)$$

where S is normalized to unity at the radiance of the featureless surface. Note that all these frames should really be averaged across at least four separate exposures so that the noise of the calibration frames does not dominate that of the data you're going to take with the system. Make sure that the long-exposure dark frames are at least as long as your longest measurement, and that the flat fields are just below half-scale to guard against nonlinearity. The long-exposure frames will be slightly corrupted by cosmic ray events, so in the averaging operation, code in a check that throws out the highest value in each pixel if it's way out of line (e.g., four times the RMS noise).

When all this work is done, a properly normalized and calibrated true image I can be computed from a raw image $R(t_R)$ of exposure time t_R :

$$I = \frac{(R(t_R) - B - t_R \cdot T)}{S}. \quad (3.21)$$

This is a flexible and convenient calibration scheme, because you can make your exposures any length you need to. You'll need to experiment to determine how often it needs to be repeated. Watch out especially for sensor temperature changes, which will make these calibrations go all over the place. (See the problems for how many of each kind of frame you need.) Above all, make sure that your sensor is always operating in a highly linear regime: if your hot pixels saturate, you can't do good dark frames; if your flat fields are above half-scale and the antiblooming or MPP is on, you'll have the nonlinearity.

The flat field is wavelength sensitive, so make sure you take the flat field frames with a source of the same spectral characteristics as your typical data. Fringing in thinned, back-illuminated CCDs makes the wavelength variation of the nonuniformity sharper than you might expect; the normalized photoresponse nonuniformity (PRNU) of a good back-illuminated CCD is 1% or so in the red, rising to 5–10% in the UV and as much as 20% near the IR cutoff.

As we discussed earlier, IR arrays are much more difficult to calibrate, because they are not as linear, are more temperature sensitive, and are less uniform to begin with.

3.9.20 Linearity Calibration

CCDs are usually pretty linear up to about half the full well capacity, and many are very linear nearly to full well. On the other hand, it is much more comfortable knowing than hoping. One good way to get a linearity calibration is to use a diffused LED source (a few frosted LEDs all around the CCD, at some distance) that provides a nice uniform illumination across the whole chip (a few percent is OK, you can remove that

[†]If you're using integer arithmetic, you'll want to normalize to some longer time interval to avoid significance loss.

mathematically). Start with the CCD reset and shift the image out normally. Flash the LEDs once after each row is shifted out, and you get a nice intensity staircase signal that will tell you a great deal about your linearity curve. This is much faster than doing many, many full frame calibrations, and so can be used as an online linearity calibration. Note that LEDs generally have a temperature coefficient of output power near $-1\%/^{\circ}\text{C}$, so for decent accuracy you have to temperature compensate them or use them in a closed-loop system with a photodiode.

3.10 HOW DO I KNOW WHICH NOISE SOURCE DOMINATES?

The most basic limit to the sensitivity of an optical measurement is set by the shot noise of the signal photons. Once other noise sources have been reduced below this level, further SNR improvements can come only from increasing the signal strength or narrowing the bandwidth. There is a certain cachet to “shot noise limited” or “quantum limited” measurements, which should not be allowed to obscure the fact that such measurements can still be too noisy, and that most of them can still be improved.

The contributions of Johnson noise, signal and background shot noise, and thermal fluctuations are easily calculated from parameters given in typical data sheets; the one major imponderable is lattice (thermal) generation–recombination (G-R) noise in IR photodetectors, which depends on the carrier lifetime, a number that is not always easily available. In most cases, one must use a seat-of-the-pants method of estimating lattice G-R noise, such as taking the published noise specification and subtracting all the other noise sources, or rely on the detector manufacturer’s assertion that a detector is background limited at given detector and background temperatures and field of view.

Formulas for noise contributions are often given in a form that explicitly includes the modulation frequency response of the detector. This seems unnecessary. Apart from $1/f$ noise, the detector noise has the same frequency response as the signal, and combining different authors’ complicated formulas is cumbersome. The frequency response is a matter of deep concern to the designer, who is unlikely to be misled by, for example, someone describing shot noise as having a flat power spectrum.

The total noise contributed by a detector will depend on just how the detector is coupled to the external circuitry. Calculating this requires some sort of circuit and noise models of the detector. Determining the overall SNR of the signal emerging from the detector subsystem is a major topic of Chapter 18. Table 3.1 is a good starting point for figuring that out.

3.10.1 Source Noise

In many measurements based on externally applied illumination, source noise is the dominant contributor. Great ingenuity is expended on reducing and avoiding it. In laser-based measurements, especially bright-field ones such as interferometry, absorption spectroscopy, and transient extinction, laser residual intensity noise (RIN) is frequently the dominant contributor. It may easily be 60 dB above the shot noise level. Fortunately, it is usually tractable; see Sections 10.6.2 and 10.8.6 for how to handle it. Laser frequency noise must usually be treated by stabilizing the laser.

Noise from incoherent sources, especially arc lamps, is more difficult to deal with since it is strongly dependent on position and angle, so that compensating for it by

TABLE 3.1. Which Noise Source Dominates?

Detector Type	Noise Source	Dominates When ^a	Noise Spectral Density
Si, Ge, InGaAs photodiodes	Photocurrent shot	$i_s R_L > 2kT/e$ (50 mV @ 300 K)	$i_N = (2ei_s)^{1/2}$
	Background shot Johnson	$i_b R_L > 2kT/e$ $(i_s + i_b)R_L < 2kT/e$	$i_N = (2ei_b)^{1/2}$ $i_N = (4kT R_L)^{1/2}$
IR photodiodes	Photocurrent shot	$i_s(R_L \parallel R_{sh}) > 2kT/e$ (50 mV @ 300 K)	$i_N = (2ei_s)^{1/2}$
	Photon (background shot)	$i_b(R_L \parallel R_{sh}) > 2kT/e$	Eq. (3.10)
	Lattice generation/recombination	Only when reverse biased	Eq. (3.23)
	R_{sh} Johnson	Always unless cryogenically cooled	$i_N = (4kT/R_{sh})^{1/2}$
IR photoconductors	Shot (photogeneration/recombination)	$i_s G(R_L \parallel R_{sh}) > kT/e$ (25 mV @ 300 K)	$i_N = (4Gei_s)^{1/2}$ (recombination doubles variance)
	Lattice G-R	$V_{DC} \tau \mu > \ell^2 kT/e$	Eq. (3.23)
	Photon	BLIP when cryogenically cooled (believe manufacturer)	Eq. (3.10)
	Johnson	Always unless cryogenically cooled	$i_N = (4kT/R_{sh})^{1/2}$
	Thermal detectors	Johnson Thermal fluctuations	Nearly always
Avalanche photodiodes	Shot	Almost never	$i_N = (2Mei)^{1/2}$
	Multiplication	Almost always	$i_N = (2M^{1+x}ei)^{1/2}$
	Johnson	Only if M is too low	$i_N = (4kT/R_L)^{1/2}$

^aHere i is the actual current from the device (after multiplication, if any).

comparing the measured signal in some way to a sample of the light from the source may not be sufficiently accurate to control the effect. It is possible to stabilize arcs using strong magnetic fields, but as this is awkward and expensive, it is seldom done; the problem remains a difficult one. Incandescent bulbs are much quieter than arcs and are often a good choice.

Source noise is multiplicative in character, since both signal and coherent background are proportional to the source intensity; the multiplication of the source noise times the coherent background level gives rise to the additive part of the source noise, while source noise multiplying the signal puts noise sidebands on the desired signal, a particularly obnoxious thing to do. See Chapter 2 for more on source noise and what to do about it.

3.10.2 Shot Noise

Shot noise is the easiest limit to calculate, but the hardest to improve upon. There is a fair amount of confusion about shot noise, where it comes from, and where it applies.

The formula, which states that if the arrival of electrons at a given circuit point is a Poisson process, the current i will have a noise current spectral density of

$$\langle i_N \rangle = \sqrt{2ei} \text{ A/Hz}^{1/2}, \quad (3.22)$$

is perfectly correct. However, much of the time the current is not Poissonian, so that the formula is inapplicable. The simplest case is in ordinary photodiodes, where all photocurrents and leakage currents exhibit exactly full shot noise, regardless of quantum efficiency.

It is easy to make currents with full shot noise, much less than full shot noise, or much more than full shot noise. A battery connected to a metal film resistor will exhibit much less than full shot noise, because any charge density fluctuations are smoothed out by electron scattering, which tends to reestablish the correlations and reduce the temperature of the electrons; the shot noise power is reduced by a factor of L/l , where L is the length of the resistor and l is the mean free path for electron–electron scattering.[†] Since $l \sim 100 \text{ \AA}$ for disordered metals, and $L \sim 1 \text{ mm}$, the suppression is pretty strong.

An avalanche photodiode's dark current will exhibit much more than full shot noise, since each electron generated in the junction gives rise to a large pulse; even if all N pulses arriving per second were exactly A volts tall, the RMS fluctuation in 1 s will be $A\sqrt{N}$. This is \sqrt{M} times larger than the shot noise corresponding to the average output current.

Shot noise in an optical beam can be viewed as the interference term between a noiseless optical signal and the zero-point *vacuum fluctuations* of the electromagnetic field. It affects both the amplitude and phase of a detected signal; by heroic preparation, in very special circumstances, the noise can be redistributed slightly between two conjugate components (sine and cosine, or amplitude and phase), but the product of the two components cannot be reduced. It is thus a fundamental feature of the interaction of light with matter.

Thermal light, at least at frequencies where $h\nu \gg kT$, is Poissonian to good accuracy; at lower frequencies, where the occupation number of the quantum states is higher, there is an additional term due to the Bose–Einstein statistics of thermal photons; this makes the variance of the photon count N go up by a factor of $1 + 1/[\exp(h\nu/kT) - 1]$.[‡] We are nearly always in the high frequency limit in practice.[§]

It is useful to concentrate on the statistics of the initially generated photocarriers, before any gain is applied, because in all quantum detectors this current exhibits full shot noise, and amplification does nothing whatever to improve the signal to shot noise ratio. Like quantum efficiency, this ratio must be kept in mind; it is a vitally important sanity check.

Carrier recombination is also Poissonian, so the shot noise variance is doubled. This double shot noise in photoconductors is sometimes called photocarrier

[†]There's been a lively literature on the details of this for the last 25 years, and there's a lot of interesting physics there: for example, see R. Landauer, *Ann. New York Acad. Sci.* **755**(1), 417–428 (1995).

[‡]For example, see E. L. Dereniak and G. D. Boreman, *Infrared Detectors and Systems*. Wiley, Hoboken, NJ, 1996, Section 5.2.2.

[§]Hanbury Brown and Twiss demonstrated that these classical fluctuations could be used to measure the angular diameter of hot thermal sources such as blue stars by cross-correlating the measured noise from two detectors whose spacing was varied—a so-called *intensity interferometer*. See Robert Hanbury Brown, *The Intensity Interferometer*. Halsted Press (Wiley), Hoboken, NJ, 1974. Hanbury Brown is one of the present author's technical heroes.

generation–recombination noise, but it passes the duck test[†] for shot noise. This renaming leads to the highly misleading statement that photoconductors do not exhibit shot noise.

A rule of thumb is that if the photocurrent from a photodiode is sufficient to drop $2kT/e$ (50 mV at room temperature) across the load resistor, the shot noise dominates the Johnson noise; in a photoconductor, the required current is reduced by a factor of 2 because of the increased shot noise from recombination.

Comparison of shot noise due to signal and background is easier; because both photocarriers generated by signal and background photons exhibit shot noise, the signal shot noise will dominate the background shot noise whenever the signal photocurrent is larger.

3.10.3 Background Fluctuations

In the absence of other modulation, the photon statistics of otherwise noiseless background light are the same as signal light. Background shot noise will dominate the Johnson noise any time that the background photocurrent drops more than $2kT/e$ across the load resistance for photodiodes, or kT/e for photoconductors.

In many instances, the background light is strongly modulated; examples include 120 Hz modulation in room light, and 15.75 kHz from television screens. Furthermore, in many measurements, the coherent background is very important; for example, the ISICL system (a short range coherent lidar) of Example 1.12 encounters unwanted reflections from the plasma chamber walls that are 10^6 time stronger than the desired signal, and which further exhibit strong modulation during scanning due to speckle. A combination of baffles, homodyne interferometry, and laser noise cancellation produce a stable measurement even so.

3.10.4 Thermal Emission

In the mid- and far-IR, say, from 5 to 20 μm , room temperature or thermoelectrically cooled quantum detectors are limited by the Johnson noise of their own shunt resistance, while a cryogenically cooled unit is generally limited by the fluctuations in the background thermal radiation, the so-called *BLIP* condition.[‡] For sufficiently strong signals, the shot noise limit may be reached, but this gets increasingly difficult as the wavelength gets longer, because the decreasing energy per photon makes the shot noise limited SNR higher, and because the thermal background gets stronger. The key design goal while using a BLIP detector is to reduce the detector area and field of view as much as possible, while keeping all the signal photons, and not doing anything silly in the electronics to add significant additional noise. Example 3.2 describes how to calculate the expected noise from an IR photodiode.

3.10.5 Lattice Generation–Recombination Noise

Photoconductors exhibit noise due to random fluctuations in their number of carriers, so-called generation–recombination noise. The last heading dealt with noise due to coupling to the fluctuations in the radiation field; noise also arises from coupling to the

[†]“If it looks like a duck and it quacks like a duck, it’s a duck.”

[‡]BLIP originally stood for “background-limited infrared photoconductor” but has come to be applied to any detector whose performance is limited by the thermal background.

fluctuations of the lattice vibrations. Since this is basically noise in the conductivity, it causes noise mainly in the bias (dark) current. Thus it does not strongly affect photodiodes, which are normally run without a bias current.

For a photoconductor with resistance R , made of a material with majority carrier lifetime τ and mobility μ , with a DC current I flowing in it, the lattice G-R noise voltage v_l is

$$\begin{aligned} v_l &= IR \sqrt{\frac{4\tau B}{N}} \\ &= IR \sqrt{\frac{4\tau\mu BR}{\ell}}, \end{aligned} \quad (3.23)$$

Unfortunately, this calculation depends on parameters that are not readily extracted from most data sheets, such as the lifetime of each carrier species, which is the majority, and so on. In spectral regions (IR) where this is a significant noise source, it is usually necessary to rely on the manufacturer's assertion that a certain detector, at a given temperature and field of view, is BLIP, and go from there. This is unsatisfactory, since it makes it difficult to make trade-offs between, say, a cooled filter with significant passband losses or no filter at all. The cooled filter will be better if the lattice G-R noise and R_{sh} Johnson noise are low enough, whereas no filter is preferable if the signal loss is enough to bring the detector into the G-R or Johnson limits. It is usually best to set a lower limit on the noise using the Johnson noise of the published shunt resistance of the device, and then quiz the manufacturer as to how low the field of view can go before the detector ceases to be BLIP.

3.10.6 Multiplication Noise

APDs and photomultipliers exhibit multiplication noise, which appears as gain fluctuations. In a PMT, the variance of the electron gain at the first dynode is the principal contributor to this noise, while in an APD, the effect is distributed. This noise source is normally specified in the manufacturers' data sheets and must not be overlooked when designing detector subsystems, as it is often the dominant noise source toward the bright end of the system's dynamic range.

3.10.7 Temperature Fluctuations

A small detector weakly coupled to a thermal reservoir exhibits local fluctuations in thermal energy density, which are sometimes described as temperature fluctuations. This is a poor name, since the idea of temperature is well defined only in the limit of large systems, but it is unfortunately entrenched. Thermal fluctuations depend on the thermal mass of the detector, and on the thermal resistance between it and the reservoir, but not directly on the area; this makes it one of the few intrinsic additive noise sources for which D^* is an inappropriate measure. For a detector connected to a reservoir at temperature T through a thermal conductance G , the RMS thermal noise power spectral density is given by

$$\langle \Delta P^2 \rangle = 4kT^2G. \quad (3.24)$$

If the two surfaces are connected by thermal radiation, then for a small temperature difference, the thermal conductance is given by the derivative of the Stefan–Boltzmann

formula (2.4), with an emissivity correction:

$$G_{\text{rad}} = 4\sigma T^3 \frac{\eta_1 \eta_2}{\eta_1 + \eta_2 - \eta_1 \eta_2}, \quad (3.25)$$

so the the fluctuation due to radiation thermal conductance is

$$\langle \Delta P_{\text{rad}}^2 \rangle = 4k\sigma T^5 \frac{\eta_1 \eta_2}{\eta_1 + \eta_2 - \eta_1 \eta_2}. \quad (3.26)$$

For $\eta = 1$ and $T = 300$ K, the thermal conductance due to radiation is about $2 \text{ W/m}^2/\text{K}$, which is equivalent to that of 1.3 cm of still air (0.025 W/m/K). Thus for well-insulated thermal detectors, radiation can be an important source of both thermal forcing and fluctuation noise. Low emissivity surfaces and cooling can help a lot.[†]

3.10.8 Electronic Noise

A good detector can easily be swamped in electronic noise from a poorly designed or poorly matched preamp. It is vitally important to match the detector's characteristics to those of the amplifier if the best noise performance is to be obtained.

One common way in which amplifiers are misapplied is in connecting a simple transimpedance amp to a detector whose shunt capacitance is significant. This leads to a large noise peak near the 3 dB cutoff of the measurement system, which (although data sheets and application notes often describe it as inevitable) is easily avoided with a few circuit tricks. See Section 18.4.4 for details.

3.10.9 Noise Statistics

It is not enough to know the RMS signal-to-noise ratio of a measurement; without knowledge of the noise statistics, it is impossible to know what the effects of noise on a given measurement will be. The noise sources listed in this section are Gaussian, with the exception of most kinds of source noise and some kinds of electronic noise. Section 13.6.12 has a detailed discussion of this and other noise and signal detection issues.

3.11 HACKS

This section contains only optical hacks, but there are a number of circuit tricks listed in Chapters 10, 15, and 18 as well, which should be considered when choosing a detection strategy. It is important to keep the complete subsystem in mind during selection of a detector element and detection strategy.

[†]See, for example, Lynn E. Garn, Fundamental noise limits of thermal detectors. *J. Appl. Phys.* **55**(5), 1243–1250 (March 1, 1984).

3.11.1 Use an Optical Filter

If your measurement is limited by the noise of the background light, it can often be improved by a filter. In mid- and far-infrared systems using cooled detectors, you usually have to cool the filter too, because it emits radiation of its own in its stopbands. Interference filters may present less of a problem; they mostly reflect the stopband light, so the detector may see a reflection of itself and its cold baffles in the out-of-band region. Make sure you mount a room temperature interference filter with its interference coating facing the detector, or else the colored glass backing will radiate IR into your detector. This helps with narrow filters, which drift a long way when cooled.

3.11.2 Reduce the Field of View

In background-limited situations, the background level can often be reduced by limiting the field of view (FOV) of the detector. Typical ways of doing this are descanning the detector in a flying-spot measurement, or by using baffles and spatial filters to reject photons not coming from the volume of interest. In the mid- to far-IR, the situation is complicated by the thermal radiation of the baffles themselves, which must often be cooled in order to afford a signal-to-noise improvement. For BLIP detectors, if the background radiation is isotropic, the (electrical) noise power will scale approximately as the solid angle of the field of view, which is a very worthwhile improvement. To control stray background light and reduce the thermal load on these cold shields, the inside should be black and the outside shiny.

3.11.3 Reduce the Detector Size

As the FOV is reduced, there will come a point at which the background ceases to dominate other noise sources, so that further reductions are no help. If the shot noise of the signal is the next-largest effect, then only the collection of more photons will improve the measurement; most of the time, however, the next-largest effect will be Johnson or lattice G-R noise, which scale as the detector area.

The spatial coherence of the incident light will set a minimum étendue $n^2 A \Omega'$ for the detector, but if this has not yet been reached, it is possible to focus the light more tightly (larger FOV) on a smaller detector. This strategy has the advantage of reducing all the noise sources, while keeping the signal strength constant; as a bonus, smaller detectors tend to be faster and cheaper. The limits on this approach are set by the available detector sizes, by working distance restrictions at higher NA, and by approaching the shot noise level, which does not depend on detector area and FOV.

3.11.4 Tile with Detectors

Gain isn't everything. To pull really weak optical signals out of significant amounts of background light (too much for photon counting), consider using detectors as wallpaper. If your measurement is limited by background noise statistics, and the signal and background have the same spatial, angular, and spectral distribution, then the tricks of reducing detector size or FOV won't help any more. There's only one way forward: collect *all* of the light.

As you increase the detection solid angle, the background noise grows as $\sqrt{\Omega}$ but the signal goes as Ω . Sometimes you can just line a box with photodetectors, such as

CCDs or solar cells, and improve your measurement statistics. In this sort of case, consider especially whether your detector really needs imaging optics. Would a nonimaging concentrator or just putting the detector up close be better?

3.11.5 Cool the Detector

Cooling a silicon photodiode below room temperature doesn't accomplish much, though cooling does reduce the dark current of CCDs quite a bit. In the IR, where we're stuck with narrow bandgap materials, cooling helps detector noise in two ways. The main one is that it reduces leakage by sharply cutting the rate of thermal carrier generation (in photodiodes) or thermionic emission (in photocathodes); this effect is exponential in the temperature. In a photodiode, this leads to an enormous increase in the shunt impedance of the device, which reduces the Johnson noise current as well as the G-R noise.

The other way cooling helps is that the Johnson noise power of a resistor is proportional to its temperature, so that even with a fixed impedance, the noise current goes down; this effect is only linear, and so contributes less to the overall noise reduction. Transistor amplifiers running at room temperature can have noise temperatures as low as 30 K (see Section 18.5.3), so that it is usually unnecessary to cool the amplifier if the detector is run at 77 K (liquid nitrogen) or above.

Cooling schemes divide into thermoelectric (TE) and cryogenic. Neither is free, but TE coolers (TECs) are much cheaper than cryogenic ones. Multistage TECs can achieve trouble-free ΔT s of 130 °C, and single stage ones 60 °C, provided care is taken not to "short-circuit" them with heavy wires or mounts. This is adequate for work at 3.5 μm and shorter, or with strong signals at longer wavelengths.

Getting BLIP performance at $\lambda \gtrsim 5 \mu\text{m}$ requires cryogenic cooling, which is much more involved. For lab use, LN₂ cooling is usually best, because simplest; in a field instrument, where LN₂ is hard to come by, some sort of mechanical cooler, such as a Joule–Thompson or Stirling cycle device, will be needed. Both alternatives are expensive.

In the extreme IR (beyond 20 μm), choices are more limited: often the choice is between a helium cooled extrinsic photoconductor such as Ge:Zn, or a room temperature bolometer or pyroelectric detector.

3.11.6 Reduce the Duty Cycle

The signal-to-noise ratio of a detection scheme can also be improved by concentrating the signal into a shorter time, as in a pulsed measurement with time-gated detection. Assuming the average optical power remains constant, as the the duty cycle[†] d decreases the electrical SNR improves as $1/d$, because the average electrical signal power goes as $1/d$, and the average noise power is constant, because the noise bandwidth goes as d^{-1} but the detection time goes as d . The limit to this is when the shot noise of the signal is reached—see Sections 10.8.2, 13.8.10, and 15.5.6.

3.11.7 Use Coherent Detection

By far the quietest and best-performing signal intensification scheme is coherent detection. It exploits the square-law properties of optical detectors to form the product of the

[†]Duty cycle is the fraction of the time the signal is active: a square wave has a 50% duty cycle.

signal beam with a brighter beam (often called the *local oscillator (LO) beam*, by analogy with a superheterodyne receiver). If the signal and LO beams have time-dependent vector electric fields \mathbf{E}_s and \mathbf{E}_{LO} , respectively, the photocurrent is given by

$$\begin{aligned}
 i(t) &= \mathcal{R} \left\{ \iint_{\text{det}} |\mathbf{E}_{LO}(t) + \mathbf{E}_s(t)| dA \right\} \\
 &= i_{LO} + i_s + 2 \operatorname{Re} \left\{ \iint_{\text{det}} \mathbf{E}_{LO}(t) \mathbf{E}_s^*(t) dA \right\} \\
 &= (DC) + 2R\sqrt{i_{LO}i_s} \iint_{\text{det}} W(t) \cos(\theta(t)) dA,
 \end{aligned} \tag{3.27}$$

where \mathcal{R} is the responsivity, i_s and i_{LO} are the photocurrents generated by the signal and LO beams alone, W is the ratio of the local value of $|\mathbf{E}_{LO}\mathbf{E}_s|$ to its average, and θ is the optical phase difference between the two beams as a function of position. If the two beams are in phase, perfectly aligned, and in the same state of focus and polarization, the integral evaluates to 1, so that the signal photocurrent sees a power gain of (i_{LO}/i_s). The shot noise is dominated by the additive noise of the LO beam, but since the amplification ratio is just equal to the ratio of the LO shot noise to the signal shot noise, the resulting total signal to shot noise ratio is equal to that of the signal beam alone, even with an E_s equivalent to one photon in the measurement time—a remarkable and counterintuitive result. This effect can overcome the Johnson noise of a small load resistor with only a milliwatt or two of LO power. This remains so for arbitrarily weak signal beams, so coherent detection offers an excellent way to escape the Johnson noise limit.

If θ is not 0 everywhere on the detector, the value of the integrals in Eq. (3.27) will be reduced. Even slight differences in angle or focus between the two beams will give rise to fringes, which will dramatically reduce the available amplification, and hence the signal-to-noise ratio. This seriously restricts the field of view of a heterodyne system, which may be undesirable. In some instances this restriction is very useful, as it allows rejection of signals from undesired locations in the sample space.

When the light beams are at exactly the same frequency, this is called homodyne detection, and when their frequencies differ, heterodyne. The SNR for heterodyne detection goes down by a factor of 2 because the signal power is averaged over all ϕ , and the average value of $\cos^2\phi$ is 0.5. Another way of looking at this is that a heterodyne detector receives noise from twice the bandwidth, since an upshifted optical beam gives the same beat frequency as a downshifted one (see Section 13.7.2). Temporal incoherence between the beams will spread the interference term out over a wide bandwidth, reducing the gain available as well (see Section 2.5.3). Examples of the use of this technique are heterodyne confocal microscopes, measuring interferometers, and coherent cw lidars.

3.11.8 Catch the Front Surface Reflection

You can get a signal-to-noise boost, in really tight spots, by arranging the first photodiode at 45° to the incoming beam and putting another one normal to the reflected light, wiring the two in parallel so that their photocurrents add (Figure 3.8). That way, a sufficiently




Figure 3.8. Catching the front-surface reflection from photodiodes. This trick improves detection efficiency and can enormously reduce the effects of etalon fringes in the photodiode windows.

low NA beam has to make three bounces off a photodiode before it can escape. (Using a smaller angle will increase the number of bounces but require bigger photodiodes.)

The front-surface reflection from an uncoated silicon photodiode is about 40%, so this trick can result in a gain of almost 4 dB in electrical signal power (2–4 dB in SNR), a highly worthwhile return from a small investment. With coated photodiodes, the gain will be smaller, but even there, this is a simple and inexpensive way to pick up another several tenths of a decibel in signal strength. Another benefit is that by collecting most of the reflected light, you can greatly reduce the signal strength drifts caused by temperature and wavelength sensitivity of etalon fringes in the photodiode windows. (This effect, which often limits the attainable accuracy of CW optical measurements, is discussed in Section 4.7.2.) Adding a third photodiode to bend the light path out of the page by 90° eliminates polarization dependence completely, and the two extra bounces improve the light trapping even more. This approach is used in high accuracy radiometry.

3.11.9 Watch Background Temperature

In mid- and far-infrared systems, the detector is normally the coldest thing, and so (although its emissivity is very high) its thermal radiation is weak. A convex lens surface facing the detector will reflect a demagnified image of the detector and its surroundings. In an imaging system, this will result in a dark region surrounded by a lighter annulus, the *narcissus effect*, from Ovid’s story of the boy who fell fatally in love with his own reflection. Narcissus is actually a good effect—it becomes a problem when the image of the detector moves, or when its magnification < 1 as in our example, so that it has hot edges and a large spatial variation. Since the detector’s emission is wideband, in general narcissus is not an etalon issue. Good baffles and careful attention to silvering are needed to ensure that radiation from the rest of the optical system doesn’t dominate the detected signal.

Besides narcissus and other instrumental emission, nonuniformity in background temperature can mask weak infrared sources, as skylight does stars. The chopping secondary mirror of Example 10.5 is one way of fixing this.

3.11.10 Form Linear Combinations

Many measurements require addition and subtraction of photocurrents. This is best done by wiring the detectors themselves in series (for subtraction) or parallel (addition). Doing this ensures that both photocurrents see exactly the same circuit strays and amplifier gain and phase shifts; this makes the addition or subtraction extremely accurate and stable, without tweaks (see the subtraction trick of Section 18.6.1) and the differential laser noise canceler of Section 18.6.5.[†] It is a bit confusing at first, but since the far ends of the photodiodes are connected to very low impedance bias points (which are basically AC ground), the series and parallel connections are equivalent for AC purposes; the noise sources and capacitances appear in parallel in both cases.

3.11.11 Use Solar Cells at AC

One problem with good quality silicon photodiodes is their cost per unit area. A 5 mm diameter photodiode can easily run \$100, although some are available more cheaply (down to \$5). There are lots of applications in which more area is better, but cost is a problem. If you have such an application, consider using solar cells. A 25 × 75 mm amorphous silicon solar cell costs \$5 in unit quantity, has a quantum efficiency of 0.5, and responds well throughout the visible. It is very linear at high currents, and surprisingly enough, if you use the cascode transistor trick (Section 18.4.4), you can get 3 dB cutoffs up to 20 kHz or so. Some smaller cells work at 100 kHz. Because of leakage, you can't usually run much reverse bias, so if you're using an NPN cascode transistor with its base and collector at ground potential, bias the solar cell's anode at -0.6 to -1 V. Besides large capacitance and leakage, solar cells have serious nonuniformity—they often have metal stripes across their faces, to reduce lateral voltage drops. On the other hand, for photoelectrons per dollar, you can't beat them.

3.11.12 Make Windowed Photodiodes into Windowless Ones

One good way of avoiding etalon fringes in photodiode windows is to use windowless photodiodes. Many types of metal-can photodiodes can be used without windows, but procuring such devices can be very difficult and expensive in small quantities. For laboratory and evaluation use, it is frequently convenient to remove the windows from ordinary devices. The methods used most are filing or cutting using a lathe. These methods often lead to metal chips or cutting oil being left behind on the die, possibly causing short circuits, scratches, or $1/f$ noise and drift.

A much more convenient and safe method is to use a big ball-peen hammer, although this may seem odd initially. Hold the diode in a vice by the leads, with the base resting on top of the jaws, and tap the glass gently with the peen (the rounded side). It will turn to powder, which can be removed by turning the diode over and tapping it against the side of the vice. The protruding face of the ball makes the blow fall on the glass, but the gentleness of its curvature ensures that it will be stopped by the metal rim of the case before any glass dust is ground into the die.

Because the glass is clean and nonconductive, it does not lead to any long-term degradation of the optical performance of the detector, and because any glass falling

[†]There's lots more on noise cancelers in Philip C. D. Hobbs, *Ultrasensitive laser measurements without tears*. *Appl. Opt.* **36**(4), 903–920 (February 1, 1997).

on the die does so reasonably gently, no scratches result. The only problem with this approach is that not all diodes are adequately passivated for windowless operation. The passivation layer used in ordinary IC chips is usually a thick layer of silica glass produced by a sol-gel process or by sputtering. Because the glass and the chip have very different refractive indices (1.5 vs. 3.4 to 4), it is not so easy to AR coat a diode processed this way, especially since the thickness of the passivation layer may be poorly controlled; for best performance, it may be necessary to AR coat the die, passivate it, and then AR coat the top of the passivation layer. Understandably, this is not often done, so that a diode with excellent performance in a hermetically sealed package may degrade very rapidly when the window is removed. The usual symptoms are a gradually increasing dark current, together with rapidly growing $1/f$ noise and occasional *popcorn* bursts. The otherwise good Hamamatsu S-1722 used in the example is in this class.

Aside: Hermetic Seals. Most ICs and other active devices are packaged in Novolac epoxy. Lots of optoelectronic parts such as LEDs and photodiodes are encapsulated in clear polycarbonate. CCD windows are often glued on with epoxy. All these plastics are great, but there's one thing they aren't: hermetic. Water vapor diffuses readily through plastic. The air inside a CCD package with an epoxied window will respond to humidity changes outside with a time constant of a week or two; this can lead to condensation inside the package in service, especially in cooled setups. If you're building a system with a cooled detector, insist on a glass-to-metal or frit-bonded seal, or else work without a window and fight the dust instead.

3.11.13 Use an LED as a Photodetector

Direct bandgap devices such as GaAs diodes have very steep long-wavelength cutoffs, which can reduce the need for short-pass filters. This can be used to good account in highly cost-sensitive applications, at least when these detectors can be had at low cost. Unfortunately, most such detectors are relatively expensive. One exception is ordinary AlGaAs LEDs. These devices are inefficient as detectors; their quantum efficiency is low, they have small areas, and the optical quality of their packages is extremely poor. Nevertheless, their long-wavelength cutoff is very steep, and it can be selected to some degree by choosing a device of the appropriate emission color and package tint. Where spectral selectivity is needed and every nickel counts, they are sometimes just the thing.

3.11.14 Use an Immersion Lens

Although the minimum étendue cannot be reduced, remember that it contains a factor of n^2 . If you contact a hemisphere of index n to the photodiode, you can reduce its area by n^2 , thereby reducing the capacitance by the same factor and increasing the effective D^* too. This of course works only for n up to the refractive index of the photodiode material, but this is 3.5 for Si and 4 for Ge. Plastic package photodiodes are a good candidate for this, because their indices are similar to glass, so that UV epoxy or index oil can be used easily. Thermoelectrically cooled HgCdTe devices really need this treatment.

3.11.15 Use a Nonimaging Concentrator

The immersion idea can be extended by using a nonimaging concentrator, as shown in Figure 3.9. Repeated bounces off the sides of the cone cause the angle of incidence to




Figure 3.9. Nonimaging concentrator for improving photon collection with a small diode. The simplest kind is the cone concentrator, shown along with its unfolded light path; unfolding is an easy way to see which rays will make it and which won't. Due to near-normal incidence, the bottom of the cone will need silvering.

increase. As shown in the figure, TIR cannot be relied on near the bottom of the cone, so it will probably have to be silvered. Don't silver the whole cone unless you have to, since TIR is more efficient. There are better ways to make a concentrator than this, for example, the compound-parabolic concentrator, which can achieve the thermodynamic limit.

3.11.16 Think Outside the Box

There are fashions in the detector business as elsewhere, and the received wisdom about how to do things is always a mixture of actual engineering experience and “professional judgment.” Sometimes it's right and sometimes it's wrong. For example, most solid state thermal cameras are built in lithographically defined arrays much like CCDs or CMOS imagers. They may be cryogenically cooled (InSb and HgCdTe) or run at room temperature (PZT, lithium tantalate, or microbolometers), but their basic physical outline is an array of small pixels on a solid surface. This has significant effects on the performance and economics of the devices—they tend to cost between \$2000 and \$40,000 and have maximum $NE\Delta T$ of a bit below 0.1 K, in array sizes from 256 to 500,000. The small pixel sizes require well-corrected lenses, which are very expensive in the infrared. Not all applications absolutely need that many pixels, and for those uses, there's an alternative method as shown in Figure 3.10.

This sensor uses large pixels (3×5 mm) made of carbon ink applied to a $9 \mu\text{m}$ film of PVDF by screen printing (T-shirt lithography). The film is freestanding in air, leading to very low thermal conductance. Interestingly, a photon budget shows that the best SNR is achieved by insulating the pixels (which makes them slow but sensitive) and recovering the bandwidth by digital filtering afterwards, as we'll do in Example 17.1. The reason is that the insulation slows the sensor down by *increasing* the low frequency sensitivity, rather than *decreasing* the high frequency sensitivity. The big pixels and low resolution (96 pixels) mean that a simple molded polyethylene Fresnel lens works well. The multiplexer is a little more of a problem, but it turns out that an array of ordinary

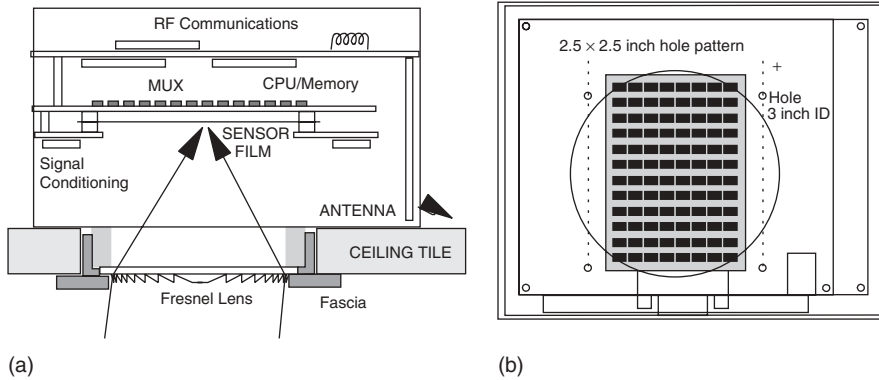


Figure 3.10. Footprints thermal infrared imager: (a) a cross-section view of the development version shows its 50 mm diameter, 0.7 mm thick HDPE Fresnel lens, and $9\ \mu\text{m}$ PVDF sensor film freestanding in air; (b) a front view without the lens shows the 8×12 array of 3×5 mm carbon ink pixels screen-printed on the PVDF.

display LEDs used as switches (see Section 14.6.1), one per pixel and driven by 5 V logic lines, does an excellent job, leading to a sensor costing about \$10, whose NE Δ T is about 0.13 K. There are some other signal processing aspects we'll look at in Section 18.7.2 and Example 17.1.[†]

[†]For more details, see Philip C. D. Hobbs, A \$10 thermal infrared sensor. *Proc. SPIE* **4563**, 42–51 (2001) (<http://electrooptical.net/www/footprints/fpspie11.pdf>) for the gory technical stuff, and Philip C. D. Hobbs, Footprints: a war story. *Opt. Photonics News*, pp. 32–37 (September 2003) (<http://electrooptical.net/www/footprints/fpwaropn.pdf>) for the war story.

Lenses, Prisms, and Mirrors

In theory, theory and practice are the same. In practice, they're different.

—Anonymous

4.1 INTRODUCTION

Although lasers and nonlinear optics get all the attention, most of the real work in an optical system consists of making the beam you've got into the one you want and routing it where you want it. This prosaic but necessary work is done by lenses, prisms, and mirrors (with an occasional cameo appearance by a grating). In this chapter, we discuss the workings of these simple devices, beginning with what they're made of and how they work. The bulk of the chapter is a heuristic treatment of what they actually *do*, and the surprisingly subtle business of how best to combine them so that they do what you want them to.

Designing a complex lens is a highly specialized art using special tools, and is beyond our scope. Using and combining lenses that others have designed, on the other hand, is a very practical skill that everyone in optics should have.

4.2 OPTICAL MATERIALS

4.2.1 Glass

Glass is remarkable stuff. It can be far more transparent than air and stronger than steel. It is useful for almost every optical purpose from lenses to torsion springs.

By and large, glass is a trouble-free material as well. It comes in a very wide variety of types. The main properties we care about in optical glass are its refractive index, its dispersion (i.e., how much the index changes with wavelength), and its transmittance. For historical reasons, these are specified with respect to certain spectral lines of common elements, which were discovered and labeled by Joseph von Fraunhofer in the solar spectrum. The index usually quoted is n_d , at the d line of helium at 587.6 nm. Older references use the slightly redder D lines of sodium, 589.3 ± 0.3 nm, but it doesn't matter much. The dispersion is quoted as N_{FC} ,

$$N_{FC} = \frac{n(486.1 \text{ nm}) - 1}{n(656.3 \text{ nm}) - 1}. \quad (4.1)$$

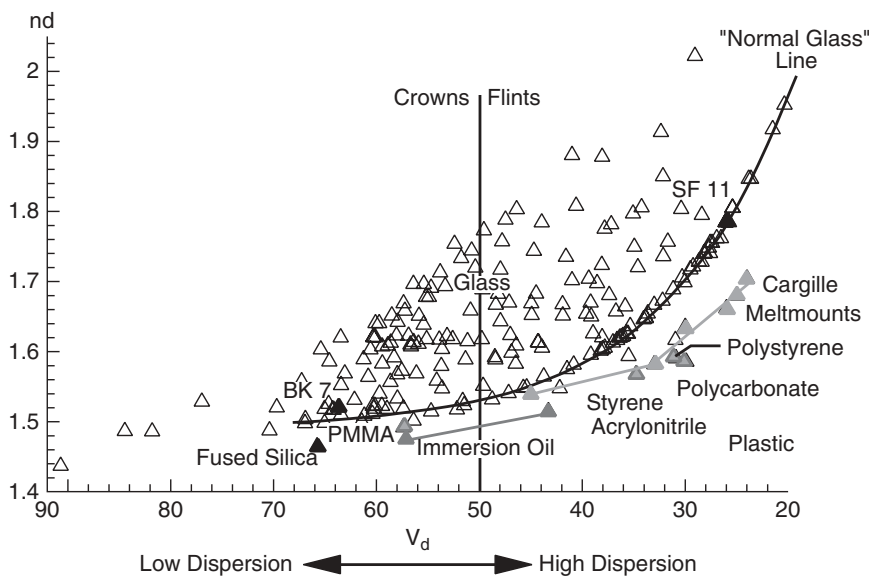


Figure 4.1. Refractive index n_d versus reciprocal dispersive power V_d for Schott optical glasses and common plastics.

The deep red C line at 656.3 nm is the Balmer α line of hydrogen, and the blue-green F line at 486.1 nm is Balmer β . The quantity $n - 1$ governs the power of a glass–air surface, so N_{FC} is the ratio of the powers of a given surface at the F and C lines. The classical way of quoting dispersion is the Abbe number V ,

$$V = \frac{n_d - 1}{n_F - n_C}, \quad (4.2)$$

also called the reciprocal dispersive power (a big V means low dispersion). Figure 4.1 is a plot of the index versus dispersion for the optical glasses manufactured by Schott.

By and large, glass has a large coefficient of dispersion in the visible, which is highly variable among different glass types, but it has low temperature coefficients of index and of expansion.

Optical glasses are traditionally divided into two types: crowns, which have low indices (1.5–1.6) and low dispersion, and flints, which have higher indices and dispersion. The classical distinction was that anything whose V was over 50 was a crown, but the two categories have become blurred over the years as new glass formulations have been developed. The most common optical glass is BK7, a borosilicate crown glass with $n = 1.517$. It is popular because it is inexpensive and works well. Glass prices span a huge range—more than 100 \times .

Glass always has residual optical defects, such as bubbles and striae (long, thin regions of slightly different optical properties). For a critical application such as laser windows, choose a grade of glass whose striae and bubbles are guaranteed to be low enough.

For high quality applications, optical elements are often made of synthetic fused silica, a very pure quartz glass made by a chemical vapor deposition process. Fused quartz, an inferior material, is made by melting natural quartz sand. Fused silica comes in several

grades, differing in the density and type of their bubbles and striae, and in their OH content. The O—H bond absorbs at $1.34\ \mu\text{m}$, $2.2\ \mu\text{m}$, and especially $2.7\ \mu\text{m}$. High-OH fused silica is essentially opaque at $2.7\ \mu\text{m}$. Low-OH fused silica, such as the Infrasil grade from Heraeus Amersil, can be pretty transparent there (depending on its thickness of course).

Fused silica and many types of glass are chemically fairly inert, but different glasses differ significantly. Hostile environments, such as continuous exposure to salt spray, will weather the surface of the glass (and any coatings present), degrading its optical performance. Severely weathered glass may appear to have a whitish crust on top. Fused silica and crown glasses resist weathering quite well, but high index glasses $n \approx 1.8\text{--}2$ are often environmentally delicate, as they have less quartz and more lead oxide and other less inert materials. Some of these will darken and weather over time even in ordinary lab conditions. The trade term for chemical inertness is *stain resistance*, and glasses are specified for it. Corrosion is not always a disaster: Fraunhofer discovered that weathering glass slightly increased its transparency—he correctly guessed that weathering the surface produced a layer of lower refractive index, which reduced the surface reflection. Pickling telescope lenses to improve their transmission was popular throughout the 19th century.

4.2.2 Temperature Coefficients of Optical Materials

The subject of temperature effects on optical elements is somewhat subtle. The temperature coefficients of expansion (CTE) and of refractive index (TCN) are both positive, but CTE is specified in normalized (dimensionless) form, whereas TCN is just $\partial n/\partial T$. The time (phase) delay through a piece of dielectric is $(n\ell)/c$, where ℓ is the length of the optical path. The normalized temperature coefficient, TC_{OPL} , is

$$\text{TC}_{\text{OPL}} = \frac{1}{n\ell} \frac{\partial(n\ell)}{\partial T} = \frac{\text{TCN}}{n} + \text{CTE}. \quad (4.3)$$

For most glass types, TC_{OPL} is approximately $10^{-5}/^\circ\text{C}$ or a bit below. Fused silica has a very low CTE—in the 5×10^{-7} range—but a big TCN, about 9×10^{-6} , so with $n = 1.46$, its TC_{OPL} is 7×10^{-6} . BK7, on the other hand, has a larger CTE, 8×10^{-6} , but a low TCN, only 1.6×10^{-6} ; its TC_{OPL} is 9×10^{-6} . There are a few glasses with negative TCN, such as that used for molded optics by Corning.[†] The only common solid with a negative TCN is magnesium fluoride (MgF_2). Air's TCN at constant pressure is about -1×10^{-6} (see below). Glass used in solid state laser rods is often specially formulated to achieve a zero TC_{OPL} . Etalons sometimes use two materials with opposite signs of TC_{OPL} , their thicknesses chosen so as to athermalize the path length.

The definition of TC_{OPL} here is for the phase delay inside the dielectric, which is relevant for discussion of fringes. In discussions of the temperature coefficients of lenses or of a mixture of dielectric and free space, another temperature coefficient is also useful, that of the differential optical path length through the element,

$$G = \text{TCN} + (n - 1)\text{CTE}. \quad (4.4)$$

[†]Mark A. Fitch, Molded optics: mating precision and mass production. *Photonics Spectra*, October 1991.

4.2.3 Air and Other Gases

The expression (4.4) is actually only an approximation, since it leaves out the effects of the refractive index of air. Air's index is nearly 1, but it has surprisingly large temperature and pressure coefficients. An ideal gas behaves strictly as a collection of isolated molecules, whose molecular susceptibilities are constant. Thus the dielectric susceptibility of an ideal gas is strictly proportional to its density, which in turn can be predicted from the ideal gas law. What this means is that χ (and hence $\epsilon_r - 1$) is proportional to pressure and inversely proportional to temperature,

$$(\epsilon_r - 1) \propto \frac{P}{T}. \quad (4.5)$$

Since $\epsilon_r \approx 1$, a binomial expansion shows that

$$\text{TCN} = -\frac{n-1}{T}, \quad \frac{\partial n}{\partial P} = \frac{n-1}{P}. \quad (4.6)$$

For dry air at $T = 288\text{K}$ (15°C) and $P = 101.325\text{ kPa}$ (1 atm), $n = 1.00028$, so that $\text{TCN} \approx -1.0 \times 10^{-6}/\text{K}$ and $\partial n/\partial P \approx 2.8 \times 10^{-6}/\text{kPa}$. Thus air's TCN is comparable in magnitude to that of BK7. These are small numbers that aren't usually much of a concern, but they become very important in the design of interferometers, especially Fabry–Perot etalons, and in the presence of temperature and pressure gradients. Humidity is a second-order effect, because the water content of moist air is only a couple of percent and the molecular susceptibilities of H_2O , N_2 , and O_2 are similar. Helium has a susceptibility about 8 times less than air, with corresponding decreases in the temperature coefficients. Note also that these are *partial* derivatives—TCN is quoted at constant P , and $\partial n/\partial P$ at constant T .

4.2.4 Optical Plastics

Plastic lenses have become very popular lately. They are lightweight, cheap (at least in high volumes), and can be molded with a mounting surface integral with the lens element, which helps a lot with assembly and alignment. Mass-produced aspheric lenses (perhaps with a diffractive element on one side) make it possible to get good performance with fewer elements.

On the other hand, plastic tends to be less uniform than glass, cannot readily be cemented, and is harder to coat. Plastic also has much larger temperature coefficients of expansion ($\approx 150\text{ ppm}/^\circ\text{C}$) and of refractive index ($\approx 100\text{ ppm}/^\circ\text{C}$) than glass. The most popular plastic used is polymethyl methacrylate (PMMA), sold under the trade names Lucite, Perspex, and Plexiglas. Others are polycarbonate (Lexan), cyclic olefin copolymer (COC, sold as Zeonex and Topas), and CR39, used for eyeglasses.

Plastics don't have the variety of optical properties found in glasses. Their indices range from about 1.44 to 1.6, with some of the newest (and very expensive) ones reaching 1.7. They have a narrow but reasonable range of V , 30 to about 58, so that plastic lenses can be achromatized. They have higher internal Rayleigh scatter due to the high molecular weight of the polymer chains. They are less transparent in both the UV and IR than most glasses and are more vulnerable to solvents and other environmental hazards.

UV exposure is especially damaging to some kinds of plastics, causing them to yellow and craze. This is especially true of the polyethylene used in far-IR Fresnel lenses, as used in automatic porch lights. Thermosets such as CR39 ($n_d = 1.50$, $V = 58$) are about the most durable optical plastics.

4.3 LIGHT TRANSMISSION

Alongside refractive index and dispersion, the transmission properties of a material govern its range of applicability. As a rule of thumb, a wide selection of good materials is available between 300 nm and 3 μm ; there is some choice between 200 nm and 15 μm ; below 200 nm and beyond 15 μm , most materials are strongly absorbing, so we take what we can get.

4.3.1 UV Materials

Optical materials don't go very far into the UV. The absolute champion is lithium fluoride, LiF, which dies at about 120 nm, right around the Lyman α line of hydrogen at 121.6 nm. The fluorides of barium (BaF_2), magnesium (MgF_2), and strontium (SrF_2) are nearly as good and are more practical materials—harder, easier to polish, and less vulnerable to water. Water absorption can destroy the UV transmittance of LiF completely.

UV grade fused silica is useful down to 170 nm, but glass craps out at about 300–350 nm. Many types of materials are damaged by exposure to short wave UV (below 320 nm or so); glass will darken, and plastics yellow and craze. Flashlamps and arc lamps have strong UV emission, so this can easily happen even in visible-light applications.

4.3.2 IR Materials

Optical glass and low-OH fused silica are useful out to 3 μm or so. Beyond there, the choices diminish considerably; the best window materials are semiconductors like silicon and germanium, both of which can be made transparent out to 15 μm or further. These materials have high indices, 3.5 for Si and 4 for Ge. This is great for lenses, because with a high index, large aperture lenses can be made with shallowly sloped surfaces, so that aberrations are minimized. It is less helpful for windows, because Fresnel losses are large, and the huge index mismatch makes AR coatings rather narrowband. These materials are opaque in the visible, which is a pain because all your alignment has to be done blind. (Don't underestimate the difficulty of this if you haven't tried it—you'll have a good deal less hair afterwards.)

There exist lower index materials with excellent IR transmission, but most of them are toxic or water soluble. The best ones are diamond (if you can afford it, it goes from 230 nm to 40 μm with a weak interband absorption at 2.6–6.6 μm), zinc selenide (ZnSe), arsenic trisulfide or triselenide glass (As_2S_3 and As_2Se_3), and sodium chloride (NaCl). Good quality synthetic sapphire and calcium fluoride are also good if you can live with their limitations (mild birefringence for sapphire and sensitivity to thermal shock for CaF_2). Others, such as thallium bromiodide (KRS-5), are sufficiently toxic that only the stout of heart and fastidious of touch should grind and polish them. These materials have the enormous advantage of being transparent in at least part of the visible, which makes using them a great deal easier.

In the far IR, some plastics such as high density polyethylene (HDPE) are reasonable window materials. Their high molecular weight and polycrystalline morphology, which make them highly scattering in the visible, are not a problem in the IR (Rayleigh scattering goes as λ^{-4}). Ordinary polyethylene or PVC (Saran) food wrap makes a good moisture barrier to protect hygroscopic window materials from humidity. These films can be wrapped around the delicate element without affecting its optical properties in the far IR very much (although each type should be tested for absorption before use).

Unlike visible-light optical materials, most IR materials have low dispersion but have huge temperature coefficients of index and expansion compared with glass (silicon's dn/dT is ~ 170 ppm/K, and while As_2S_3 's dn/dT is only about 10 ppm/K, its CTE is 25 ppm/K. Near- and mid-IR absorption depends on molecular vibrational modes.

4.4 SURFACE QUALITY

It isn't just the material that matters, but the surface quality too. Ray bending happens at the surfaces, so they have to be accurate to a fraction of a wavelength to maintain image quality. Total scattered light tends to go as $[4\pi(\text{rms roughness})/\lambda]^2$, so optical surfaces have to be smooth to $\sim \lambda/1000$. The figure error is the deviation of the surface from the specified figure, without regard for small localized errors, which are divided into scratches and digs (small craters or pits). Scratches are usually just that, but digs are often the result of a bubble in the glass having intersected the surface of the element.

A commodity colored glass filter might have a scratch/dig specification of 80/60, which is pretty poor. An indifferent surface polish is 60/40, a good one is 20/10, and a laser quality one (good enough for use inside a laser cavity) is 10/5. The scratch/dig specification largely determines the level of scatter that can occur at the surface and also affects the laser damage threshold and the weathering properties of the glass. It tells something about how numerous and how large the defects can be but is a subjective visual check, not anything that can be easily converted to hard numbers.

Figure error and scattering from scratches and digs are not the only ways that manufacturing variations can mess up your wavefronts. Striae and bubbles in the glass can result in significant wavefront errors, even in a perfectly ground and polished lens, and will produce a certain amount of internal scattering. Rayleigh scattering from the dielectric sets the lower limit.

Lenses are more forgiving of surface errors than mirrors are, because the rays are bent through a smaller angle at the surface. Tipping a mirror by 1° deflects the reflected light through 2° , regardless of incidence angle, whereas tipping a window or a weak lens has almost no effect at all.

An alternative way of saying this is that a figure error of ϵ wavelengths produces a phase error on the order of $\epsilon(n - 1)$ waves in a lens and 2ϵ waves in a mirror at normal incidence. If the error is a tilt or a bend, even this reduced error tends to cancel upon exiting the other side of the lens. If the mirror is operated at grazing incidence, to reproduce the ray-bending effect of the lens, the surface roughness sensitivity is reduced equivalently, because k_z is smaller.

Aside: Surface Error Sensitivity. The effects of roughness or surface error in a high index material such as germanium ($n = 4.0$) can be even worse than in a mirror, when quoted in terms of wavelengths. Since these materials transmit only in the IR, however,

the actual sensitivity in waves per micron of surface error is not much different from glass in the visible.

Small-scale roughness produces a phase front with random phase errors. Since $\exp(i\phi) \approx 1 + i\phi$, this appears as additive noise on the wavefront and produces a scattering pattern that is the Fourier transform of the surface error profile. In Section 13.6.9, we'll see an exactly analogous case in which additive noise produces random phase shifts in signals. For the optical case, compute the deviation from perfect specular reflection. Anything that doesn't wind up in the main beam is scatter, so in the thin object approximation (Section 1.3.9) the total integrated scatter (TIS) is

$$\text{TIS} \equiv P_{\text{scat}}/P_{\text{refl}} = 1 - \exp(-(2k_z \langle \Delta z \rangle)^2). \quad (4.7)$$

Quantities like $\langle \Delta z \rangle$ are implicitly filtered to spatial frequencies below k , since evanescent modes don't reduce the specular reflection.

Lenses are often poorly specified for wavefront error, since most of the time simple lenses would be limited by their aberrations even if their figures were perfect. Mirrors are usually specified with a certain degree of flatness ($\lambda/10$ at 633 nm is typical).

4.5 WINDOWS

Although windows are very simple optical elements, they are unusual in that their purpose is not usually optical at all, but environmental: a window separates two volumes that are incompatible in some way. Ordinary windows in a house or an airplane are examples, as are windows in a vacuum chamber, laser tube, or spectrometer cell. Microscope slides and cover slips are also windows. A window is the simplest lens and the simplest prism.

The trick with windows is to make sure that their environmental function does not detract from the working of the optical system. A beam of light is affected by any surface it encounters, so windows should be specified with the same sort of care as lenses and mirrors. There is a large range of choice among windows: material, thickness, coatings, wedge angle, and surface quality. For work in the visible, in a benign environment, choose windows made of an optical crown glass such as BK7 or its relatives. In a corrosive environment, quartz, fused silica, or sapphire are better choices. Filters sold for cameras, for example, UV or Video filters (not skylight or color correcting ones), are pretty flat, have good multilayer coatings for the visible, and are very cheap. Be sure you buy the multilayer coated ones.

4.5.1 Leading Order Optical Effects

The leading order optical effects of windows are shown in Figure 4.2. Rays entering a window are bent toward the surface normal, so images seen through a window appear closer than they are; a window of thickness d and index n shifts the image a distance $\Delta z = d(1 - 1/n)$ but does not change the magnification. The window thus looks from an imaging point of view like a negative free-space propagation. This effect on images is opposite to its effect on the actual optical phase; because light is slower in the window, the phase is delayed by the presence of the window, as though the propagation distance had *increased*. These effects are very useful in imaging laser interferometers, where they allow a nonzero path difference with identical magnification in the two arms—tuning the laser slightly shifts the phase of the interference pattern.




Figure 4.2. The leading order effect of a window is a shift in image position and a phase delay.

4.5.2 Optical Flats

A window that has been polished very flat ($\lambda/20$ or so) is called an *optical flat*. Flats are used as phase references (e.g. in Newton's rings and Fizeau wedges). Most of them have a small built-in wedge angle of 30 arc minutes or so to reduce etalon effects, but you can get them with two flat parallel surfaces.

4.6 PATHOLOGIES OF OPTICAL ELEMENTS

It is normally possible to get lenses, windows, and mirrors of adequate surface quality in a suitable material. The troubles we get into with optical elements don't come so much from random wavefront wiggles due to fabrication errors, but from aberrations, front-surface reflections, and birefringence. Aberrations we'll talk about beginning in Section 9.2.2, but briefly they are higher-order corrections to the paraxial picture of what optical elements do. As for the other two, let's do birefringence first—it's simpler.

4.6.1 Birefringence

(See Section 6.3.2 for more detail.) Birefringence in good quality lenses and windows comes from material properties, as in sapphire or crystalline quartz, and from stress. The stress can be externally applied as in pressure-vessel windows or lenses with poor mounts (e.g., tight metal set screws). It can also be internal to the glass due to nonuniform cooling from the melt or surface modification such as metallizing. The birefringence produced by stress is $n_{\perp} - n_{\parallel}$ with respect to the direction of the stress:

$$n_{\perp} - n_{\parallel} = K \cdot S \quad (4.8)$$

where S is the stress in N/m^2 (tensile is positive) and K is the piezo-optic coefficient (or stress-optic coefficient). Most optical glasses have piezo-optic coefficients of around $+2 \times 10^{-12} \text{ m}^2/\text{N}$, so that a compressive stress of 200 N on a 1 mm square area will produce a Δn of around 0.0004. A few glasses have negative or near-zero K values, for example, Schott SF57HHT, whose K is two orders of magnitude smaller ($+2 \times 10^{-14} \text{ m}^2/\text{N}$). Note that the K value is wavelength dependent, which matters at these low levels.

Stress birefringence comes up especially when we use tempered glass, in which the residual stress is deliberately kept high in order to force surface cracks to remain closed. Colored glass filters are usually tempered, so that it is dangerous to assume that the polarization of your beam will remain constant going through a glass filter.

Good optical glass is carefully annealed to remove most of the residual stress. Fine-annealed glass, the best commercial grade, has less than 12 nm of residual birefringence in a 100 mm thickness. In a normal window of perhaps 3–5 mm thickness, this is usually negligible, but in a large prism it may not be.

Intrinsic birefringence is encountered in windows made of sapphire and some other crystalline materials. It is often vexing because these windows are usually chosen for a good physical reason—sapphire has unusually high mechanical strength and unsurpassed chemical and thermal inertness. Examples are flow cell windows in liquid particle counters, which may encounter liquids such as hydrofluoric acid (HF) solutions, which rapidly destroy quartz and glass, or windows in the plasma chambers used for reactive ion etching. In such cases, we must be grateful that such windows even exist, and make the best of their optical deficiencies.

Like most common birefringent optical materials, sapphire is *uniaxial* (two of its indices are the same). Its birefringence is fairly small and negative ($n_{\perp} - n_{\parallel} = 0.008$) and its index is around 1.8. The phase difference due to randomly oriented sapphire amounts to a few dozen waves for a typical 3 mm thick window in the visible, which is large enough to be very awkward. If we just need to get a beam in or out near normal incidence, we can use so-called *c-axis normal* sapphire windows, where the optic axis is normal to the surfaces, and light entering near normal incidence sees no birefringence. If this is not possible, we must usually choose optical designs that are insensitive to polarization, or compensate for the polarization error with appropriate wave plates or other means. Should this be unreasonably difficult, it is often possible (though usually painful) to use polarization diversity—doing the same measurement at several different polarizations and choosing the ones that work best. Fortunately, it is now possible to obtain quartz windows with a thin coating of transparent amorphous sapphire,[†] which has the inertness of crystalline sapphire without its pronounced birefringence. There's more on these effects beginning with Section 6.3.1.

4.7 FRINGES

The good news about fringes is that they are very sensitive to many different physical effects—displacement, frequency, temperature, air speed, and so on. Fringes are the basis of a great many highly sensitive measurements, as we saw beginning in Section 1.6. The bad news is that they are very sensitive to many different physical effects. The power of interference fringes to turn small changes into big signals is not limited to the ones we make on purpose, but extends to all the incidental fringes we create by putting things in the way of light as it propagates.

4.7.1 Surface Reflections

All surfaces reflect light. An uncoated glass-to-air surface at normal incidence reflects about 4% of the light reaching it, the proportion generally increasing for higher incidence angles (depending on the polarization). This leads to problems with stray light, etalon fringes, and multiple reflections.

[†]Research Electro-Optics, Inc.

4.7.2 Etalon Fringes

There are lots of different kinds of fringes, associated with the names of Fizeau, Newton, Haidinger, Fabry and Perot, and so on. All are useful in their place, all will occur accidentally, and all have baleful effects on measurement accuracy, independent of nomenclature. The usual informal term for these unloved stripes is *etalon fringes*. The name should not be allowed to conjure up visions of beautiful uniform beams blinking on and off as the length of a carefully aligned Fabry–Perot cavity changes by $\lambda/2$ —the fringes we’re talking about are not pretty, not uniform, and usually not visually obvious.

Fringes arise from the linear superposition of two fields. In Chapter 1, Eq. (1.68) shows that for two beams having the same shape, and whose phase relationship is constant across them, the combined intensity of two beams is

$$i_{AC}|_{\text{peak}} = 2\sqrt{i_{LO}i_{\text{sig}}}, \quad (4.9)$$

where i_{LO} and i_{sig} are the detected photocurrents corresponding to the two beams individually, and i_{AC} is the amplitude of the interference photocurrent (half the peak-to-valley value). In the case of a single window, etalon fringes arise from the interference of fields reflected from its front and back surfaces. Their interference causes large modulations in the total reflectance of the window, which vary extremely rapidly with wavelength and temperature. If the two reflections are comparable in strength, the net reflectance varies from twice the sum of the two (16% in the case of uncoated glass) to zero. (Since $T + R = 1$ for lossless elements, the transmittance changes from 84% to 100%.) The magnitude of this problem is not commonly realized, which is part of what gives optical measurements their reputation as a black art.

Since the size of the fringes depends on the *amplitude* of the stray reflections, it does not decrease as rapidly as you might imagine with multiple reflections. A two-bounce reflection, whose intensity is only 0.16% of the main beam, can cause a 3.2% p-v change in the reflectance, and even a completely invisible five-bounce beam ($i_5/i_0 \approx 10^{-7}$) can manage a p-v change of 6 parts in 10^4 (with a long path length to boot, which multiplies its sensitivity). In a more complicated system, the possibilities go up very rapidly; the number of possible five-bounce beams goes as the fifth power of the number of elements.

Example 4.1: Polarizing Cube. Consider an ordinary 25 mm polarizing beamsplitter cube, made of BK7 glass ($n_d = 1.517$), with a broadband antireflection (BBAR) coating of 1% reflectance. We’ll use it with a HeNe laser at 632.8 nm. As shown in Figure 4.3, if the beam is aligned for maximum interference, the peak-to-valley transmission change is about 4% due to etalon fringes. The cube is 60,000 wavelengths long (120,000 round trip), so it goes from a peak to a valley over a wavelength change of 4 parts in 10^6 —0.0026 nm, or 2 GHz in frequency. If the HeNe has a mode spacing of 1 GHz, then a mode jump can produce a transmission change of as much as 2.8% from this effect alone. The temperature effects are large as well. BK7 has a temperature coefficient of expansion of about $+7 \times 10^{-6}/^\circ\text{C}$, and its TC of index is $+1.6 \times 10^{-6}/^\circ\text{C}$. Thus the TC of optical path length is 8 ppm/ $^\circ\text{C}$. At this level, a 1° temperature change will cause a full cycle fringe shift, and for small changes on the slope of the curve, the TC of transmittance is $\pi(8 \times 10^{-6})(120,000)(4\%) \approx 12\%/^\circ\text{C}$ for one element alone. You see the problem—a 1 millidegree temperature shift will make your measurement drift by 120 ppm. Fortunately, we’re rarely that well aligned, but being anywhere close is enough to cause big problems.




Figure 4.3. Normal-incidence transmission of a polarizing cube for 633 nm HeNe laser light, as a function of temperature.

4.7.3 Getting Rid of Fringes

Since we obviously can't make decent narrowband measurements in the presence of these fringes, we'll just have to get rid of them. Two basic strategies are commonly used: get rid of the fringes altogether, or cause them to smear out and average to some constant value. There is a long list of approaches people use to do these things, because everybody has to do them. None of them works that well, so a highly precise instrument usually relies on a combination of them—wear a belt and suspenders.

Add Strategic Beam Dumps. This sounds like an arms reduction program, but really it's like emergency roof maintenance: put buckets in the living room to catch the leaks. By calculating or finding experimentally where the front-surface reflections go, it is possible to put an efficient beam dump to catch them. See Chapter 5 for more on beam dumps—black paint is *not* usually good enough by itself. The main problem with this is that if you don't catch the stray beams before they hit other optical elements, they multiply in numbers to the point that it is nearly impossible to catch all of them.

Cant the Window. Tipping the element so that the front- and back-surface reflections are offset laterally from one another can be helpful. If the beam is narrow enough that the two will miss each other completely, this is a tolerably complete solution, assuming that there are no higher-order reflections that land on top of each other.

Apply Coatings. Front-surface reflections can be reduced by coating the glass. This is less effective than we would wish, as we saw in the egregious polarizing beamsplitter example above. A really good multilayer coating such as a V-coating can reduce Fresnel reflections to the level of 0.25%, but rarely any further unless the optical materials

are specially chosen for the purpose. Such coatings are normally narrowband, but it's narrowband applications that really need them. V-coating the cube of Example 4.1 would reduce the slope to a mere 3%/K.

Come in at Brewster's Angle. Windows and prisms can be used at Brewster's angle, so that the p -polarized reflections go to zero. With highly polarized light and great care in alignment, the reflections can be reduced to very low levels. To reduce the amplitude reflection coefficient of an air–dielectric interface to less than ϵ requires an angular accuracy of

$$\Delta\theta_1 < \epsilon \frac{2n_1n_2}{n_2 - n_1}, \quad (4.10)$$

so if we require $\epsilon = 10^{-3}$ (reflectivity = 10^{-6}) for a glass–air interface, we find that the incidence angle must be within 9 milliradians, or about 0.5 degree. Since the s -polarization reflection coefficient is

$$r_s|_{\theta_B} = \frac{n_2^2 - n_1^2}{n_2^2 + n_1^2}, \quad (4.11)$$

which is larger than the normal incidence value, the polarization must be really pure. With fused silica, $r_s = 13.8\%$, rising to 27.9% for a flint glass with $n = 1.8$.

The other advantage of Brewster incidence is that the residual surface reflection goes off at a large angle, where it can be controlled easily. For a prism, it is more important to control the internal reflection than the external one, because you usually can't put a strategically placed absorber inside an optical element. Thus if you have to choose, AR coat the entrance face and put the exit face at Brewster's angle. You will have to accept the residual stress in the prism causing slight polarization funnies. Brewster angle incidence is never quite what it promises to be.

Cement Elements Together. Elements whose refractive indices are similar can be cemented together, which reduces the surface reflections. If the match is good, for example a plano convex lens and a prism made of the same glass, or a colored glass filter ($n \approx 1.51$ – 1.55) and a BK7 window, the reflections can be reduced by factors of 100 or more. Another convenient feature is that when the two indices are closely similar, Brewster's angle is 45° , which is often a convenient incidence angle to work at.

Remember that the cement has to be index-matched as well! If the cement has an index midway between the two glass elements, the reflectivities are reduced by an additional factor of 2 on average, but anywhere in between is usually fine. Index oil can be used in lab applications; it comes in standard index increments of 0.002 and can be mixed or temperature controlled to better accuracy than that. Note its high dispersion (see Figure 4.1) and high TCN.

Use Noncollimated Beams. By using a converging or diverging beam, all interfering beams will be in different states of focus. This leads to lots of fringes across the field, so that the average reflectance is stabilized. The average interference term drops only polynomially with defocus, even with Gaussian beams, so more may be needed. With careful design, a reflection can be eliminated by bringing it to a focus at the surface of another element, and placing a dot of India ink or opaque there to catch it—baffle

design in its minimalist essence. Remember the astigmatism, spherical aberration, and chromatic aberration caused by putting windows in nonparallel light. (Note: Parallel includes collimated but is more general; it refers to any place where the image point is at infinity.)

4.7.4 Smearing Fringes Out

Use Canted Windows and Low Spatial Coherence. Sometimes the beam is wide, the element thin, or the allowable tilt small, so the reflections won't miss each other completely. Nonetheless, if the spatial coherence of the beam is low enough, canted windows can still be useful. If the reflections are laterally separated by several times λ/NA , where NA is the sine of the minimum divergence angle of the beam at the given diameter, the fringe contrast will be substantially reduced. This works well for narrowband thermal light such as a mercury line.

Use Time Delays and Low Temporal Coherence. The same trick can be played in time. Fringes can be smeared out by taking advantage of the limits of your beam's temporal coherence, with a window several times $c/(n\Delta\nu)$ thick, so that the different optical frequencies present give rise to fringes of different phase, which average to near zero (for narrowband sources, this will not be your thin delicate window). Do remember the coherence fluctuation problem too.

Modulate the Phase. If the source is too highly coherent, etalon fringes can be reduced by wobbling the window or the beam rapidly, as with rotating wedge prisms or rotating diffusers, so that fringe motion is large and rapid compared to the measurement time. Unless this is done really well, it is less effective than it looks. The strength of the fringes depends on the time autocorrelation of the field at a delay corresponding to the round-trip time. If the phases of the two are being varied continuously but slowly compared to optical frequencies, what we get is fringes moving back and forth.

The kicker is that these fringes won't in general average to zero. For instance, take a triangle wave phase modulation of ± 10.5 cycles, where the unmodulated fields exhibit a bright fringe. Over a modulation cycle, the pattern will traverse 20 bright fringes and 21 dark ones, so that the average fringe amplitude is down by only a factor of 20. If the amplitude changes by 5%, to 10.0 or 11.0 cycles, the average fringe amplitude is 0—assuming that your measurement time is an exact multiple of the modulation period.

Modulate the Frequency. Another way of applying phase modulation is to tune the source wavelength rapidly (e.g., current tuned diode lasers). For modulation slow enough that the entire apparatus sees nearly the same optical frequency, this is nearly the same as the previous case, except that the range of phase shifts attainable is usually lower due to the limited tuning range of inexpensive diode lasers.

If the tuning is fast with respect to the delay between the reflections, the two reflections will be at a different frequency *most of the time*. Since the average frequencies of the two are the same, if the laser is turned on continuously the two instantaneous frequencies have to be the same twice per cycle of modulation (a stopped clock is right twice a day). The autocorrelation thus falls off more slowly than you might hope as the modulation amplitude increases, but nevertheless, this is a very useful trick, especially since by adjusting the modulation frequency and phase, you can put an autocorrelation null at the

position of your worst reflection. It is especially good in reducing another etalon effect: mode hopping from diode lasers used in situations where backscatter is unavoidable. (Gating the laser can also improve matters sometimes.)

Put in a Wedge. Fringes are tamer if the two beams are not exactly collinear. Replacing the parallel surfaces with wedged ones makes sure this will be the case, leading to fringes of higher spatial frequency. These fringes will average out spatially if a large detector or a spatial filter is used. If there is enough space available, or the angle can be made large, a simple baffle will separate the two reflections completely. The key to this trick is to make sure that there is no low-spatial frequency component in the fringe pattern. Interference between two unvignetted Gaussian beams is an excellent example; the integral over the detector of the interference term goes to zero faster than exponentially with angular offset.

Allow the beams to be vignetted, or use less smooth pupil functions (e.g., uniform), and all bets are off. Since the pupil function and the vignetting both enter multiplicatively, they work like a frequency mixer in a superhet radio, moving the energy of your harmless high frequency fringe down toward 0, defeating the spatial averaging. Detector nonuniformity is also a problem here. Nonetheless, if this strategy is carefully applied, rejection on the order of 10^4 (electrical) is easily possible.

4.7.5 Advice

You will note that all of these strategies require care, and that all will become less effective very rapidly as the number of optical surfaces becomes larger, or their spacing smaller. Keep your system as simple as possible and, in an instrument where high stability is needed, be prepared to sacrifice aberration correction, which multiplies the number of surfaces. It is usually preferable to make a measurement at half the spatial resolution with 10^3 times better stability.

In a monochromatic system, eliminate all closely spaced, parallel planar surfaces, cement things together a lot, and use mirrors rather than prisms for beam steering.

4.8 MIRRORS

4.8.1 Plate Beamsplitters

These useful devices are just windows with an enhanced reflection coating on one side and often an antireflection coating on the other. They were once made with a very thin (10 nm or less) coating of metal on a glass plate, but such coatings are very lossy and so are no longer widely used. Good beamsplitters use a half-wave stack coating instead. Their problems are the same as those of windows, and the treatment is similar as well. They are available in a wide range of splitting ratios, from 90:10 to 10:90, with a very rough 50:50 being most common.

Beamsplitters are often used to allow illumination and viewing along the same optical path, as in microscope vertical illuminators. In that case, the light encounters the beamsplitter twice, once in reflection and once in transmission. The optimal efficiency is only 25% and occurs with a 50:50 beamsplitter. This maximum is very flat; 60:40 and 40:60 both give you 24%, and even 80:20 gives you 16%. Thus the poor accuracy of the 50:50 beamsplitter is not much of a worry.

Beamsplitters are *always* polarization sensitive. Polarizing plate beamsplitters are available at wavelengths of common high power pulsed lasers: 694 nm ruby and 1064 nm Nd:YAG. The cement used in polarizing cubes is easily damaged by high peak powers. These rely on multilayer $\lambda/2$ stacks oriented near Brewster's angle, so that one polarization is passed nearly completely and the other nearly completely reflected. In these devices, the transmitted beam is much more thoroughly polarized than the reflected one.

A dielectric plate inserted in a converging wave produces significant spherical aberration and, for off-axis points, astigmatism and coma as well (see Section 9.4.1). Either use the first-surface reflection for imaging, and illuminate in transmission, or put the beamsplitter before the objective lens, where the NA is low.

4.8.2 Pellicles

A pellicle is a plate beamsplitter on a diet. It consists of a 2–5 μm thick membrane (typically made of nitrocellulose) stretched across a very flat metal ring, sometimes coated. A pellicle is sufficiently thin that (at least in transmission) the aberrations it introduces are small enough to ignore. It is surprisingly robust mechanically, providing nothing actually hits the membrane.

Pellicles reduce the drift due to etalon fringes by making the etalon very thin, so that the fringe period is large, and its variation with temperature relatively slow. This works well enough for moderately narrowband sources, such as mercury tubes; with lasers, it may or may not, depending on your accuracy requirements. In fact, with broader band sources, pellicles tend to be a nuisance, as their broad fringes make the transmission nonuniform on the scale of thousands of wave numbers (hundreds of nanometers wavelength). Their main benefit in white-light systems is that they essentially eliminate ghost images due to the two reflecting surfaces.

Pellicles are not very flat—1 wave/cm typically, far worse than a decent plate beamsplitter. What's more, pellicles are very sensitive to vibration and air currents, which make them deform. A deformed or vibrating pellicle will not reflect a faithful replica of the incident wavefront; the transmitted beam is almost unaffected by the vibration but still suffers from the nonuniformity. The reflection from a pellicle is strongly angle dependent, varying from about 16% to 0% with angle and wavelength. Cleaning pellicles is difficult—you obviously can't use compressed air or lens paper, but in addition, nitrocellulose softens in ethanol. You can get away with detergent and deionized water or with isopropanol. As with gratings, it's best not to get pellicles dirty in the first place.

4.8.3 Flat Mirrors

Flat mirrors are conceptually the simplest optical elements available and often are the simplest to use, as well. Excellent quality mirrors are available at low cost, for a wide range of wavelengths, and from many suppliers. There are three main dangers in using them: neglect, leading to underspecifying or careless mounting; worst-case design, which although commendable in most fields, is something of a vice in optical systems, since it leads to very expensive overspecification; and blunders such as thumb prints. Mirrors are often more sensitive to contamination than lenses and more difficult to clean. Mirror coatings are discussed in detail in Section 5.2.

Some situations require high surface quality in mirrors: interferometers, flying-spot systems, and the best quality imaging systems are examples. Even there, however, there

are lots of places in the system where the best mirrors are not needed. Before the two beams of an interferometer have been split, and especially after they have been recombined, the precision required of the mirrors is less than that of the end mirrors of the interferometer arms. Mirrors made of ordinary float glass (as used in domestic windows) are flat to better than 1 wave per cm of diameter. These very inexpensive mirrors are good enough to pipe light into detectors, for example. If there are more than a few mirrors in use, the total photon efficiency starts to drop dramatically if ordinary protected aluminum is used. Depending on how valuable your photons are, you may be much better off buying more expensive mirrors (or at least more expensively coated ones) if you can't simplify your optical system.

4.9 GLASS PRISMS

Glass prisms are used for dispersing light into its color spectrum, but most often for beam bending and image erection, both of which usually involve internal reflections off one or more faces of the prism. These reflections can be of two sorts: total internal reflection (TIR), in which the light hits the face at an incidence angle greater than the critical angle; or an ordinary reflection from a mirror-coated surface. Which of the two is superior depends on the application. TIR prisms are often used because their efficiency is superior to that of any mirror coating (provided the entrance and exit faces of the prism are AR coated sufficiently well, and the TIR face is really clean). Mirror coatings such as silver or aluminum are used in applications where the TIR condition is violated, where the reflecting surface cannot conveniently be kept clean, or where the phase and polarization shifts on total internal reflection are unacceptable. Some of the more common types of glass prism are shown in Figure 4.4.

4.9.1 Right-Angle and Porro Prisms

Right angle prisms are used for bending a beam through roughly 90° as in Figure 4.4a. Their performance is similar to a simple mirror oriented parallel to the hypotenuse of

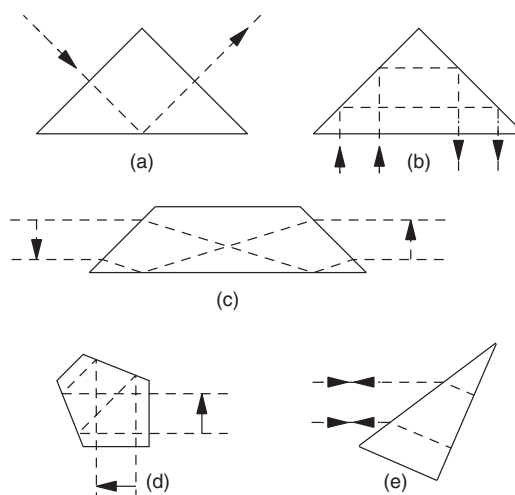


Figure 4.4. Types of glass prisms: (a) and (b) right angle, (c) Dove, (d) penta, and (e) Littrow.

the prism. Light enters normal to one face, bounces off the hypotenuse (either through total internal reflection or by means of a mirror coating), and exits through the other face. This arrangement is commonly used in microscopes, where prisms are more easily aligned and cleaned than mirrors, and where the high efficiency and spectral flatness of TIR reflectors or metallic silver coatings is important. Another advantage of prisms for microscopes is that the bulky optic is mostly inside the optical path, whereas the thick substrate of a mirror is entirely outside it. This makes a prism system mechanically more compact, an important attribute of parts that must slide in and out of the optical path.

The other way of using a right angle prism is shown in Figure 4.4b, which is the typical operating mode of the Porro prism (which is just a gussied-up right angle prism). Here the beam is reflected through 180° in one axis and undeviated in the other. The 180° angle is constant irrespective of rotations of the prism about an axis coming out of the page. This property is purely geometrical—the 180° is made up of two 90° reflections that add; rotating the prism through an angle ϕ will decrease the effect of the first reflection by ϕ while increasing the second one by exactly the same amount. Even the surface refraction cancels out, since the light emerges through the hypotenuse at exactly the same angle it entered at.

Porro prisms have a big enough incidence angle for TIR, so they are usually uncoated. There is a polarization shift on TIR, since the s and p polarizations have different phase shifts. Porro prisms are usually used in L-shaped pairs, one for up–down reversal and one for left–right, so as to produce an erect image. Provided that the prisms are at right angles to one another, s polarization for one bounce is p for the other, so the polarization shift cancels out.

Right angle prisms have one major pathology, which is the one that runs through this chapter (and much of the first part of this book in fact): severe etalon fringes due to the coincidence of the incident light and the front-surface reflections. We all have a natural desire for the beam path to be as simple as possible, without complicated alignment at odd angles. Unfortunately, this desire leads to lots of collimated beams and perpendicular surfaces, which makes for lots of etalon fringes. It is thus in direct conflict with our other natural desire, namely, to have our gizmos work when they're done.

4.9.2 Dove Prisms

The Dove prism is an image rotating prism, which also inverts the image left–right. It is a cylinder of square cross section and has faces cut at Brewster's angle to its central axis. Light propagating along the axis is refracted at the faces, bounces off one side, and exits through the other face, being refracted back to axial propagation in the process. If the prism is rotated about its axis, the image rotates at twice the speed of the prism. Interestingly, the polarization of the light does *not* rotate—it stays more or less the same (see Section 6.2.4 for why). Because of the two symmetrical refractions, there is no *angular* chromatic aberration, as in a simple prism, but there is *lateral*—different colors emerge moving in parallel directions but offset laterally from each other.

4.9.3 Equilateral, Brewster, and Littrow Prisms

An equilateral prism is commonly used for spectral dispersion. For use with polarized light, employment of a Brewster prism, in which the light enters and leaves near Brewster's angle, is generally superior. A Littrow prism Figure 4.4(e) is one in which the light

enters at Brewster's angle and is reflected at normal incidence from the second face of the prism. Light of a certain wavelength thus returns along the path of the incident light, with other wavelengths dispersed on one side or the other. Such a prism is nice because it avoids having the beam direction pass through inconveniently arbitrary angles, and leads to a compact light path with few bends. Such prisms are commonly used as cavity mirrors in argon ion lasers. The laser can oscillate only at the wavelength at which the light path retraces itself. Littrow prisms are not particularly vulnerable to etalon fringes, because the front-surface reflections go off at large angles. The residual external reflection can be got rid of easily with a beam dump, and the internal one controlled with a patch of black wax or other index-matched absorber placed where it hits the prism surface.

There are several types of compound dispersing prisms, of which the Amici prism is representative. It has alternating triangles of high and low dispersion glass cemented together, oriented like the teeth of a bear trap with the low dispersion prisms forming one row of teeth and the high dispersion ones the other. This allows multiplication of the dispersing power without the beam having to go in circles. The cemented construction allows the internal surfaces to work near grazing, for high dispersion, without the large surface reflections and sensitivity to figure errors. Such high dispersion prisms have been superseded almost entirely by diffraction gratings, except in oddball applications where polarization sensitivity or overlap of grating orders is a problem and the linearity of the dispersion is not critical.

4.9.4 Pentaprisms

A pentaprism (Figure 4.4(d)) is an image erecting prism that maintains a constant 90° deviation between incoming and outgoing rays, independent of their incidence angle. The physical basis of this is two reflections in one plane, as in the porro prism. The beam undergoes two reflections from mirrors that are at an accurate 45° to one another (as shown in Figure 4.4d). Unless the prism is made of a very high index material ($n > 2.5$ or so), the steep incidence angle makes TIR operation impossible, so the reflecting faces are normally silvered. The entrance and exit facets are both normal to the beam direction, so pentaprisms produce etalon fringes but don't exhibit much chromatic distortion in parallel light. (See Section 4.10.)

4.9.5 Other Constant-Angle Prisms

There's nothing special about 90° or 180° as far as constant deviation prisms are concerned—the only requirement is to have two reflections from surfaces rigidly held together. To make a 60° constant deviation, for example, you can use a 30-60-90 degree prism. Send the beam in near normal to the short face. It will bounce off the hypotenuse (by TIR) and the long side, then exit through the hypotenuse at normal incidence, deviated by 60° exactly. The incidence angle on the long side is only 30° , so it must be silvered unless $n > 2.0$.

4.9.6 Wedges

A wedge prism is used for performing small deviations (up to perhaps 20°) in the pointing of a beam. Two such identical wedges mounted coaxially and independently rotatable can be used to point a beam anywhere in a cone of 40° half-angle (except for a small

zone around 0° , caused by the inevitable slight mismatch between the angles of the two prisms). This adjustment is considerably more compact and robust than a reflection from two mirrors, but is somewhat less convenient to adjust and (like all refracting devices) more prone to etalon fringes.

The tendency to produce fringes is better controlled in wedge prisms than in most other refracting devices, since the surfaces are not parallel and the surface reflections can often be isolated by applying black wax or other index-matched absorbing material in strategic locations, or merely by making sure that none even of the high-order surface reflections can reenter the main optical path.

4.9.7 Roof Prisms

Most of the standard prism types are occasionally modified by replacing one face with a *roof*—a pair of surfaces at 90° to one another. The effect is identical to cementing a right angle prism, hypotenuse-first, on top of the standard prism: an additional left–right inversion takes place. The most common type is the Amici roof prism, which is a right angle prism with a roof. It has the useful property of reflecting a beam through 90° without inverting it left-to-right. In imaging applications, the roof prism must be made very carefully, because the ridge of the roof appears right in the middle of the field of view; any imperfections will be quite obvious.

4.9.8 Corner Reflectors and Cats' Eyes

A corner reflector (aka corner cube or retroreflector) is a constant 180° deviation prism. These useful devices come in two varieties: hollow ones, which are built up from three flat mirrors accurately perpendicular to one another, and solid ones, which are ground from a block of glass. They have the threefold symmetry of a cube about its body diagonal, but since the beam comes in and out at opposite sides, the optical symmetry is sixfold—looking into the corner cube, you see your own eye cut by six radial segments like an equatorial slice of an orange.

Solid retroreflectors may use TIR or may be coated with metal. The hollow ones tend to work better (stock items are available with 2 arc seconds tolerance, vs. 20 to 50 for solid). Solid ones have poorer transmitted beam quality and suffer from etalon fringes and multiple reflections; those that rely on TIR also cause large polarization shifts. On the other hand, solid retroreflectors are considerably easier to clean, very rugged, and can be used as a vernier adjustment of polarization by rotating them slightly. The polarization changes where the beam crosses a segment boundary, and the shift is big enough to cause weird apodizations in polarizing applications. The phase is also not continuous across the boundary, which will mess up focused or interferometric measurements. The net of all this is that corner cubes work great, but if you want to do anything fancy with the returned beam, it has to fit completely within one of the six 60° orange segments. It follows that the displacement of the returning beam axis has to be at least two beam diameters or so.

A retroreflector successively inverts k_x , k_y , and k_z of the incoming beam on each reflection. The amplitude profile is reflected through the axis of the cube, and $\mathbf{k}_{\text{out}} = -\mathbf{k}_{\text{in}}$. For collimated beams, this is identical to the action of a *cat's eye*—a lens with a mirror surface at its focus. It is different for converging or diverging beams, of course, since the back focus of the lens arrangement is reimaged at the back focal plane, whereas the retroreflector looks like a free-space propagation, so that the light would converge or diverge

considerably before returning. This approximate equivalence is useful in building focused beam interferometers such as the ISICL sensor of Example 1.12. With a lens in one arm and a corner reflector in the other, no fine alignment is necessary, apart from collimation.

4.9.9 Beamsplitter Cubes

A beamsplitter cube works the same way as a plate beamsplitter, except that the reflective coating is deposited on the hypotenuse of a right angle prism, and another one is cemented on top of it, forming a cube with a reflective surface at its diagonal. They are most commonly polarizing, so that one linear polarization is reflected and the other transmitted, similarly to the polarizing type of plate beamsplitter.

Cube beamsplitters are very widely used, much more widely than their merits deserve. The advantages of the cube, namely, no beam deviation and easy mounting, are nowhere near sufficient to make up for the deficiency we saw in Example 4.1: severe etalon fringes in both the transmitted and reflected beams. If you do laser-based measurements, these infernal devices will make your life miserable.

On the other hand, with broadband sources and low spectral resolution, etalon fringes are not normally a problem, so cubes are a good choice. Even with lasers a bit of tweaking can help; if your beams are narrow, canting the cube slightly will help by making the reflections miss each other laterally, but you can't go too far or the polarization quality degrades badly. For experiments where you want good polarization and don't mind tweakiness, a cube mounted on a three-axis tilt table (such as a small prism table) can often be adjusted to give polarization purity of 1 part in 10^5 or even better in the transmitted beam.

Relaxing the requirement for no beam deviation can improve matters quite a lot more. If the faces are polished at an angle of a few degrees to one another, much better control of the reflections can be achieved. Such near-cubes are not catalog products, unfortunately, but can be synthesized for lab purposes by removing the coating from one or two faces and attaching wedge prisms with UV epoxy or index oil. There is only a very small index discontinuity at the surface, so the repolished surface doesn't have to be perfect, making hand work feasible if you don't have an optical shop.

Like all reflection type polarizers, polarizing cubes have problems with the polarization purity of the reflected light; more on that appears in Chapter 6.

4.9.10 Fresnel Rhombs

In Section 1.2.6 we saw that beyond the critical angle, the reflection coefficients of p - and s -polarized light both have magnitude 1 but have different phases, and that the phases depend only on n and θ_i , as given by (1.15),

$$\delta = \delta_s - \delta_p = -2 \arctan \left[\frac{\cos \theta_i \sqrt{\sin^2 \theta_i - n_2^2/n_1^2}}{\sin^2 \theta_i} \right]. \quad (4.12)$$

To get a 90° retardation requires that $n_2/n_1 > 1/(\sqrt{2} - 1) = 2.41$, which while not impossible is inconveniently high. A Fresnel rhomb, shown in Figure 4.5, does the trick by using two TIR reflections from a glass-air interface to produce an achromatic quarter-wave retarder.




Figure 4.5. The Fresnel rhomb functions as an achromatic quarter-wave retarder.

Retarders in general are discussed in Section 6.9, but briefly, a quarter-wave retarder can be used to change linear into circular polarization and back again. Most retarders produce phase shifts by time delaying one polarization with respect to the other, so that the retardation depends strongly on the wavelength. Fresnel rhombs do not, so that apart from material dispersion, their phase retardation is constant with wavelength. This property makes them unique. The two reflections make this another constant-deviation prism— 0° this time. Two rhombs cemented together like a V make an achromatic half-wave retarder with no beam deviation.

The retardation of a rhomb depends mildly on field angle, though less than that of a birefringent retarder (since when one angle goes down the other goes up). The main problem is the the long path length in glass. A 10 mm aperture rhomb made of BK7 has a path length in glass of $(10 \text{ mm})(2 \sec(54^\circ)) = 34 \text{ mm}$, so that material nonuniformities produce retardation variations and phase wiggles across the aperture. Big rhombs are thus less accurate than other retarders for narrowband applications, besides being heavy.

4.10 PRISM PATHOLOGIES

Glass prisms are pretty trouble-free devices. They share the normal problems of any thick piece of dielectric, namely, residual birefringence, material nonuniformity, and etalon fringes. In polychromatic measurements, chromatic effects are also important.

A thick piece of glass is not as uniform as an air space of the same size, so that the waveform quality is poorer. Apart from polarization funnies, a prism has the same effect on an image as a window whose thickness is the length of the optical path inside the prism, unfolded (see Section 3.11.15). If the entrance and exit angles from the prism are not equal, the equivalent window has a wedge angle as well. Reflecting prisms such as pentaprisms and Fresnel rhombs can unfold to a very long path in glass. This means, for example, that a pentaprism placed in a converging or diverging beam will introduce large amounts of spherical aberration if care is not taken.

A big chunk of dispersive dielectric will cause lots of chromatic aberration if either the entrance and exit angles are different or the incident light is not parallel (focused at infinity).

4.11 LENSES

The sign conventions used in lens design are simple, but tend to be hard to remember, because they are completely arbitrary. Here are the four rules.

Sign Conventions in Lens Design

1. The object is at the extreme left of the drawing, and the image at the right (not so for mirrors of course).
2. The radius of a curved surface is positive if it is convex toward the left.
3. Distances along the direction of ray propagation are positive. If the ray would have to back up to get from the object to the lens or from the lens to the image, the distance is negative. Both d_o and d_i are positive when the image is real (true also for mirrors).
4. Phase increases with extra propagation distance; a ray that has to travel further than it should to get to a reference surface has a more positive phase, and so a positive aberration coefficient.

Glass lenses have been used for over a thousand years, since transparent glass became available. The fact that useful lenses could be made in the early days of glassmaking is an excellent indication of their forgiving qualities; for such precise artifacts, lenses are remarkably hard to get wrong. The main danger to the beginner is getting the signs backwards.

A lens images one space (the *object space*) into another, the *image space*. Since light propagation is time-reversal symmetric, lenses work fine backwards too; thus the choice of which is the object and which the image is somewhat arbitrary, so the two are often lumped together as *conjugate points*, or just *conjugates*. We covered the paraxial thin lens case in Section 1.3; here we go into the more general case.

4.11.1 Thin Lenses

The simple lens, usually made of glass and having spherical surfaces, is the most useful basic optical component. Although they are not in themselves adequate as imaging devices, except well into the infrared or at very low numerical aperture, they can be built up into lens systems that perform remarkably well. The simplest approximation to what a lens does is the *thin-lens approximation*, where the total thickness of a lens is vanishingly small. How thin is thin in real life? The thickness of the lens has to be small compared to the depth of focus of the beam you're using,

$$d < \frac{\lambda}{\text{NA}^2}, \quad (4.13)$$

so that not knowing just where the ray bending is really going on doesn't affect the results.

A thin lens is characterized by its focal length f or equivalently by its *power* P , which is $1/f$. If its radii of curvature are r_1 and r_2 , it has a focal length f (in air) given by the so-called *lensmaker's equation*,

$$\frac{1}{f} = (n - 1) \left[\frac{1}{R_1} - \frac{1}{R_2} \right]. \quad (4.14)$$

From the rules in the text box, both radii are measured from the right side of the lens, so that for a biconvex lens r_1 is positive and r_2 negative; thus they both contribute positive power in (4.14). The powers of thin lenses placed in contact add.




Figure 4.6. A thick lens acts like a thin lens, but operating between the *principal planes*. In the thin-lens limit, the principal planes collapse at the center of the lens. (Adapted from Kingslake.)

This approximation is convenient for initial layout of an optical system, but the effects of thickness must be put in before the real optomechanical design is done; the condition (4.13) is extremely stringent, requiring that a 1 mm thick lens of 10 mm diameter have a focal length $f \gg 800$ mm when used with a collimated red HeNe beam.

4.11.2 Thick Lenses

Fortunately this ferocious restriction can be got round easily; Gauss himself recognized that a lens of finite thickness has imaging properties very similar to those of a thin lens, except for the location error, and that this error can be eliminated by splicing in a small section of hyperspace, as shown in Figure 4.6[†].

The planes P_1 and P_2 are the *principal planes* of the lens and intersect the lens axis at the *principal points*. Notionally, a paraxial ray coming from the left is undeviated until it hits the first principal plane P_1 . It is then magically translated to P_2 at exactly the same lateral position (*height*) and then bent as if a thin lens of the same focal length were at P_2 . The focal points F_1 and F_2 are the front and back foci, respectively. The axial distance from the left vertex of the lens to F_1 is the *front focal distance* or *working distance*, and that from the right vertex to F_2 is the *back focal distance* (also confusingly called the *back focus*). These are what you'd measure with a caliper and are tabulated by the lens manufacturer. The back focus is nearly always less than the focal length ($BF < FL$).

If the refractive indices of the media on both sides of the lens are the same, then $f_1 = f_2$. If not, the Lagrange invariant (see Section 9.2.9) can be used to show that $n_1/f_1 = n_2/f_2$.

The lensmaker's equation can be generalized to the case of a single thick lens;

$$P = \frac{1}{f} = (n - 1) \left[\frac{1}{R_1} - \frac{1}{R_2} + \frac{t(n - 1)}{n R_1 R_2} \right], \quad (4.15)$$

where t is the thickness of the lens, measured from vertex to vertex. The front and back focal distances are

$$WD = f_1 \left[1 - \frac{t(n - 1)}{n R_1} \right] \quad (4.16)$$

[†]Rudolf Kingslake, *Lens Design Fundamentals*. Academic Press, Orlando, FL, 1978, p. 49.

and

$$\text{BF} = f_2 \left[1 - \frac{t(n-1)}{nR_2} \right], \quad (4.17)$$

and the separation 2δ between the principal planes is

$$2\delta = t + \text{BF} - \text{WD} = \frac{t(n-1)}{n}. \quad (4.18)$$

It is reasonable to take the center of the lens to be halfway between the front and back foci.

The axial distance from the vertex of a surface to its rim is called the *sagitta* or *sag* for short. The same term is used for an off-axis point, but there it is not the sag of the surface but the sag of that point, so confusion seldom arises. For a spherical surface of radius R and element diameter d ,

$$\text{sag} = R \left[1 - \sqrt{1 - \frac{d^2}{2R^2}} \right] \approx \frac{d^2}{4R}. \quad (4.19)$$

Example 4.2: Biconvex Lens. Consider a 100 mm $f/4$ symmetrical biconvex lens (25 mm diameter), made of BK7. From the lensmaker's equation, $R_1 = 200(1.517 - 1) = 103.4$ mm. Over a 25 mm diameter, each surface will have a sag of about $25^2/(412)$ or 1.5 mm. If we leave 1 mm edge width, then $t \approx 4$ mm. If we take that thickness, then the thick-lens equation gives us (in units of meters and diopters)

$$10 \text{ diopters} = 0.517 \left[\frac{2}{R} + \frac{0.004(0.517)}{1.517R^2} \right], \quad (4.20)$$

where the second term is expected to be a small perturbation. We can either use the quadratic formula or just use the approximate value of 0.1034 m we got before to plug into the correction term; either way, we get $R = 104.7$ mm. Since the refractive index is uncertain at the level of ± 0.002 , and most of the time the tolerance on focal length is a couple of percent, the iterative method works fine. Note also that we have discovered a useful rule of thumb: for glass of $n = 1.5$, the surface radii are equal to f for an equiconvex lens ($f/2$ and ∞ for a plano-convex).

Note that the temperature coefficients of index and of expansion (both positive) fight each other in (4.15); as T increases, the radii, thickness, and index all normally go up. This is in contrast to what we saw earlier for the temperature coefficient of optical path length. Since the net effect can be made positive or negative, it is possible to *athermalize* even a single element lens, so its focal length is nearly constant with temperature. Of course, the mounting and other mechanical parts must be considered in an athermalized design.

A thick lens can easily be put into the $ABCD$ matrix formulation. Consider a thick lens of focal length f whose principle planes are at $\pm\delta$. The $ABCD$ matrix for this is

composed of a thin lens $L(f)$ with a (negative) free-space propagation operator $Z(-\delta)$ on each side:

$$LT(f; \delta) = Z(-\delta)L(f)Z(-\delta) = \begin{bmatrix} 1 + \frac{\delta}{f} & -\left(2\delta - \frac{\delta^2}{f}\right) \\ -\frac{1}{f} & +\frac{\delta}{f} \end{bmatrix} \quad (4.21)$$

One subtle but important point: you might think that the symmetry of the operator form in (4.21) would mean that the lens can be put in backwards without any problem, but that isn't so. For an asymmetric lens, the front and back focal distances are different, so putting the lens in a mount backwards will move the center of the lens, and so cause a focus shift. It also changes the aberrations. This somewhat subtle effect leads to a huge gotcha if the lens is nearly, but not quite, symmetric; the asymmetry may not be immediately obvious, leading to blunders in assembly.

4.11.3 Fast Lenses

A lens with a short focal ratio (diameter/focal length) produces a bright image, due to concentrating light from a large angular range. In photographic applications, this allows a short exposure time, so that the lens is said to be *fast*. Fast lenses bend rays sharply, for which they need highly curved (*steep*) surfaces. Unfortunately, aberrations increase rapidly with increasing incidence angles of the rays; making a fast lens with good image quality is challenging. A rule of thumb to minimize spherical aberration is to minimize the maximum incidence angle of any ray on any surface. Thus when using a single element lens to focus a beam, use a plano-convex one with its flat side toward the focus. Use a double-convex lens for 1:1 imaging, and a meniscus lens in a converging beam.

Aside: Short and Fast. If you're wondering why we bother with *fast* and *short* instead of, say, *large*, it's another one of those historical quirks. A rather slow lens whose focal length is 8 times its diameter is said to be an $f/8$ lens, pronounced "eff eight." The aperture rings of camera lenses just say "8." People who think this is $1/8$ would say that the ratio is small, while those who think it's 8 would say it was large. Everybody knows what *fast* and *short* mean—fast exposures and a short focal length for the lens diameter. Increasing and reducing the aperture with a diaphragm are always known as *opening up* and *stopping down*, also from camera lore; the detentes on an aperture ring are known as *stops* and go in integral powers of $\sqrt{2}$: 1.4, 2, 2.8, 4, 5.6. ... Doubling the aperture would be "opening up 2 stops." Don't confuse this with the other use of *stop*, as in *field stop* and *aperture stop*—this could get muddled, since when stopping down the aperture, you're adjusting the aperture stop.

4.11.4 Lens Bending

The lensmaker's equation shows that the power of a lens can be distributed between the two surfaces by increasing one radius while decreasing the other. This procedure is called *lens bending* and leads to lenses of the sorts shown in Figure 4.7. Lenses of different types have identical paraxial properties, but their aberrations differ when the focal ratio is shorter. Lens bending is the most important degree of freedom in lens design.




Figure 4.7. Lens types: (a) double convex, (b) plano-convex, (c) positive meniscus, (d) double-concave, (e) plano-concave, and (f) negative meniscus.

4.11.5 Dependence of Aberrations on Wavelength and Refractive Index

A single element 1-inch $f/2$ glass lens bent for minimum spherical aberration has about 10 waves RMS error at 588 nm. If it had an index of 4, that would be 1 wave; at $10.6\ \mu\text{m}$, it's 0.05 waves—diffraction limited.[†] The time-delay differences between different components do not change with wavelength, apart from dispersion, but as the wavelength gets longer, these differences become small compared to a wavelength, which is what *diffraction limited* means. Another way of looking at it is that as λ increases, the diffraction spot grows until it eventually dwarfs the geometric errors.

4.11.6 Aspheric Lenses

Because of the limitations of simple spherical lenses, it is natural to consider two possible ways of improving their performance: using them in combination, and relaxing the seemingly astrological constraint of spherical surfaces. Aspheric lenses can indeed perform better than simple spheres, and in low precision applications such as condensers, or large volume applications such as disposable cameras (where the lenses are made by a sophisticated plastic-over-glass molding process), they can be an excellent solution. The plastic-over-glass approach minimizes problems with the temperature coefficient and poor transparency of the plastic. Aspheric lenses are also commonly made by hot-pressing a few waves of asphericity into a glass preform (Corning) and by a sol-gel process based on tetramethyl orthosilicate (TMOS), which can be turned into a gel consisting of pure silica and then cured and baked to make it fully dense (Geltech). One-off custom aspheres are difficult to make and so are too expensive for most purposes. Molded glass aspheres can have good optical performance (e.g., a single element 4 mm 0.55 NA laser collimating lens (Corning 350160) with $\lambda/20$ RMS wavefront error—a Strehl ratio of around 0.95).

It is not really that a particular asphere is so very awkward to fabricate, at least not more so than another one; rather, what is at work is the strong tendency of any surface being polished to become spherical. This tendency is largely responsible for the fact that a small optical polishing shop producing components with surface accuracies measured in nanometers usually looks rather lower tech than an auto garage. The precision comes from the lens grinder's skill, the ease of testing the particular property sought (i.e., focusing), and from the surface's own seeming desire for sphericity.[‡]

Making aspheric lenses or mirrors requires resisting this desire, either by generating the surface with computer numerically controlled (CNC) diamond machining, or by

[†]*McGraw-Hill Encyclopedia of Lasers and Optical Technology*, p. 530.

[‡]Large optical shops nowadays have big surface generating machines that can polish many lenses at once.

nonuniform grinding and polishing, combined with iteration after iteration of polishing and careful interferometric measurement using a precisely made null corrector. Both procedures are expensive, and diamond machining has the additional disadvantage that the tool marks left behind tend to scatter large amounts of light when the element is used with visible light (it is much less of a problem in the IR).

4.11.7 Cylinder Lenses

The most common type of asphere is the cylindrical lens. These are widely available and relatively inexpensive, but their optical quality is notoriously poor. Grinding a lens with one accurately circular cross section and one accurately rectangular one is nontrivial.

Few applications of cylindrical lenses really require high accuracy, fortunately. Cylinders are often used as light buckets, directing light to a slit, as in a spectrograph, or to a linear photodiode array. Their one common, moderate accuracy application is in anamorphic pairs for correcting gross astigmatism or distortion, as in diode laser collimators; a better choice for this application is the combination of anamorphic prisms and controlled misalignment of the collimator. Despite the nice catalog pictures, cylinder lenses are lousy, so don't design systems requiring accurate ones.

4.12 COMPLEX LENSES

4.12.1 Achromats and Apochromats

For nearly all dielectric materials at nearly all wavelengths, the dispersion coefficients are positive; that is, n increases as λ decreases. The only exceptions are deep within absorption bands, where you won't want to use the stuff anyway. Thus it is not possible to color correct merely by sandwiching two plates of opposite dispersion.

On the other hand, you can color correct lenses by putting a positive lens of a high dispersion material next to a negative lens of low dispersion material, or vice versa. Let's take the positive lens case. A powerful positive lens made from crown glass next to a weaker negative lens made from flint glass produces a weakened but still positive two-element lens. If the powers of the two elements are adjusted correctly, then as we go to shorter λ , the increasing positive power is balanced by the increasing negative power, so that the net power of the combination is nearly constant. With two elements, the lens can be color corrected at two wavelengths and is called an *achromat*. Exploiting the different shapes of the dispersion curves of different glasses, color correction can be extended to three or more wavelengths, giving much lower average error over the wavelength band of interest; such a lens is called an *apochromat*.

A side benefit of needing two lenses to get color correction is that the extra degrees of freedom can be used to improve the monochromatic aberrations as well; a commercial achromat has so much less spherical aberration that its wavefront error will generally be 10 or more times better than a single-element lens of the same diameter and focal length.

Example 4.3: Achromatic Doublet. Suppose we want to make an achromatic 200 mm $f/8$ lens, corrected so that the F and C wavelengths come to a common focus. We'll use BK7 ($n_d = 1.51673$, $n_F = 1.52224$, $n_C = 1.51432$, so $N_{FC} = 1.01539$ and $V = 65.24$) for the front element, and SF11 ($n_d = 1.78446$, $n_F = 1.80645$, $n_C = 1.77599$, so $N_{FC} =$

1.03925 and $V = 25.75$) for the rear. The crown glass is usually more durable than the flint, so it is put on the exposed side unless there is a compelling reason not to. Real lens designers do this by exact ray tracing. We'll do it paraxially with the lensmaker's equation:

$$\begin{aligned} P_{\text{tot}} &= P_1 + P_2 = (n_1 - 1) \left(\frac{1}{R_1} - \frac{1}{R_2} \right) + (n_2 - 1) \left(\frac{1}{R_3} - \frac{1}{R_4} \right) \\ &= P_{1d} \frac{(n_1 - 1)}{(n_{1d} - 1)} + P_{2d} \frac{(n_2 - 1)}{(n_{2d} - 1)}, \end{aligned} \quad (4.22)$$

where R_1 and R_2 are the surface radii of the first element, and R_3 and R_4 are those of the second element.

The achromatic condition requires $P_F = P_C$, which leads to

$$\frac{P_{1d}}{P_{2d}} = -\frac{V_1}{V_2} = -2.534 \quad (4.23)$$

and hence $P_{1d} = 1.652P$, $P_{2d} = -0.652P$. It is slightly simpler to express this in terms of the ratio of P_{1C}/P_{2C} and N_{FC} . For this lens, we require a positive element 1.65 times stronger than the combination, and a negative element 0.65 times as strong. We haven't specified anything about the bending of this lens, so we can put most of the power in the buried surfaces; nevertheless, it is difficult to make a cemented achromat faster than $f/1.4$ this way. It is possible to distribute the chromatic correction across the groups of a more complicated design to achieve good color correction at short focal ratios.

4.12.2 Camera Lenses

Camera lenses are a wonderful and inexpensive resource for optical instrument builders. For a hundred bucks or so, you get a very well corrected lens, preadjusted, mounted, and tested, on its own focusing ring and with an aperture diaphragm. Mating bayonet mounts are available to attach them to your system.

Bargains like that are not common in optics, so take advantage of it while it lasts. Camera lenses tend to have lots of elements, so their stray light and etalon fringe performance is not as good as simpler designs. Use the slower, fixed focal length designs rather than the superduper fast ones, the extreme wide angles, or the zooms; they have better image quality and a lot fewer elements to scatter and make fringes. Ordinary camera lenses are best for distant objects; the macrolenses are better for magnifications of 1:10 to 1:1. For higher magnifications, turn the lens around—for 1:1 to 10:1, use a macrolens backwards, and for magnifications greater than 10 \times , use an ordinary lens backwards (watch out for the soft coatings on the rear elements—they're not as durable as the hard front-surface coatings). For large magnifications, you can also use enlarger lenses, which are somewhat cheaper.

Camera lenses are often described as having, say, 13 elements in 6 groups (about typical for a zoom lens in 35 mm format). That means that there are 13 bits of glass, but that some of them are cemented together, leaving only $6 \times 2 = 12$ air-glass surfaces. Since the air-glass surfaces scatter more light, this is worth knowing; a 13-element lens with 13 groups would have 26 surfaces, so its surface reflections would likely be much

worse. The etalon fringes in a lens like this would daunt the bravest, but even in a white-light application with good coatings, around a quarter of the total light would be bouncing around inside the lens barrel, much of it eventually arriving at the image plane to reduce the contrast and introduce artifacts. The 13/6 lens would be about half that bad.

4.12.3 Microscope Objectives

A lot of optical instruments can be cobbled together from microscope parts. Microscopes are high image quality systems built in modular fashion. Their optical trains must be designed to support this use, and those are just the qualities we want in instrument prototypes. Usually we just use the objectives, but sometimes an eyepiece or a trinocular head is useful too—you can't align what you can't see (see Section 11.8).

Microscope objectives are specified by magnification and numerical aperture. To find out what the actual focal length is, divide the magnification into the *tube length*, which is nearly always 200 mm, or 160 mm for old designs. Thus a modern 20 \times , 0.5 NA objective normally has a focal length of 10 mm. The working distance will be significantly less than this, which is sometimes very inconvenient in instruments; we often have to get other things in between the sample and the lens. Since this gets worse with shorter focal lengths, and since the NA and not f controls the resolution, 20 \times , 0.5 NA is the most all-round useful microscope objective for instruments.

Long working distance objectives are available; the longest ones come from Mitutoyo and are big enough to use as anti-aircraft shells. They get rapidly more expensive as the aperture and working distance increase.

Some high NA microscope lenses come with an adjustable correction for cover glass thickness, which dials in a controllable amount of spherical aberration to correct for that introduced by the cover slip. This can be useful in other situations as well.

Microscope lenses exhibit severe longitudinal chromatic aberration; different colors come to focus at different depths. This is a trade-off based on the characteristics of the human visual system, which has poorer spatial resolution in the blue, but is obnoxious in some instrument applications, such as white light scanning microscopy. For such applications, and for use in the UV or IR, where microscope lenses are tough to get, you can use an all-mirror microscope objective, the *Schwarzschild objective*.

In choosing a microscope objective, know what it is you need. For applications not requiring the highest quality, such as the condenser side of a spatial filter, use a commodity objective such as the cheap ones from American Optical, Newport, Swift, or several others. For imaging, or on the collection side of a spatial filter, a good objective such as a Nikon, Olympus, Reichert, Leitz, or Zeiss will work much better. Japanese objectives tend to be nicely corrected in the objective itself, which makes them useful for other purposes where you don't want the microscope attached.

4.12.4 Infinity Correction

A lens designed to operate with its image at infinity is said to be *infinity corrected*. Most modern microscope objectives are infinity corrected, because the resulting parallel light exhibits no chromatic errors when being piped through bending prisms, and the system aberrations do not depend strongly on where subsequent optical components are located. These properties make infinity corrected lenses extremely useful in building instruments.

Camera lenses run backwards are another example of infinity correction—most of them are designed to have the object at infinity, but by time reversal symmetry, this is

exactly equivalent to having the object at the film plane and the image at infinity. A pair of camera lenses operated nose to nose makes a good transfer lens, for example, to image the center of one scan mirror onto the center of another one, or to image an acousto-optic cell at the pupil of a microscope objective to make a scanning microscope. Note that the antireflection coating on the back of the lens is often much softer than the one on the front, and so much more easily damaged in cleaning. Similarly, the glass itself is often more delicate, and of course all the mechanical works are exposed to damage or corrosion.

4.12.5 Focusing Mirrors

Curved mirrors can do nearly anything lenses can but have a completely different set of design trade-offs. A focusing mirror causes much larger ray bending than a lens of the same focal length; this makes it much more sensitive to surface inaccuracies and misalignment than a lens, but also can lead to a more compact optical system, due to the opportunity for folding the light path. Folding leads to two problems, however; obscuration, as mirrors partially shadow one another, and light leakage, as far off-axis light can often get into the detector without having traversed the entire optical system. Baffles can eliminate leakage, but eliminating obscuration requires the use of off-axis aspheric mirrors, which are very difficult to align among other faults.

Mirrors exhibit no dispersion, so chromatic aberration is eliminated in all-mirror systems; the expected improvement in optical quality is not always realized, since it is much easier to make multielement lens systems than multielement mirror systems. On the other hand, for UV and IR use, it is delightfully easy to be able to focus the system with a HeNe laser and be sure that it will be perfectly focused at 335 nm or 1.06 μm . Most lenses have poorly corrected longitudinal chromatic aberration, so that different wavelengths come to focus at slightly different depths. Where spatial filtering is used with broadband illumination (e.g., real-time confocal scanning optical microscopy, where a disc full of pinholes is spun to make an array of moving spots), longitudinal chromatic is extremely objectionable, so mirror systems make a lot of sense.

Because there are no transparent surfaces in an all-mirror system, there is no opportunity for etalon fringes to form, which can be a very important advantage.

Where there are fewer surfaces, there are fewer degrees of freedom to optimize, and mirrors are much more seriously affected by surface errors than lenses are. For these reasons, aspheric mirrors are much more common than aspheric lenses. Overall, mirrors are wonderful in special situations such as working with invisible light, but there is nothing like an optical system based on focusing mirrors to make you appreciate the forgiving qualities of lenses.

Aside: Off-Axis Performance of Focusing Mirrors. Fast mirrors have amazingly bad oblique aberrations, and off-axis ones are the worst. For example, a 25 mm diameter, $f/1$, 90° off-axis paraboloid exhibits a spot size of approximately 30% of the off-axis distance—if you go 100 microns from the axis, the spot grows from the diffraction limit to 30 μm diameter. You really can't run any sort of field angle at all with those things.

4.12.6 Anamorphic Systems

An anamorphic system is one whose magnifications in x and y differ. The main uses of these are to correct for perspective distortion caused by oblique imaging, as in a picture

of the ground taken from the side of an aeroplane, and to circularize elliptical beams from diode lasers. There are two main types: prism or grating systems, and telescopes made from cylindrical lenses. When a beam encounters a surface, its edges define an illuminated patch. If the beam comes in near grazing incidence, its illuminated patch will be greatly elongated. The law of reflection guarantees that a reflected beam will have the same shape as the incident one, but if the beam is refracted or diffracted at the surface, this is no longer the case. On refraction, a beam entering near grazing incidence will leave near the critical angle, elongated. Two prisms are often used together, with the second one bending the beam back to its original propagation direction. On diffraction, a beam entering near grazing can be made to leave near normal, so that a properly chosen grating can substitute for a 90° folding mirror. This idea is used in commercial beam circularizers based on anamorphic prisms.

Cylindrical telescopes can easily be made in any desired magnification, can be slightly defocused in order to correct astigmatism, and do not offset the beam axis, but this is as far as their advantages go. It is a great deal easier to make good prisms than good cylindrical lenses, and the astigmatism correction can be done by mildly misaligning the collimating lens.

4.12.7 Fringe Diagrams

Fringe diagrams are very nearly useless for good quality ($< \lambda/4$) optics. A good quality optical component will have a surface accuracy of a small fraction of a wavelength. Such a small deviation produces only very small irregularities in the spacing and direction of the fuzzy fringes in the diagram. Localized errors are reasonably easily spotted, but global ones, such as astigmatism, are very hard to see by eye, especially since the strong visual asymmetry caused by the fringes running in one direction masks asymmetric errors—you can see a $\lambda/4$ kink by sighting along the fringes, but good luck seeing a $\lambda/4$ spacing error spread out over 20 periods. For evaluating the quality of an optical component, use a phase shifting measuring interferometer if at all possible. Failing that, a Foucault knife-edge test using a razor blade is good for qualitative work, for example, examining a batch of lenses for bad units. If you really need to get quantitative data from a hard copy fringe diagram, either scan it into a computer and digest it in software, or use a parallel rule and a pencil to construct the axes of the fringes, and measure their straightness and separation. It really would be much more useful if optics vendors would ship their fringe diagrams on CD or make them available for downloading, but the author is not hanging about waiting for the day.

4.13 OTHER LENS-LIKE DEVICES

There are other devices besides lenses and curved mirrors that can make beams converge or diverge. These are based on refractive index gradients or on diffraction.

4.13.1 GRIN Lenses

Section 9.2.6 shows how a refractive index gradient causes light to bend. It turns out that a circular cylinder whose index decreases parabolically moving out from the core works like a lens. The ray bending happens continuously throughout the bulk, rather




Figure 4.8. A GRIN lens exhibits periodic foci along its length.

than happening abruptly at the surface, as shown in Figure 4.8. Such a device is called a graded-index (GRIN) lens, or *GRIN rod*. Because the bending happens continuously, a length of GRIN rod exhibits periodically spaced images down its axis, alternating between erect and inverted. If it is cut to an even number of half-periods, a point on one surface is imaged onto the other surface; if it is a quarter period shorter, the image is at infinity. Thus making a GRIN lens of any desired focal length (up to a quarter of a period) is just a matter of cutting the rod to the right length.

At one time, GRIN lenses were quite poor—good enough for coupling light in and out of fibers but not for imaging. Recently, they have been developed to the point where their imaging properties are quite respectable. Fiber coupling is probably still the most common application, but borescopes and endoscopes are now often made from long GRIN rods instead of many sets of relay lenses. Besides simplicity and robustness, GRIN rods avoid the accumulation of field curvature that plagues designs with cascaded relays (e.g., periscopes).

Another approach to using index gradients is to fuse together slabs of optical glass of slightly different index.[†] When a spherical surface is ground into the high index side, the power of the margins of the lens is automatically weakened by the gradual decrease of n . This weakening can be chosen so as to cancel the spherical aberrations of the surface, and so aspheric-quality images can be obtained with a single spherical element.

Aside: Birefringence of GRIN Rods. GRIN rods are usually made by diffusing dopants in from the outside of a plain glass rod. This results in an axially symmetric pattern of residual mechanical stress, so that GRIN rods are actually birefringent, with a pattern not found in nature. This is sometimes important.

4.13.2 Fresnel Zone Plates, Diffractive Lenses, and Holographic Optical Elements

Recently, as a result of improvements in optical modeling software and in the molding of plastic and sol-gel glass, it has become possible to fabricate not only aspheric lenses but lenses with diffractive properties: for example, an aspheric lens with a phase grating on the other surface (zone plates and general holographic elements are discussed in Chapter 7). These diffractive lenses can have unique properties. Although the power

[†]Gradium glass, made by LightPath Technologies, Albuquerque, NM.

of the refractive surface decreases as the wavelength increases, that of the diffractive element increases; thus it can be used to cancel the chromatic aberration of the refractive surface, resulting in a novel element, an achromatic singlet. The possibilities inherent in such a capability have only begun to be assimilated, and such elements should be considered any time more than 10,000 units are needed. Before becoming too breathless with excitement, however, remember the drawbacks: plastics have huge temperature coefficients of index and of thermal expansion; getting good aberration correction over a large field is very difficult with only two surfaces of a low index material; and you absolutely must have very high diffraction efficiency, which is very difficult to maintain over a wide bandwidth (e.g., the visible spectrum). Also see Section 7.9.9 for reasons to keep the diffractive power small. The effective V number of a DOE can be found from the grating equation: $V = \lambda_d / (\lambda_f - \lambda_c) = -3.452$.

4.13.3 Fresnel Lenses

A mirror does all its work at the surface—all that stuff underneath is only there to keep the surface accurately in place. You can make a lightweight mirror by removing unnecessary material from the blank. This isn't quite as easy to do with a lens, since the material in the middle is there to preserve phase coherence by making all the rays arrive with the same delay. For crude light bucket applications, however, an idea of the same sort leads to the *Fresnel lens*. A Fresnel lens is a whole bunch of concentric annular lenses, with the same focal length, as shown in Figure 4.9. Because of all the sharp edges, there's lots of scatter, and because of the loss of phase coherence, the image quality is very poor.[†] Fresnel lenses can't normally be coated, either. Thus their efficiency is poor—as low as 50% in some cases.

On the other hand, they are light and compact, and nonimaging applications such as condensers aren't sensitive to their poor image quality. Projectors are the largest use of Fresnel lenses, but you can make a solar furnace this way, too—up to 2500 times solar concentration has been demonstrated. Off-axis Fresnel lenses are available, or can be made easily from a big one with snips or a razor knife.

When using a fast Fresnel lens, make sure to put the side with the ridges toward the more distant conjugate. Otherwise, the outer rings will exhibit TIR, and no light will get through them. Really steep conventional lenses exhibit this behavior too. The image quality of a Fresnel lens, nothing much to start with, gets dramatically worse with increasing field angle. (This can often be helped by shifting the aperture stop well out ahead of the Fresnel lens.)

4.13.4 Microlens Arrays

Lenses are sometimes used in arrays, to produce an array of very small, poor images. Microlenses as small as 10 μm in diameter are often deposited on top of interline transfer CCDs, to corral light into the small active area of each pixel. Somewhat larger lenses can be used to measure wavefront tilt as a function of position (the Shack–Hartmann technique), from which the wavefront can be reconstructed, more or less. (Even in thermal light, where the etalon fringes aren't so bad, there are never enough pixels per microlens

[†]Since there's no point knocking oneself out in a lost cause, Fresnel lenses are also manufactured to very loose tolerances.

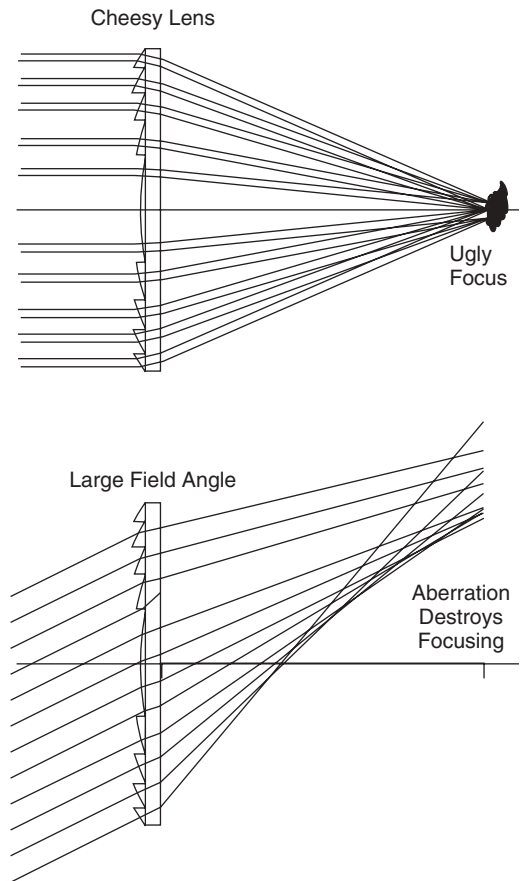


Figure 4.9. Fresnel lens.

to do a decent job of the reconstruction, unfortunately; the Shack–Hartmann measures the local slope of the wavefront, so to reconstruct it you have to extrapolate, which is always fraught with problems.)

Another interesting class of microlens array applications relies on the moiré pattern between two microlens arrays of slightly different pitch, which can synthesize the equivalent of a single short-focus lens. Standard microlens products are available (e.g., from WaveFront Sciences). They are typically 0.2 mm to a few millimeters in diameter, with focal lengths of 1–100 mm. They are generally plano-convex singlets, of course, and so their numerical apertures are limited even given their small Fresnel numbers. The existence of standard products makes microlens arrays good candidates for building into real systems.

4.13.5 Axicons

In addition to flat surfaces and spheres, there exists a class of conical prisms called *axicons*, as shown in Figure 4.10. They are generally made by single-point diamond turning, because otherwise it's difficult to get an accurate surface. Typical uses of an axicon are




Figure 4.10. An axicon converts between filled and annular beams, or between a collimated beam and a J_0 Bessel beam.

sending a laser beam through a Schwarzschild (Cassegrain) microscope objective without hitting the central obstruction, turning annular beams from unstable laser resonators into uniform beams, and, less respectably, making J_0 Bessel beams (misnamed “nondiffracting”) from uniform ones. Aligning axicons is very fiddly, especially in angle. The cone beam dump of Section 5.6.10 is also an axicon.

Coatings, Filters, and Surface Finishes

Bash to fit, file to hide, and paint to cover.

—Anonymous

5.1 INTRODUCTION

An optical element is just a chunk of stuff put in the way of a light beam. Nearly all the action happens right at the surface, which means that controlling the strength and path of surface reflected and transmitted waves is most of optics. In this schematic view, Chapter 4 is about controlling the path, and this one is about controlling the strength.

The jumping-off point for the discussion of coatings is a more detailed consideration of the Fresnel formulas of Section 1.2.4. From there, we can develop a simple way of calculating the behavior of an arbitrary plane-parallel coating stack.

Besides lenses and mirrors, optical systems use white surfaces, for diffusion, and black ones, for stray light control.

5.1.1 Refraction and Reflection at an Interface

We saw in Section 1.2.4 that the Fresnel formulas (1.8)–(1.11) predict the amplitude and phase of reflected and transmitted plane waves at planar interfaces. Here we'll go into a bit more detail about their behavior. Figures 5.1 and 5.2 show the magnitude and phase of the reflection coefficients at planar interfaces between several pairs of lossless dielectrics, as a function of incidence angle in the higher index medium. It isn't usually plotted this way, because it makes a complicated picture, but there's some useful physics here.

Below the critical angle θ_C , reflections at interfaces between lossless isotropic dielectrics always have phase angles of 0 or π , so that linear polarization stays linear after encountering such a surface. Above there, the phase goes all over the place, as you can see; the polarization change on TIR depends on the phase difference δ between s and p , so keep careful track of polarization when using TIR prisms. (This isn't all bad; in Section 4.9.10 we use this effect for polarization control with Fresnel rhombs.) Note especially that r_p goes negative between the Brewster angle θ_B and θ_C .




Figure 5.1. Modulus of the reflection coefficients $|r_p|$ and $|r_s|$ at a dielectric boundary, for several choices of material. The light is assumed to propagate from higher n to lower.




Figure 5.2. Phase of the reflection coefficients r_p and r_s at a dielectric boundary, for several choices of material. The light is assumed to propagate from higher n to lower.

5.2 METAL MIRRORS

5.2.1 Lossy Media

The most familiar mirror coatings are metals such as silver, aluminum, and gold. From an optical point of view, the distinguishing feature of a metal is the very large imaginary part of its refractive index, that is, its extremely high loss. It is slightly paradoxical at first blush, but the highest reflectance coatings are made from the highest loss materials. (See Figure 5.3)




Figure 5.3. Theoretical normal incidence reflectance of gold, silver, and aluminum mirrors as a function of wavelength, for external and internal ($n_{\text{glass}} = 1.52$) reflections.

The Fresnel formulas in terms of θ_i are also valid for absorbing media, even metals. Aluminum has an index of about $0.76 + i5.5^\dagger$ for green mercury light (546 nm), so a clean air–aluminum interface at normal incidence has a reflection coefficient of

$$r = \frac{n_2 - n_1}{n_2 + n_1} = \frac{-0.24 + i5.5}{1.76 + i5.5} = 0.953 \angle 160^\circ, \quad (5.1)$$

and so the reflectivity $R = |r|^2 = 0.91$. An internal reflection (e.g., an aluminized BK7 glass prism) has an even lower R of 0.87.

From a designer’s point of view, just bouncing a light beam once from an aluminum mirror costs 0.8 dB in detected (electrical) signal power for a first surface mirror and 1.2 dB for Al–glass. You can’t do that too many times and still have a signal; this is one of the reasons for the popularity of folding prisms based on TIR, especially in complicated optical systems like microscopes. Since n_1 is real, the reflection coefficient is maximized for a given $|n_2|$ when n_2 is purely imaginary. Metals work a whole lot better in the IR.

5.2.2 How Thick Does the Metal Have to Be?

We saw that metals make reasonable although not fabulous mirrors. Since we usually don’t want to make mirrors from solid metal, we need to know how thick to make the coating. The transmitted light typically has an initial amplitude $|E''| \approx 0.2|E|$, so we can’t just ignore it—we have to make the coating thick enough to be opaque.

Because of the hugeness of $\text{Im}\{n\}$, the wave equation in metals behaves like the diffusion (heat) equation, so that the electromagnetic fields are diffusive in character[‡]; the amplitude of a plane wave in a metal dies off by $\exp(-2\pi)$ per cycle. Also, $|n|$ is so large that \mathbf{k} is directed essentially normal to the interface, regardless of incidence angle.

In order for the light transmitted through the film not to significantly reduce the reflected power, we choose a film thickness $d > \lambda/\text{Im}\{n\}$, so that the light making a round trip through the film is attenuated by at least $e^{-4\pi}$, or 3×10^{-6} . A thousand angstroms of aluminum makes a good mirror in the visible, but 200 Å is getting a bit see-through. The optical constants of metals change so rapidly with wavelength that the thickness required can easily vary 2:1 across the visible, so light transmitted through a thin metal layer is often strongly colored. Metallic neutral density filters are made of chromium, rhodium, Inconel, or other metals whose properties are less wavelength dependent.

Aside: Free-Electron Metals. In the infrared, metals such as copper, silver, and gold exhibit *free-electron* behavior. That is, their dielectric constants behave as though the electrons were undamped harmonic oscillators. Since the harmonic oscillator equation is of second order, the response to a sinusoidal \mathbf{E} at a frequency above the resonance is multiplied by two copies of $j\omega$ —in other words, these metals have *negative real dielectric constants*. This leads to all sorts of useful effects such as *surface plasmons*, which are true electromagnetic surface waves that may turn out to have lots of useful technological applications.

[†]With our Fourier transform sign convention, absorbing media always have positive $\text{Im}\{n\}$. Why?

[‡]*Diffusive* is used in its mathematical rather than optical sense here.

5.2.3 Designing Metal Films

Silver is the best metal coating in the visible but tarnishes so fast that it is nearly useless for first-surface reflections. It works much better than aluminum ($R = 0.96$) for internal reflectors (e.g., pentaprisms), where the back of the silver can be protected with a plated layer of less active metal (e.g., copper, Inconel, or nickel) and a coat of paint. On the other hand, if silver is protected from sulfides in the air, it lasts for months, and it can be applied chemically in the lab, which is sometimes a very important property.

Gold is the best mirror coating in the NIR and the red ($\lambda > 633$ nm). It is often convenient for lab use, as many laboratories have small gold sputtering chambers intended for electron microscope samples, so the turnaround time is unbeatable. It doesn't stick all that well, though, so be careful with it. A very thin adhesion layer of chromium or titanium makes gold stick to glass very well. Rhodium is a noble metal whose reflectivity holds up well far into the UV (that's the good news—the bad news is that $R \approx 0.8$ in the visible and 0.4 in the UV). Good metal films (e.g., Cu, Ag, Au, Al) in the IR are essentially perfect conductors, so their efficiency is excellent.

5.3 TRANSMISSIVE OPTICAL COATINGS

Most optical surfaces are not bare glass or plastic, but are coated with one or more thin layers of another material to modify their transmission and reflection properties. The most common are antireflection (AR) and mirror coatings, but sometimes we want a beamsplitter, a polarizer, or a filter, all of which can be made by appropriate coatings.[†]

One can treat coatings in different ways; because we're after physical insight rather than, say, efficient simulations, we'll keep it simple and use the Fresnel formulas, assuming a plane-parallel geometry and homogeneous, isotropic films and substrates.

5.3.1 Dielectric Coating Materials

The theoretical performance of coatings is limited by the available materials and by the difficulty of getting good coating morphology and adhesion. In the early days of coated optics (the 1940s), the best available method for putting down dielectric coatings was vapor phase reactions in air. This produced surprisingly good coatings, so good that Carl Zeiss gave up on vacuum coaters completely for a while (see Anders).

At present, most coatings are deposited in a vacuum, by evaporation or sputtering. These techniques work well, but almost all coatings are a bit less dense than the bulk material, owing to the formation of small voids during deposition. These voids reduce the refractive index slightly and, by adsorbing air and water, cause drift in n with temperature and humidity. They also reduce the corrosion protection the coating affords. This is not especially serious with substrates that form thin oxides with good properties (e.g., Al), but is more of a problem with silver and copper, which do not.

Films are highly variable, depending on deposition conditions such as stoichiometry, humidity, choice of substrate, temperature, pressure, and other things, and there is significant variation from run to run. The microstructure of the film may be quite different from the bulk material's; a good quality coating is amorphous, which will influence its

[†]The discussion of optical coatings is indebted to the admirable small monograph by Hugo Anders of Carl Zeiss, Oberkochen, *Thin Films in Optics*, Focal Press, London, 1967 (J. N. Davidson, tr.).

refractive indices and transmission bands significantly (especially for anisotropic materials). Besides porosity and microstructure, coatings often have stoichiometry errors that can make n go up or down. For example, one maker quotes a range of 1.34–1.38 for its MgF_2 coatings at 550 nm. The idea of a table of optical constants of thin films is thus something of an oxymoron, so don't take the values in Table 5.1 too seriously. Above all, don't expect to get exactly the same refractive index and transparency range as the bulk material. The high variability of the indices of the films, and the moderate

TABLE 5.1. Common Coating Materials

Material	Index	λ (nm)	Comments
Cryolite (Na_3AlF_6)	1.35	Visible	Lowest n among dense coatings; water soluble; soft
Magnesium fluoride (MgF_2)	1.38	Visible	Lowest index hard coating; popular
Quartz (SiO_2)	1.46	Visible	
Silicon monoxide (SiO)	1.5–1.9	Visible	Nonstoichiometric; high index SiO absorbs in the blue
Silicon nitride	2.02	500	Absorbs strongly below 300 nm
Sapphire (Al_2O_3)	1.75	Visible	
Titanium dioxide (TiO_2)	2.5	Visible	Varies between 2.2 and 2.7
Zinc sulfide (ZnS)	2.35	Visible	
Lead fluoride	1.75	Visible	
Indium–tin oxide	2.0	500	Electrically conductive; good ITO transmits >87% in the visible
Silicon (Si)	3.5	1300	
Gold (Au)	1.66 + i 1.956 1.24 + i 1.80 0.61 + i 2.12 0.31 + i 2.88 0.16 + i 3.80 0.19 + i 5.39 0.27 + i 7.07 7.4 + i 53.4	400 477 517 564 689 827 1030 10,000	Needs Cr or Ni adhesion layer; best metal for $\lambda > 900$ nm
Silver (Ag)	1.32 + i 0.65 0.17 + i 1.95 0.13 + i 2.72 0.12 + i 3.45 0.15 + i 4.74 0.23 + i 6.9 0.65 + i 12.2 10.7 + i 69	310 400 476 564 729 1030 2000 10,000	Corrodes rapidly; best metal in the visible; much poorer below 400 nm; can be applied chemically in the lab
Aluminum (Al)	0.13 + i 2.39 0.49 + i 4.86 0.76 + i 5.5 1.83 + i 8.31 2.80 + i 8.45 2.06 + i 8.30 1.13 + i 11.2 25.4 + i 67.3	207 400 546 700 800 900 1130 10,000	Reasonably stable in dry air; best all-round metal; reflectance dips badly in the deep red and near IR (700–1000nm)

difficulty of obtaining highly uniform films of the correct thickness, must influence the way we design with thin films—a design needing three-figure accuracy in refractive index and thickness will be totally unmanufacturable. Theoretical elegance must be sacrificed to the exigencies of coating manufacturing. (This is not a blanket dismissal of fancy coatings—some highly multilayer coatings have been designed precisely to be very tolerant of certain classes of coating errors.)

Beyond the electromagnetic properties of coatings, their mechanical ones, such as adhesion, residual stress, and environmental sensitivity, must be considered. Not every coating material sticks to every other, or to every substrate (a couple of nanometers of Cr, Ti, or Ti_3N_4 can help a lot with metals). Materials with a coefficient of thermal expansion (CTE) very different from that of our glass will experience severe stress upon cooling, which may make the film craze or delaminate. Surface preparation is vitally important too—coating really is a bit of a black art. Detailed design of thin film coatings is beyond our scope, because most instrument builders buy parts already coated; nevertheless, you may easily require a custom coating occasionally and so need some understanding of the difficulties involved.

5.4 SIMPLE COATING THEORY

Grinding through the algebra for multiple-layer coatings becomes less fun very rapidly. What's more, the formulas so obtained provide no insight and are useless in practice due to their specificity. You can formulate the problem as a band-diagonal matrix equation based on the matching of tangential \mathbf{E} and perpendicular \mathbf{D} at each interface, with a phase delay of $\exp(\pm ik_z z)$ for waves going in the positive and negative z direction, but there's an easier way: for a plane-parallel geometry, the Fresnel formulas can be cobbled together to give us the full electromagnetic solution, assuming that all the materials are isotropic. We'll add up all the reflections and find the result, taking account of the propagation phase delays. As usual, we take the z direction to be the surface normal directed from layer j to layer $j + 1$.

A wave $\exp(i\mathbf{k} \cdot \mathbf{x} - \omega t)$ delayed by propagating through a layer of index n_j and thickness d_j acquires a phase delay of $k_{Zj}d_j$. The value of k_{Zj} depends on n_j , so we have to find it using the phase matching condition.[†] Phase matching states that \mathbf{k}_\perp is preserved across a boundary, as is necessary to preserve translational invariance. Thus in the j th layer, k_{Zj} obeys

$$k_Z^2 = n_j^2 k_0^2 - |\mathbf{k}_\perp|^2. \quad (5.2)$$

All the forward waves in the j th layer have this k_Z , and the reverse waves (i.e., those reflected an odd number of times) have $k'_{Zj} = -k_{Zj}$.

A point of terminology: film thicknesses are universally specified in waves and not nanometers: a half-wave film is one whose optical thickness is $\frac{1}{2}$ wave, that is, $d_j = \lambda/(2n_j)$. In coatings such as beamsplitters, intended to be used off-normal, a “half-wave” coating is usually specified as $d_j = \pi/k_{Zj}$, that is, thicker by a factor of $\sec \theta_j$, where θ_j is the angle of incidence in medium j .

[†]We're doing wave propagation here, so we'll stick with the physicists' sign convention, where a plane wave is $\exp[i(\mathbf{k}\mathbf{x} - \omega t)]$, and a time delay τ contributes a factor of $\exp(+i\omega\tau)$.

5.4.1 Multilayer Coating Theory

Multilayer coatings appear much more complicated, and in fact the explicit formulas for their reflection and transmission coefficients are ugly enough to give small children nightmares.[†] We therefore proceed recursively, calculating the effect on an existing coating stack of adding another layer, using the total reflection coefficient \hat{r} of all lower layers in place of the simple r from the Fresnel formula on the bottom side (see Section 1.2.4).

First we need to restate the Fresnel formulas ((1.8)–(1.10)) in terms of \mathbf{k} instead of θ_i ; we especially must keep the index 12 for propagating from n_1 into n_2 , since as we saw the Fresnel coefficients are not the same going in and coming out, and in dealing with multiple bounces we'll need to keep them straight:

$$\begin{aligned}
 r_{p12} &= \frac{n_2^2 k_{z1} - n_1^2 k_{z2}}{n_2^2 k_{z1} + n_1^2 k_{z2}}, & r_{s12} &= \frac{k_{z2} - k_{z1}}{k_{z2} + k_{z1}}, \\
 t_{s12} &= \frac{2k_{z1}}{k_{z1} + k_{z2}}, & t_{p12} &= \frac{2n_1 n_2 k_{z1}}{n_2 k_{z1} + n_1 k_{z2}}.
 \end{aligned}
 \tag{5.3}$$

Referring to Figure 5.4, we observe that the multiple bounces form a geometric series with common ratio $\hat{r}_{23} r_{21} \exp(+i2k_{z2}d_2)$, so that the total reflection coefficient \hat{r}_{12} can be expressed as the total reflection coefficient \hat{r}_{23} of the lower system with the additional effect of the layer n_2 , yielding

$$\hat{r}_{12} = r_{12} + \frac{t_{12} t_{21} \hat{r}_{23}}{\exp(-i2k_{z2}d_2) - \hat{r}_{23} r_{21}}.
 \tag{5.4}$$




Figure 5.4. Geometry of the plane-parallel coating problem, showing the geometric progression of multiple bounces from the stack below.

[†]Besides, the amount of CPU time spent calculating them is minuscule compared with the time spent getting the coating recipe debugged.

For the bottom layer, we take $\hat{r}_{12} = r_{12}$ from the Fresnel formula, and for all subsequent ones we use the total reflection of all lower layers, obtained from repeated application of (5.4), which gives us a recursive algorithm for calculating the total \hat{r} for the stack. We can compute t the same way, but now we need to go from the top of the coating stack toward the bottom, instead.

5.4.2 Lossless Coating Examples

Here are some typical examples of the uses of coatings. The goal is physical insight, not detailed coating recipes, so we neglect dispersion, material absorption, adhesion problems, and interface effects, for example, the 10–50 nm of Al_2O_3 that grows immediately on top of deposited aluminum, even in vacuum. Don't take this to mean that these effects are negligible in practice.

Example 5.1: Single-Layer AR Coating. For a single interface, the strength of the Fresnel reflection depends on n_2/n_1 and θ_i . At normal incidence, we can get reflections of the same strength and phase from both sides of a coating if $n_3/n_2 = n_2/n_1$, that is, when n_2 is $(n_1 n_3)^{1/2}$. By choosing the layer to be $\lambda/(4n_2)$ thick, the two reflected waves are put π out of phase (there and back) and thus cancel perfectly. The problem is that low index glass requires that n_2 be around 1.2 or 1.25, and there aren't any solid materials in that range (people have used silica aerogels with some success in the red and IR). The lowest index material that is hard and insoluble is MgF_2 , at 1.38. This is an excellent AR coating for high index glass, but it's no great shakes with garden-variety borosilicate such as BK7 (1.517), as you can see from Figure 5.5 (We'll take 1.37 as its index, a reasonable value closer to the center of the distribution than the bulk value of 1.38.) Note the angle tuning; at normal incidence the coating is centered around 514.5 nm, but at larger angles it shifts significantly to the blue. Higher index materials have more constant k_z (since $k_{\perp} \leq k_0$). The coating is also polarization sensitive, which means that




Figure 5.5. Single layer MgF_2 coating on BK7. Note the angle tuning and polarization dependence.

the polarization of your incoming beam will be changed somewhat by passing through the coated surface. A number of such surfaces together (e.g., in a camera lens) can easily cause serious polarization shifts with position and angle.

Example 5.2: Protected Aluminum Mirrors. First surface metal mirrors are too soft to clean vigorously and (except for gold) are also vulnerable to corrosion. The usual solution is to put a dielectric overcoat on top for protection. We saw that a glass–metal interface was not as good a mirror as air–metal, so protected metal mirrors start out with poorer performance even than Al–air. We can adjust the relative phases of the two reflected waves by changing the thickness of the coating; if we make the layer $\lambda/2$ thick, the reflections will add in phase. This effect is used to make hybrid mirrors, where the coating partly offsets the poorer reflectance of an internal reflection from aluminum. The most common coating choice is $\lambda/2$ of SiO over Al, the classical *protected aluminum* coating of Figure 5.6. Over a relatively narrow bandwidth, the result can be as good as a bare aluminum coating (the “internal reflection” curve is calculated for Al/SiO). These mirrors are OK for simple systems, or ones in which you have photons to burn; be careful how many bounces you use, though, or you may burn more than you can spare. The coating is cheap, which is important, but that and physical robustness about exhaust its virtues. For polarization-sensitive systems such as spectrometers, the fact that a significant proportion of the reflectance comes from the thin film coating means that protected aluminum mirrors are partially polarizing when used off normal. This is not so in the IR, where metals are more or less perfectly conducting; there, a dielectric coating of any thickness does not change R , which is always 1; protected gold coatings make great IR mirrors.

5.4.3 Angle Tuning

Because of the variation in k_z with incidence angle, the tuning of coatings shifts to shorter λ as θ_i increases, a phenomenon called *angle tuning*. It’s an easily calculated effect that can be reduced by using high index materials and reducing the field angle.




Figure 5.6. The protected aluminum mirror: 0.5 wave at 520 nm of SiO ($n = 1.7$) over Al.

Tuning with angle is generally a minor nuisance, as in Example 5.1, because k_z is a weak function of k_\perp at small angles. The difference between r_p and r_s and the increase in $|r|$ at high angles usually cause us worse problems. There are exceptions to this rule, such as polarizing beamsplitters, in which selectivity degrades with angle, and sharp interference filters, which can angle-tune your signal right out of the passband.

5.4.4 Examples of Multilayer Coatings

A good coater can put down layers of different materials without breaking vacuum, so it is convenient as well as worthwhile to put down stacks of many layers. Excellent antireflection coatings require several layers, and more complicated multilayer coatings can have excellent performance as mirrors, filters, and beamsplitters. A good AR coating can achieve 0.2% reflectance over a 50% bandwidth, with 0.4% being a typical guaranteed spec. Coatings for use across the whole visible usually achieve <1%.

Most optics manufacturers have a variety of standard coatings available, which may or may not be stock items. The price list will specify how much the coated elements cost, but be aware that lead times are frequently rather long (several weeks is not unusual). If you want different coatings on different surfaces, or a custom coating on parts you supply, be prepared to pay for it: even apart from design time, you'll be charged a setup fee of perhaps \$1200 for a coating run, plus a per-piece charge of around \$100 per surface, and will get no guarantee that your parts will survive.

Example 5.3: V-Coating. Single-layer AR coatings on plastic and low index glass don't work too well, because there are no materials available with n near 1.25. If we care about a narrow range of wavelengths (e.g., in a laser system), a quarter-wave of MgF_2 over a quarter-wave of SiO (with $n = 1.71$) can fix this problem, as Figure 5.7 shows. Here the coating's reflectance is calculated at normal incidence, and at 30° in both s and p . The coating angle-tunes to shorter λ with increasing θ_i , as we expect. Note how the p




Figure 5.7. Two-layer V-coating: quarter-wave MgF_2 over quarter-wave SiO (1.70) on BK7.




Figure 5.8. BBAR coating.

polarization is more tolerant of angular shifts; the decrease in the r_{pi} values partially compensates for the error in k_z .

Example 5.4: Simple Broadband AR Coating. For crown glass (1.46–1.65), a four-layer stack consisting of 0.50λ of SiO (1.60), 0.54λ MgF₂, 0.25λ SiO (1.60), and a final 0.27λ of MgF₂ on top makes quite a reasonable broadband AR (BBAR) coating; Figure 5.8 shows the calculated results on BK7 (1.517) at normal incidence, and 30° s and p .

Example 5.5: Enhanced Aluminum. The idea of metal with a dielectric overcoat can be improved by using an LH pair over a metal surface: Figure 5.9 shows the result of using a quarter-wave each of ZnS over MgF₂ on top of the aluminum (as before, the “internal reflection” curve is calculated for the index of the film adjacent to the aluminum). Another LH pair on top produces an even better mirror, with $R > 95\%$ even in the 800 nm dip. Enhanced aluminum mirrors are nearly as cheap as protected aluminum and are the minimum quality you should consider for high sensitivity instruments. (See Figure 5.9.)

Example 5.6: Quarter-Wave $(HL)^m H$ Stack. A sufficiently tall stack of $\lambda/4$ layers of alternating high and low index materials makes an excellent mirror. It works better with high index layers on both ends of the stack, so that the optical prescription is $(HL)^m H$. Figure 5.10 shows an 11-layer ($m = 5$) stack tuned to 514.5 nm. Note that it’s now the p polarization that droops at high angles, since the r_p values are dropping as θ_i increases.

Example 5.7: Stagger-Tuned HL Stack. Although the HL stack high reflector becomes broader as the number of layers increases, this happens only as \sqrt{N} , which is wasteful, and even this eventually stops due to absorption in the coating. By putting two HL stacks,




Figure 5.9. Replacing the half-wave of SiO with quarter-waves each of ZnS over MgF₂ yields a substantially improved metal mirror for the visible, the enhanced aluminum coating.




Figure 5.10. Eleven-layer (HL)⁵H stack, ZnS/MgF₂, centred at 514.5 nm.

tuned slightly differently, we can get a broader bandwidth with very high efficiency, as we see in Figure 5.11. You usually use a spacer layer between them. This idea is called *stagger tuning*, and it is broadly useful as a bandwidth-increasing device, in circuits as well as optics.




Figure 5.11. Two $(HL)^5H$ stacks (ZnS/MgF_2), tuned to 463 and 600 nm, with an 0.21 wave L spacer layer.

5.4.5 Polarizing Beamsplitters

Beamsplitters used to be made with thin metal films, and cheap ones still are. Inconel is the most popular choice due to its spectral neutrality, but all such coatings are very lossy—typically 40–50% gets absorbed, which is a lot, considering that even a lossless beamsplitter wipes out 75% of our light if we go through and back. Modern beamsplitters are made from dielectric stacks for narrowband or polarizing applications (see Figure 5.12), and from even thinner metal films with dielectric overcoats for wideband applications. You typically lose 15–25% of your light in an enhanced metal beamsplitter, and 5% or less in a good polarizing one.

Because r_s and r_p are so different at high incidence angles, the performance of coatings changes drastically with polarization at high angles. This is how broadband polarizing beamsplitter cubes are made: you come in at 45° to an $(HL)^mH$ stack, with enough layers that T_s is very low[†] (11 layers ($m = 5$) of ZnS and cryolite will get you to 0.0001, not counting absorption), for example, the one in Figure 5.11.

You choose n_{glass} so that light coming in at 45° to the hypotenuse (i.e., normal to the cube face) is refracted into the coating stack at Brewster's angle for the HL interface; this guarantees Brewster incidence at lower layers, because there are only two different indices involved. Cementing another 45° prism to the top of the coating stack makes a cube, where the transmitted light is undeviated and the reflected light comes off at 90° . This makes a good broadband beamsplitter, whose main problem is the first-surface reflections at the glass–coating and coating–cement interfaces. These pollute the reflected light with as much as 5% of the p polarization (it varies with λ because the two p reflections interfere). Adding extra AR coatings top and bottom can make a very good beamsplitter,[‡] marred only by its narrow angular acceptance (and, of course, the horrendous etalon

[†]Remember that you have to rejigger the coating thicknesses a bit to make $d_j k_{zj}$ equal to $\pi/4$ at each layer.

[‡]Not even close to a Wollaston prism for p -polarization quality, of course, but lower in absorption and cheaper—pretty good for a chunk of glass with a few films on it.

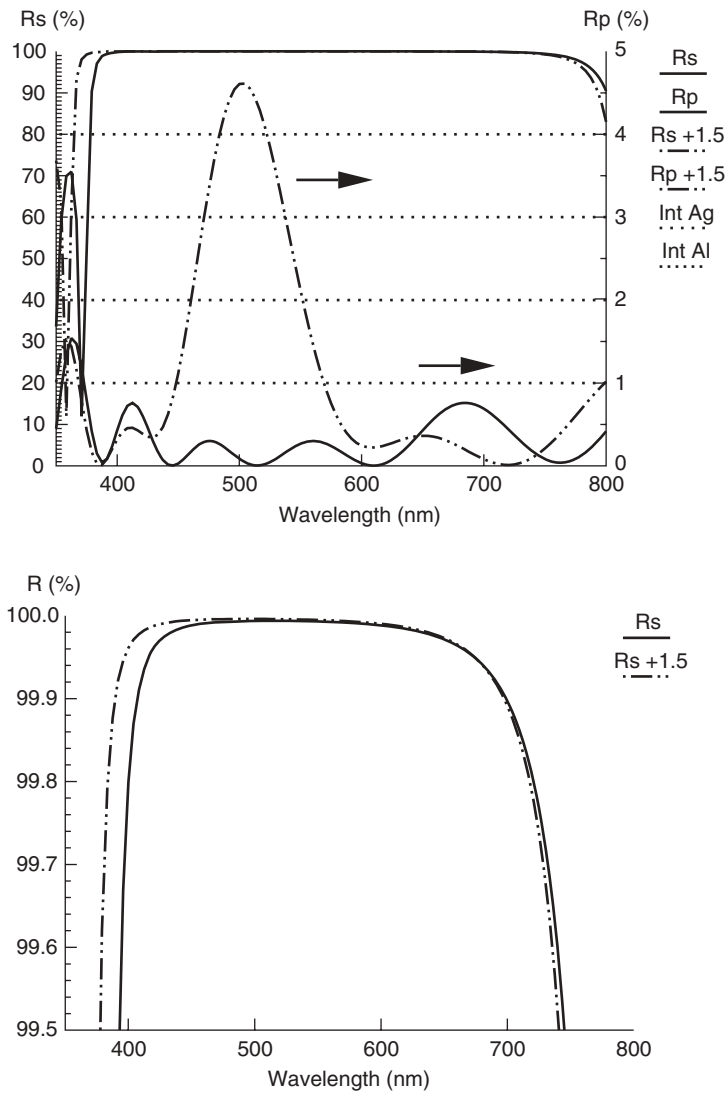


Figure 5.12. Polarizing beamsplitter: $A(HL)^5HA$, where A is SiN (2.0), and H and L are ZnS and cryolite. The AR layer A suppresses the reflection from the glass (1.66). Note the selectivity reduction (to 20:1 from 100:1) due to coming in at only 1.5° off normal.

fringes due to the manufacturer's insisting on choosing 45° , which we've alluded to in Section 4.7.2).

You can also make wider angle, narrower band beamsplitters by working off Brewster's angle, at a wavelength where r_p has fallen way off but r_s is still large. By careful control of the sidelobes of the $(HL)^m H$ stack's reflectance, you can make good beamsplitters for laser applications this way.

The polarization purity is still worse in the reflected light, only 25 or 50:1, whereas in the transmitted light it can be 1000:1 in a narrowband device or 100:1 in a wide-band one.

Aside: Unintentional Polarizing Beamsplitters. Some coatings are pretty good polarizers, including yours if you're not careful. Polarization effects in AR coatings cause major problems with lasers and in high accuracy applications.

5.4.6 Interference Filters

Two HL stacks separated by a spacer layer make a Fabry–Perot etalon, which has a sharply peaked passband near where the spacer is an integral number of half-wavelengths thick (phase shifts from the HL stack move the passbands around a bit). The bandwidth and free spectral range can be traded off by changing the spacer thickness, from $\lambda/2$ up.

A complex structure composed of two of these etalons deposited on top of each other, so that their passbands coincide at only one peak, makes an excellent filter with adjustable parameters. As usual, stagger tuning can flatten out the peak and suppress the sidelobes, yielding a flat-topped bandpass filter with steep edges and few artifacts.

Interference filters are fiddly to develop and need long coating runs that must be precisely controlled. Thus they tend to be available only for commonly desired wavelengths (e.g., laser lines) and spectral bands of wide interest (e.g., Balmer α at 656 nm). They are normally backed by colored glass, to suppress unwanted sidelobes, so that out-of-band light hitting the top of the filter is mainly reflected, while that hitting the bottom is largely absorbed. Sometimes it matters which way round you put the filter, for example, in the infrared, where the absorbing glass radiates but the mirror coating doesn't, and with high powered light sources, which may overheat the filter.

The stopband rejection of interference filters isn't always that great. Besides the occasional spurious peak, they sometimes have only 30 dB (optical) typical rejection. That might not be too bad in a color compensating filter for a camera, where the passband is wide and the rejection requirements modest. On the other hand, if you're looking at solar $H\alpha$ absorption at 656 nm with a 0.5 nm passband, you're in trouble—30 dB rejection means that each nanometer in the stopband will be attenuated by 30 dB, but there are a lot more nanometers in the stopband than the passband, so the background light will dominate. Make sure you calculate and measure the total out-of-band leakage in your filters. A quick test is to cant the filter enough to angle-tune your desired signal into the stopband, and see how much the signal level changes. This isn't really precise, because the L layers angle-tune more than the H , so the shape of the curve will change with angle too.

If you use interference filters in your instrument, be aware that they drift with temperature and time. They normally drift toward longer λ with increasing T , generally with $\Delta\lambda/\lambda \approx 10\text{--}30$ ppm/K; time variations are usually associated with hydration (or even corrosion) of the coatings (see Section 12.13.2). Get detailed specs from the manufacturer.

5.4.7 Coating Problems

Coatings usually have a columnar morphology, which makes them porous, chemically somewhat unstable, nonstoichiometric, and often having properties significantly different from the bulk material. Lots of work has gone into fixing these problems, but the solutions are different for different coatings. Sometimes depositing at an angle of 30° or so, or using ion bombardment during deposition, can reduce porosity and produce a fully dense coating. The columnar morphology can be exploited (e.g., by rotating the substrate eccentrically at an angle to the deposition source), so as to make the columns helical—that makes an optically active coating (see Section 6.3.6).

5.5 ABSORPTIVE FILTERS

Optical filters are used to get rid of light we don't want. (They always get rid of some of the desired light as well, but not too much with a bit of luck.) Here we'll talk about absorbing materials and scattering from small particles.

5.5.1 Filter Glass

Glass makers can exploit the absorption characteristics of different materials to make intentionally colored glass filters. Filter glass comes in a wide range of colors and characteristics, but the two most used are long pass and short pass, with the long pass glasses having generally better performance. The coloring can come from one of two sources: colloids, which are formed by a carefully controlled heat treatment process (*struck* or *colloidally colored* glass), or by the formation of color centers due to ionic doping (*ionically colored*). Ionically colored glass is a great deal more stable with time and thermal history. Color centers are not easily bleached by optical dose either, so ionically colored glass is pretty stable all round. In addition, it can be annealed to eliminate stress birefringence.

The data sheet for the filter glass usually tells how the color has been achieved. The transmission spectrum of the glass does shift somewhat with time and exposure to the air, so that your design should have a safety factor. It is usually a mistake to expose glass filters to severe weathering; in a hostile environment, make the objective (outermost element) out of something more robust (e.g., quartz or borosilicate crown glass).

Glass filters are often called upon to absorb large amounts of optical power, for example, in color-correcting a tungsten bulb with a blue filter to approximate daylight. Glass has poor thermal conductivity, so the temperature does not equilibrate rapidly; this leads to large temperature gradients and consequently large tensile stress in the cooler areas, to the point of shattering the filter. Filter glass therefore is usually tempered to increase the amount of heat that can be dumped into it before it breaks. This is useful but causes severe stress birefringence, which only gets worse with nonuniform heating (see the Schott filter glass catalog). For large heat loads, consider sending the heat away from the filter using a hot or cold mirror, or spreading the load by using a gentle filter to absorb half the heat, then a dense one to absorb the rest.

In long pass filters, the band edge shifts toward the red as the temperature increases, at a rate between 0.02 nm/K for deep UV filters to 0.3 nm/K for NIR filters; it tends to go as

$$\frac{\partial \lambda_c}{\partial T} \approx (5 \times 10^{-7} \text{nm}^{-1}) \lambda_c^2. \quad (5.5)$$

This shift is linear for reasonable temperature excursions, and large enough (hundreds of ppm/°C) to be very obnoxious. There is also a significant shift in passband absorption, which tends to be very large proportionately, since the wings of the exponential are very sensitive to slight changes in kt/e . These shifts are of course sensitive to field angle and NA and so are generally difficult to compensate for in software. If you're trying to do accurate photometry with filters, control their temperature carefully.

Filter glass is usually fluorescent, with a peak about 200 nm to the red of the absorption edge. The usual way of fixing this is to use a series of filters of different cutoff wavelength in series. Unfortunately, the order matters—putting them in the wrong order can cost you a factor of 1000 in leakage. This can be a big effect if you're looking for dim light

in a bright background (e.g., Raman spectroscopy), where it looks just like a light leak, so it'll have you chasing your tail—see Section 10.7.4.

Colored glass filters can have extremely high absorption in their stopbands and are inexpensive; these virtues make up for the gradualness of their absorption versus wavelength compared to interference filters. Unfortunately, due to low sales Schott has reduced the number of filter glasses in their catalog, so that finding a glass just right for your application is significantly more difficult than it once was.

5.5.2 Internal and External Transmittance

Some of the light incident on a dielectric surface is reflected, so that even if the material itself is completely lossless, not all the light hitting a dielectric plate makes it through. We distinguish the two sources of loss by speaking of *internal* and *external transmittance*. Internal transmittance excludes the Fresnel reflections at the surfaces, whereas the external transmittance includes them. For purposes of definition, the filter is assumed thick enough that interference effects and multiple reflections can be ignored.

One benefit of making this distinction is that the dependence of the internal transmittance on the thickness of the element is very simple; it follows *Beer's law*,

$$T_{\text{int}}(\lambda; d) = \exp[-\kappa(\lambda)d], \quad (5.6)$$

which allows us to predict T_{int} of an arbitrary thickness from a single data point:

$$T_{\text{int}}(\lambda; d_2) = [T_{\text{int}}(\lambda; d_1)]^{d_2/d_1}. \quad (5.7)$$

Due to this very strong thickness dependence, the shape of the curve of T_{int} versus d changes with thickness; the width of absorption features increases and the valley transmittance decreases as the element becomes thicker. Filter glass transmittance is usually plotted in terms of the *diabatie*

$$\Theta(\lambda) = 1 - \log_{10} \log_{10}[1/T_{\text{int}}(\lambda; d_0)], \quad (5.8)$$

where $T_{\text{int}}(\lambda; d_0)$ is the internal transmittance of a filter with standard thickness d_0 . A plot of diabatie does not change shape with thickness, but merely moves up and down; a common way of presenting the spectral characteristics of filter glass is to plot the diabatie in black on red coordinate axes, then give you a red transparent sheet with the grid and graduations on it. You line up the thickness scale so that the design thickness lines up with the fiducial on the plot, and presto, a plot of internal transmittance versus wavelength for your particular thickness (albeit on a weird vertical scale). Because numerical values of diabatie don't convey much, the scales are labeled with internal transmittance. Neutral density filters are usually thought of in terms of their optical density D ,

$$D(\lambda; d) = \log_{10}[T_{\text{ext}}(\lambda; d)]. \quad (5.9)$$

5.5.3 Holographic Filters

Another class of filters is based on holograms. Unlike interference filters, these tend to be available in bandstop types, perhaps 10–20 nm wide, with 40–80 dB (optical) rejection at the stopband center. These devices angle-tune as coatings do, but because of the depth of the null we're talking about here, it leads to a stiffer restriction: you have to use these filters with normally incident collimated beams. A strong beam at the stopband center produces weak surface scatter and stray junk that get through the filter, leading to a doughnut of speckle around the original beam direction. Since they are offset in angle, this is not a critical problem, but you have to put in a baffle after the filter.

5.5.4 Color Correcting Filters

In a tunable system, such as a monochromator driven by a tungsten source, it is often tempting to use color compensation filters, which are fairly gentle colored-glass filters intended to flatten the source spectrum by attenuating the red more than the blue. This should be avoided if possible, for a number of reasons. An optical filter cannot improve the flux of blue light, so that even in the blue, it will decrease the signal-to-noise ratio. The filter response will never be accurately inverse to the instrument function, if for no other reason than that the two change differently with time and temperature, so that a calibration will be necessary anyway. Sometimes there are good technical reasons for using such a filter, for example, a sample or detector that may be damaged by higher intensity in the red, a CCD that blooms badly when saturated, or a digitizer whose dynamic range is insufficient, but these are not as common as might be thought. A slightly more subtle problem is that these filters are mostly designed for color correction with photographic film and do not necessarily make the spectrum even approximately flat. The general idea of *whitening* a signal to improve the SNR is more useful in signal processing—see Sections 13.3.8 and 13.8.10.

5.6 BEAM DUMPS AND BAFFLES

Designing good beam dumps and baffles is a subtle business, which absolutely must be part of the early stages of your instrument design. A common error is to think about baffles last, when even trivial design changes are expensive, and then run for a pricey high performance coating such as Martin Black to fix it.

It must be clearly understood from the outset that the stray light performance of your system is controlled primarily by the geometry rather than by the quality of the black coatings themselves. You can lose a factor of 10^5 by allowing a large expanse of black painted lens barrel, illuminated at grazing incidence, to appear in your detector field of view, and you'll only gain back a factor of 10 or so by replacing the black paint with a fancy and expensive black dendritic finish. The rules are pretty simple.

1. Everything is specular at grazing incidence, so don't allow any grazing bounces to hit your detector.
2. The more illuminated area your detector sees, the more stray light it will receive, so keep the baffles and lens barrel out of the field of view as far as possible.

3. Multiple reflections from dark surfaces will rapidly eliminate stray light, so trap it and then absorb it. Don't allow any one-bounce paths to hit your detector (see rule 2).
4. Sharp edges of baffles will diffract light, so adjust the relative apertures of the baffles to catch it (i.e., later baffles should have slightly smaller inner diameters).

Instruments that search for extrasolar planets need about the best baffles going, which has led to the development of *band-limited* baffles. The idea here is just that of data windowing (see Section 17.4.9), in which a carefully chosen, gradual cutoff of the light leads to greatly reduced diffraction rings and consequently to improved sensitivity at small separations.[†]

5.6.1 What Is a Black Surface?

We call a surface black when it doesn't reflect light. The Fresnel formulas predict significant reflection from any discontinuity in \tilde{n} , in either its real or imaginary part. Accordingly, a black surface has an absorption depth of many wavelengths, but much less than its thickness, and is a good index match to the incoming wave. Black surfaces in air are doomed from the outset by the size of the index mismatch at the surface, but if the wave is coming in via a medium such as plastic or glass, the situation is much less dire.

5.6.2 Black Paint

Because of the aforementioned index mismatch, flat black paint in air has a diffuse reflectance of a few percent, which is blacker than TiO_2 but nothing to write home about. (Volume II, Chapter 3 of the *OSA Handbook* has a useful list of black finishes.)

Flat black is useful as a last ditch solution to a scattered light problem, where the stray light is not highly directional. The biggest problem with it is that a major fraction of the light will escape after only one bounce off the black surface, so that the ultimate performance of a flat black paint baffle that is in the detector's field of view is not that much better than that of a single flat black surface. The next biggest is that near grazing incidence, even flat black paint is a quite reasonable reflector. On the other hand, you have to coat the inside of your optical system with something, and flat black is at least less bad than the alternatives.

Internal reflection is a different matter; the improved index match and enforced smoothness of the glass-paint interface improve the qualities of paint enormously (flat black looks shiny from underneath). For example, garden-variety ultraflat black spray paint (Krylon #1602) is a spectacularly good index match to fused quartz, very useful for getting rid of internal reflections from unused areas of quartz prisms. Over the visible, the reflectance of such a quartz-Krylon interface is on the order of 0.01%, which is very impressive for hardware-store spray paint. Remember that paint has environmental limitations and tends to outgas and shed particles.

[†]See, for example, K. Balasubramanian, *Appl. Opt.* **47**(2), 116 (2008).

5.6.3 India Ink

India ink is an aqueous suspension of very fine carbon particles. It is pretty black when dry, but *really* black when liquid, especially if you can get rid of the first-surface reflection—in the visible, the absorption length in India ink is less than 1 cm even at a dilution of 1:10⁴.

5.6.4 Black Anodizing

Anodizing is a surface treatment for aluminum, which consists of electrochemically oxidizing a clean aluminum surface to produce a porous Al₂O₃ (alumina or sapphire) layer. The porosity is inherent—otherwise no current would flow after a very short while, as sapphire is an excellent insulator. The resulting porous matrix is then filled with something else, for example, aluminum hydroxide in the classical anodizing process, or fluoropolymer in some proprietary systems such as Tufram and Polyube. The color comes from dye that is part of the bath. Anodizing is less black than paint because of the high index of sapphire, and in the IR it may not be black at all, since organic dyes do not have the broad absorption spectrum of, say, carbon. Check before relying on its IR performance.

5.6.5 Dendritic Finishes

Coatings made of closely spaced, shiny black peaks or ridges are better absorbers than paint or anodizing. Lockheed Martin makes a dendritic black finish, Martin Black, based on this principle, which is one of a whole range of “designer blacks” (the *OSA Handbook* has a chapter on them). They’re useful but far from a complete solution and are very, very expensive. Dendritic finishes tend to reflect a bit more near grazing incidence. Recently, some dendrite-type blacks using oriented carbon nanotubes have got down to about 0.05% reflectance in the visible, but that still won’t save a system with lousy baffles.

5.6.6 Black Appliques

There are also a variety of stick-on black finishes, of which the flocked sticky paper sold by Edmund Optics deserves special mention. In the visible, it is comparable in performance to an advanced black coating such as Martin, but costs about 100 times less, and can be cut with scissors. It has a low damage threshold and probably outgasses somewhat due to the adhesive. Because its blackness comes from organic dye, it is much less impressive in the mid-IR, whereas Martin holds up quite well.

5.6.7 Black Plastic

Black plastic is optically similar to glossy black paint, although most of the time its surfaces are not as smooth on small scales. Like black anodizing, some types of black plastic are near-IR transmitting—in the 1970s, some mysterious offset drifts in plastic-packaged op amps were traced to photocurrents induced in the die by light getting through the phenolic plastic. (Modern packages are made of Novolac epoxy, which is very opaque.) If you’re in doubt, hold a big sheet of it up to an incandescent light and look through it with an IR viewer (don’t use the sun unless you’ve made sure the plastic is at least opaque enough to be safe for your eyes and your IR viewer).

5.6.8 Black Wax

Carbon black in grease or wax is very black indeed—Wood made his horn with lampblack (candle soot), which was extremely effective as well as convenient. There are a number of hard waxes that can be filled with carbon particles to make a very black material of appropriate refractive index for mounting prisms, especially calcite ones, where the mechanical weakness of wax avoids overstressing the soft crystals. (Apiezon W is a common choice that doesn't outgas.) It isn't a complete solution, though, because black wax doesn't actually stick that well. Prisms are mounted in metal cells that fit the prisms loosely, so the wax is thin and is never loaded in tension. Still, delamination due to mechanical or thermal stress is the major cause of death in calcite prisms.

5.6.9 Black Glass

Various types of very dark colored glass are available, including beer bottles[†] and Schott glasses. These can be UV-epoxied to the faces of prisms or windows in strategic positions, to collect and dissipate stray reflections before they go anywhere. This approach can take a lot more power than spray paint; a thick BK7 flat with a piece of black glass UV-epoxied to the back makes a nice laser reflection attenuator.

Black glass used straight is less satisfactory, as its surface finish is often poor, causing scatter. More subtly, if the laser beam power is high, the temperature gradients in the glass will cause it to become locally convex in the region of high intensity, which will defocus the reflected light. You don't need a kilowatt class laser to see this; 50 mW CW is easily enough. In transmission, this sort of effect is called thermal lensing.

5.6.10 Designing Beam Dumps and Light Traps

Assuming that you've designed the system sensibly, stray light will have to make at least one bounce to get to the detector. Thus controlling stray light involves two steps: reducing the amount of illuminated area in the field of view of the detector, and reducing the illumination intensity there. All that stray light has to go somewhere, and that somewhere is usually a baffle of some sort. We send unwanted beams into a *beam dump* and corral ambient light and scatter into a *light trap*. The two look more or less the same.

The best way to design beam dumps is to use *shiny* black surfaces, where the surface reflection is specular. A specular reflection can be directed onto another black surface, and another and another. . . . With care, you can make the light take many bounces before it can escape. The job is foremost to *trap* the light, and then to dispose of it.

5.6.11 Wood's Horn

The original beam dump is Wood's horn,[‡] a gently curved and tapered tube of glass coated outside with lampblack or black paint, shown in Figure 5.13a. It is a really good design, which works well over a reasonable range of incidence angles. The gentle taper traps the specular reflections, and due to the curved shape, most of the diffusely scattered light also has to make multiple bounces before escaping. Because the beam rattles around

[†]Beer's law is named after Dr. Beer.

[‡]Named after Robert Williams Wood, debunker of *N*-rays (*Nature* **70**, 530–531 (1904)), pioneer of grating spectroscopy, and author of *How to Tell the Birds from the Flowers*, among other stellar contributions.

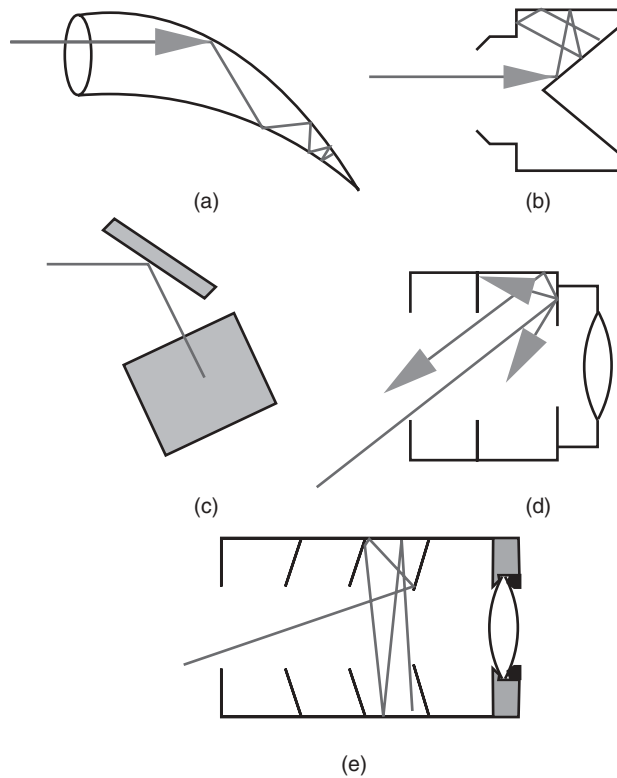


Figure 5.13. Assorted beam dump and baffle designs: (a) Wood's horn, (b) cone dump, (c) black glass at Brewster's angle, (d) knife-edge baffles, and (e) barbed baffles. Designs (a)–(c) use shiny black surfaces, and (d) and (e) shiny or flat black.

between surfaces making some angle with each other, it tends to be sent back out after some number of reflections, so the length of the horn has to be at least a few times its diameter.

5.6.12 Cone Dumps

A more convenient beam dump is the conical type, which fits optical breadboards and erector sets such as Microbench. As shown in Figure 5.13b, the cone absorbs most of the light and directs the rest into a simple trap arrangement that tends to confine the light. Light has to make at least three bounces from shiny black surfaces to escape, and most makes many more. These are easy to build in the lab if you have a lathe.

5.6.13 Black Glass at Brewster's Angle

You can combine black glass with Brewster angle incidence to get rid of an unwanted collimated beam (e.g., the residual pump laser beam in a photoacoustic measurement). As shown in Figure 5.13c, the first piece of black glass at Brewster's angle gets rid of one polarization, and the second one, at Brewster's angle for the other polarization (partly reflected from the first one), completes the job. This approach can take more peak power

than Wood's horn, but is restricted to well-collimated beams coming from one particular direction, and requires attention to the control of surface scatter from the black glass.

5.6.14 Shiny Baffles

A barbed pattern, with the barbs pointing toward the light slightly, is probably the best sort of baffle for long lens barrels and other narrow cylinders. Make the channels narrow enough that a beam needs several reflections to escape. Coat the inside with shiny black paint or make it from a shiny black material. The disadvantage of barbed patterns is that the light will eventually be reflected back out, and that the number of reflections required for this is a strong function of the angle of incidence.

5.6.15 Flat Black Baffles

In some situations, strong stray light can come from far off axis and may enter at any angle, as with sunlight in outdoor applications, so this angular dependence is inconvenient. In cases like this, we may resort to flat black surfaces and just reduce the illuminated area in the field of view. Optical design packages have stray light analysis that relies on the bidirectional reflectance distribution function (BRDF), which predicts the amount of light scattered into \mathbf{k}_2 from \mathbf{k}_1 . Do yourself a favor and design decent baffles.

An example of a flat black surface approach that works well is the knife-edge baffle, consisting of a series of black apertures lying normal to the axis of the optical system. Knife edges are easy to fabricate, being planar structures. The inside diameters decrease slightly coming toward the detector, so that ideally the earlier baffles in the series are out of the detector's field of view entirely, and thus light escaping from the flat black surfaces must either go back out the objective or hit another baffle. You do need to make the edges sharp, though, because if they're too blunt you'll have nice grazing incidence reflectors sending light into your detector. Knife edge baffles are lightweight and highly effective when properly designed.[†] Figure 5.13 shows some popular baffles and beam dumps.

5.6.16 Combinations

If the optical system is sufficiently long and narrow that the first few knife-edge baffles are really out of the detector FOV, you can use shiny black surfaces there and apply simple graphical ray tracing (starting from the detector) to estimate what the stray light intensity will be. A ray hitting baffle n will be near enough to normal incidence that it will rattle around between baffles n and $n - 1$ several times before exiting, which will get rid of it very effectively. If baffle n is smaller in aperture, and none of baffle $n - 1$ is visible to the detector, then a lot of the escaping light will exit the way it came in, which helps. Subsequent baffles, which may be in the FOV, can be flat black if necessary.

[†]See, for example, A. Buffington, B. V. Jackson, and C. M. Korendyke, Wide-angle stray-light reduction for a spaceborne optical hemispherical imager. *Appl. Opt.* **35**(34), 6669–6673 (1996).

5.7 WHITE SURFACES AND DIFFUSERS

A white surface is one that scatters incident light of all colors efficiently. We usually want diffuse white surfaces, whose scatter pattern is nearly Lambertian. White surfaces are made by suspending small particles of high index, nearly lossless dielectric in a low index medium, to scatter light as strongly and as often as possible.

5.7.1 Why Is It White?

When light enters a white surface, it is scattered in all directions; thus much of it does a random walk in the material. When it encounters the surface, most of it will escape (all of it, apart from Fresnel losses and TIR). Mathematicians call this phenomenon *gambler's ruin*—due to that boundary, eventually all your money diffuses out of your pocket. Considering how many times it will be scattered, and the geometric dependence of intensity on the number of scatterings, any absorption will be enormously enhanced; the same metals we use for mirrors appear black in powder form. Some white surfaces are better at depolarizing light than others, so measure yours carefully if it matters to you.

5.7.2 Packed Powder Coatings

The best diffuse, high reflection surface is a packed powder of pure TiO_2 , $\text{Ba}(\text{SO}_4)$, or MgO in air. Their reflectance is over 99% through most of the visible, and it is highly Lambertian (TiO_2 's drops off badly near 410 nm, and in commercial grades it also tends to fluoresce). What's more, unlike paint it is easily renewed if it gets dirty. It is primarily useful as a reflectance standard, because the surface is easily disturbed and really has to lie nearly horizontal. Barium sulfate's claim to fame is that its reflectivity is very constant with time and very flat with wavelength. Packed polytetrafluoroethylene (PTFE) powder will stick to surfaces well enough to use it in integrating spheres and has similar reflectance properties. Because of its lower refractive index, it needs a thicker section (at least 6 mm or so) to reach its peak reflectance, but can achieve values above 0.996 in the visible, and maintains its properties into the UV and NIR.[†]

5.7.3 Barium Sulfate Paint

The most popular diffuse white coating is barium sulfate paint (available from Edmund Optics as "Munsell white reflectance coating"). Paint is a collection of various particles in a dielectric binder. There is a front-surface reflection from the binder, which makes painted surfaces non-Lambertian, though they're closer than most other things. Compared with BaSO_4 powder in air, the dielectric is lossier, and the index mismatch at the surfaces is smaller, so the total reflectivity is also lower—about 98% in most of the visible and NIR. Regular white paint is loaded with TiO_2 and is closer to 90% in the visible. Barium sulfate paint is especially good in places like the interior of integrating spheres, where you need a nearly Lambertian reflector that's very opaque in a fairly thin layer (1–2 mm), and there's a lot of nonhorizontal area to cover.

[†]Victor R. Weidner and Jack J. Hsia, Reflection properties of pressed polytetrafluoroethylene powder. *J. Opt. Sci. Am.* **71**, 7 (July 1981).

5.7.4 Spectralon

Spectralon is a sintered PTFE sold by Labsphere, with properties similar to packed PTFE powder. It can be machined into odd shapes and is stable and cleanable. (Avian Technology sells similar stuff as Fluorilon-99W.) It is highly reflective and, although not so Lambertian as fine powders, it is very convenient for applications needing high efficiency diffuse reflectors that can absorb some punishment. The stuff is very expensive, though, so don't go milling an integrating sphere from a solid block of it.

5.7.5 Opal Glass

Opal glass is very inefficient at transmission (1%) but very Lambertian. It is used only for those applications for which diffuse illumination is vital. It produces very small speckles when used with lasers.

5.7.6 Magic Invisible Tape

Matte finish translucent tape is intended for mending torn papers, but it works pretty well as a diffusing material for light duty use, for example, putting a piece on the entrance slit of a spectrometer to fix spectral artifacts due to a weird pupil function, or to homogenize the ugly beam patterns of LEDs and liquid light guides. It won't take much power, and it leaves a slight residue behind, but it lasts for years, so it's just the right medicine sometimes.

5.7.7 Integrating Spheres

Light reflected from a white coating loses most of its directional information in a single bounce. A closed cavity with sufficiently high reflectance can bounce the light dozens of times before absorbing it, so that the illumination of a point on the wall becomes Lambertian to high accuracy; this is the idea of an integrating sphere. There are two main applications: measurement of optical power and building Lambertian light sources. The photon efficiency of an integrating sphere is a great deal higher than that of opal glass, so you can make a good light source by putting a small bulb inside the sphere, screened by a small white shield so that no unscattered light can reach the exit hole. The hole has to be fairly small, no more than 1/6 of the sphere diameter, to get the best performance. The same homogenizing property makes integrating spheres the best optical power sensors available; a photodiode in place of the bulb (still behind the shield) will see almost the same total flux regardless of the incident angle of the optical beam, assuming it isn't vignetted by the aperture, and furthermore the spatial and angular variations of the responsivity is homogenized out. Residual angular variation is at the 0.1% level unless the ports are too large.

The average number of bounces required for light to escape is equal to the total area of the apertures (projected on the sphere) divided by the area of the sphere. Assuming this is small, we can treat it like reflection loss, so by the geometric series formula,

$$\eta = \frac{P_{\text{out}}}{P_{\text{in}}} \approx \frac{A_{\text{out}}}{4\pi r^2(1 - R) + (A_{\text{out}} + A_{\text{in}})R}, \quad (5.10)$$

where R is the reflectance of the coating, r is the inside radius of the sphere, and A_{out} and A_{in} are the areas of the output and input ports (actually the areas of their projections on

the sphere). This doesn't take account of the few bounces it takes for the light to become Lambertian inside the sphere, the effect of baffles inside the sphere, or the deviation of the steady state illumination from a Lambertian condition due to the losses through the ports, all of which are generally small effects. For a perfectly reflecting sphere with equal sized ports, $\eta = 0.5$, and in real spheres, it is generally no more than 0.3 and is commonly much lower.

Similarly, a δ -function light impulse will be spread out into a roughly exponential pulse of time constant

$$\tau \approx \frac{4r}{3c} \frac{1}{(1 - R) + (A_{\text{out}} + A_{\text{in}})R/(4\pi r^2)}, \quad (5.11)$$

which is on the order of 10–50 ns for most spheres.

From (5.10), we see that the efficiency of a sphere is a very strong function of the coating reflectance, particularly if the coating is very good. Spheres made with the very highest reflecting coatings are therefore somewhat less stable with time and conditions than those made with ordinary white paint; on the other hand, they do a better job of diffusing the light and waste fewer photons. Keep your integrating spheres clean, plug any unused apertures with the white caps provided, and use them in their intended wavelength interval. This sensitivity can also be used to advantage in multipass measurements: see Section 10.6.5 for an example. The many bounces taken by a typical photon before it is absorbed or lost unfold to quite a long optical path, as we saw, and this can be very helpful in reconciling the demands of fast pulses to the capabilities of photodiodes, as in Section 3.5.4.

5.7.8 Ping-Pong Balls

You can make a rough-and-ready integrating sphere from a ping-pong ball. Paint the outside white, drill two holes at 120° from each other, and put a photodiode in one of the holes. This is good enough to show some of the advantages of spheres, but not for real measurements. (Ping-pong balls are also good scatterometers—see Section 9.8.7.)

5.7.9 Ground Glass

Ground glass is much more efficient at light transmission than opal glass but has a big specular component as well as a diffuse component. Use it where the specular component is not a critical problem, because besides being cheaper than opal glass, it is dramatically more efficient (30–70% vs. 1%). Because of TIR, it matters which way round you put the ground side; light crossing the ground surface from inside the glass is somewhat more diffuse but dimmer. Objects whose light scattering occurs at a surface tend to produce a constant spread of u and v , centered on the unscattered \mathbf{k} vector. Thus the angular spectrum is not independent of the incidence angle, which complicates diffuser design. If you have something like a video projector with a short focal length lens, shining on ground glass, it will scatter light only a bit around the original \mathbf{k} vector, so at the edges most of the light will still be going up and away rather than straight out, as you would probably want. Software can correct for this at a single viewing angle, but not for everyone in the room. This sounds like a job for a field lens (see Section 12.3.14)—the big Fresnel lens of Figure 5.14. straightens out the light before it hits the ground glass, which makes the brightness look much more uniform with viewing angle (though no closer to Lambertian than before).




Figure 5.14. Ground glass and other mild diffusers tend to scatter light into a cone about the incident ray direction as in (a). Adding a Fresnel field lens as in (b) can improve the apparent uniformity.

5.7.10 Holographic Diffusers

A better controlled version of ground glass can be made with a holographic element, the *holographic diffuser*. These are touted as being useful for laser beam shaping, but in reality the strong speckle they produce limits their application to low coherence sources such as LEDs. One very good application is to homogenize the output of fiber bundle illuminators. Nowadays you can get holographic diffusers that are nearly as Lambertian as opal glass, or have other angular patterns such as top-hat or square, without the high losses of opal glass diffusers. Not all holographic diffusers have the other special properties of opal glass (e.g., independence of illumination angle).

5.7.11 Diffusers and Speckle

A light source such as a HeNe laser, which is highly coherent in both space and time, is very difficult to use with diffusers. Shining a HeNe into an integrating sphere produces an optical field that is very Lambertian on a broad-area average, but that has very strong small-scale structure called *speckle*. All rough surfaces illuminated with lasers produce speckles that are a complicated function of position, but whose size is characteristic of the material and of the size of the illuminated region. More diffuse materials produce smaller speckles; the angular extent of the smallest ones is on the order of λ/d , where d is the incoming beam diameter. Speckle consists of a mass of unorganized interference fringes, caused by the coherent summation of fields from everywhere in the sphere. At each point, these random reflections produce a certain optical amplitude and phase in each polarization component, which vary all over the place. The best diffusers, such as integrating spheres, produce speckles with characteristic size $\lambda/2$. Due to speckle statistics, the relative standard deviation of the photocurrent will be on the order of $\sqrt{(\text{PD area})(\text{speckle area})}$, which isn't that small, and any vibration will smear that out into a huge noise PSD in the low baseband. Thus diffusers aren't always the way to get good measurements, at least with lasers. See Section 2.5.1 for more discussion.

Polarization

He flung himself on his horse and rode madly off in all directions.

—Stephen Leacock, *Gertrude the Governess*

6.1 INTRODUCTION

Optical polarization is the main way the vector wave nature of light manifests itself in practical problems. We've encountered plane waves, which are always perfectly polarized, and the Fresnel formulas, which predict the intensity and polarization of plane waves leaving a dielectric surface. Here we go into the measurement and manipulation of polarization, and how not to get in trouble with it. Polarization components such as retarders and Faraday rotators are mysterious to lots of people, but are actually fairly simple devices unless you try to get deeply into the physics of how they do what they do. Being practical folk, we'll stick with their phenomenology and keep their inner workings pretty well out of it.

The major uses of polarization components in optical systems are to control reflections, as in sunglasses and fiber isolators, and to split and combine beams without the heavy losses caused by ordinary beamsplitters.

6.2 POLARIZATION OF LIGHT

6.2.1 Unpolarized Light

If you send thermal light through an analyzer,[†] twist the control ring as you may, the same proportion of the light comes through. This remains true if you put any sort of lossless polarization device ahead of it; a wave plate or a Faraday rotator doesn't change the polarization at all. Thus we say that thermal light is *unpolarized*. This is a poser, because we know that any optical field can be decomposed into plane electromagnetic

[†]Analyzers and polarizers are physically identical, but an analyzer is thought of as detecting the polarization state produced by the polarizer—in communications terms, the analyzer is part of the receiving section, and the polarizer is part of the transmitting section.

waves. Since all such waves are perfectly polarized, how can thermal light be unpolarized?

The key is that we're really measuring the *time-averaged* polarization rather than the instantaneous polarization. The light at any point at any instant does in fact have a well-defined \mathbf{E} vector, because if it didn't, its energy density would be 0. In an unpolarized field, though, the direction of \mathbf{E} varies extremely rapidly with time, changing completely in a few femtoseconds in the case of sunlight. Thinking in terms of modulation frequency (see Section 13.3), the polarization information is not concentrated at baseband the way it is with lasers, but instead is smeared out over hundreds of terahertz of bandwidth. It is spread so thin that even its low frequency fluctuations are hard to measure.

In \mathbf{k} -space terms, the polarizations of different plane wave components are completely uncorrelated, for arbitrarily close spacings in \mathbf{K} . This is in accord with the entropy-maximizing tendency of thermal equilibrium—any correlation you could in principle use to make energy flow from cold to hot is always 0 in thermal equilibrium.

6.2.2 Highly Polarized Light

If we pass thermal light through a good quality polarizer, we get highly polarized thermal light. The plane wave components are still uncorrelated in phase but are now all in the same polarization state. If such light does not encounter any dispersive birefringent elements, its polarization state may be modified but it will remain highly polarized. Its polarization can be changed achromatically with TIR elements such as Fresnel rhombs, so that we can have thermal light with a well-defined circular or elliptical polarization.

6.2.3 Circular Polarization

We've encountered circular polarization before, but there's one property that needs emphasizing here, since so many useful polarization effects depend on it: the helicity changes sign on reflection. Left-circular polarization becomes right circular on reflection, and vice versa— \mathbf{E} keeps going round the same way, but the propagation direction has reversed, so the helicity has reversed too. This is also true of ordinary screw threads viewed in a mirror, so it's nothing too mysterious. Although linear polarization can be modified on oblique reflection from a mirror (if \mathbf{E} has a component along the surface normal), circular polarization just switches helicity, over a very wide range of incidence angles.[†] Since linear polarization can be expressed in terms of circular, this should strike you as odd—there's a subtlety here, called topological phase, that makes it all come out right in the end.

6.2.4 An Often-Ignored Effect: Pancharatnam's Topological Phase

When light traverses a nonplanar path, for example, in two-axis scanning, articulated periscopes, or just piping beams around your optical system, its polarization will shift.

[†]If the reflection occurs at a dielectric interface (where $r_p \neq r_s$), the polarization will become elliptical, at θ_B the ellipse degenerates into linear polarization, and beyond θ_B , the helicity no longer reverses. (Why?)

For reflection off mirrors, this isn't too hard to see: since \mathbf{E} is perpendicular to \mathbf{k} , a mirror whose surface normal has a component along \mathbf{E} will change \mathbf{E} . Make sure that you follow your polarization along through your optical system, or you may wind up with a nasty surprise.

A much more subtle fact is that the same is true for any system where light travels in a nonplanar path (e.g., a fiber helix). Left- and right-circular polarizations have different phase shifts through such a path, giving rise to exactly the same polarization shift we get from following the mirrors; this effect is known as Pancharatnam's topological phase[†] and is what accounts for the puzzling difference in the polarization behavior of linear and circularly polarized light upon reflection that we alluded to earlier (the corresponding effect in quantum mechanics is Berry's phase, discovered nearly 30 years after Pancharatnam's almost-unnoticed work in electromagnetics). This sounds like some weird quantum field effect, but you can measure it by using left- and right-hand circular polarized light going opposite ways in a fiber interferometer.[‡] These polarization shifts are especially important in moving-mirror scanning systems, where the resulting large polarization shift may be obnoxious.

It sounds very mysterious and everything, but really it's just a consequence of spherical trigonometry; the \mathbf{k} vector is normal to a sphere, and \mathbf{E} is tangent to the sphere throughout the motion; depending on how you rotate \mathbf{k} around on the surface, \mathbf{E} may wind up pointing anywhere. Equivalently, 2×2 rotation matrices commute, but 3×3 ones don't.

If you follow your \mathbf{k} vector around a closed loop enclosing a solid angle Ω , the relative phase of the right- and left-circular polarizations gets shifted by

$$\Delta\phi = \pm 2\Omega. \quad (6.1)$$

6.2.5 Orthogonal Polarizations

We often describe two polarization states as *orthogonal*. For linear polarizations, it just means perpendicular, but what about circular or elliptical ones? The idea of orthogonal polarizations is that their interference term is 0, that is,

$$\mathbf{E}_1 \cdot \mathbf{E}_2^* = 0. \quad (6.2)$$

Two elliptical polarizations are thus orthogonal when their helicities are opposite, their eccentricities equal, and their major axes perpendicular (i.e., opposite sense of rotation, same shape, axes crossed). It's an important point, because as we'll see when we get to the Jones calculus in Section 6.10.2, lossless polarization devices do not mix together orthogonal states—the states will change along the way but will remain orthogonal throughout. One example is a quarter-wave plate, which turns orthogonal circular polarizations into orthogonal linear polarizations, but it remains true even for much less well-behaved systems such as single-mode optical fibers.

[†]S. Pancharatnam, Generalized theory of interference and its applications. Part 1. Coherent pencils. *Proc. Indian Acad. Sci* **44**, 2247–2262 (1956).

[‡]Erna M. Frins and Wolfgang Dultz, Direct observation of Berry's topological phase by using an optical fiber ring interferometer. *Opt. Commun.* **136**, 354–356 (1997).

6.3 INTERACTION OF POLARIZATION WITH MATERIALS

6.3.1 Polarizers

A polarizer allows light of one polarization to pass through it more or less unattenuated, while absorbing or separating out the orthogonal polarization. Any effect that tends to separate light of different polarization can be used: anisotropic conductivity, Fresnel reflection, double refraction, walkoff, and the different critical angles for *o*- and *e*-rays (related to double refraction, of course).

Polarizers are never perfectly selective, nor are they lossless; their two basic figures of merit at a given wavelength are the loss in the allowed polarization and the *open/shut ratio* of two identical polarizers (aligned versus crossed) measured with an unpolarized source, which gives the polarization purity. The best ones achieve losses of 5% or less and open/shut ratios of 10^5 or even more.

6.3.2 Birefringence

The dielectric constant $\underline{\epsilon}$ connects the electric field \mathbf{E} with the electric displacement \mathbf{D} ,

$$\mathbf{D} = \underline{\epsilon}\mathbf{E}. \quad (6.3)$$

For a linear material, $\underline{\epsilon}$ is a tensor quantity (in isotropic materials the tensor is trivial, just ϵ times the identity matrix).[†] (See also Section 4.6.1.) Tensors can be reduced to diagonal form by choosing the right coordinate axes; the axes that diagonalize $\underline{\epsilon}$ are called the *principal axes* of the material; symmetry requires that they be orthogonal in this case. (The refractive index also of course may depend on polarization but is not a tensor, because it does not express a linear relationship.)

Some common birefringent optical materials are crystalline quartz, sapphire, calcite (CaCO_3), and stretched plastic films such as polyvinyl alcohol (PVA) or polyvinylidene chloride (Saran Wrap). All these, along with most other common birefringent materials, are *uniaxial*[‡]; two of their three indices are the same, $\epsilon_x = \epsilon_y = \epsilon_{\perp}$; light polarized in the plane they define is an ordinary ray (*o*-ray), so called because it doesn't do anything strange. The third index, which defines the *optic axis*, may be larger (positive uniaxial) or smaller (negative uniaxial) than the *o*-ray index; if \mathbf{E} has a component along the optic axis direction, strange things occur, so that the beam is called an *e*-ray, for "extraordinary." Things get stranger and less relevant for absorbing birefringent materials and for biaxial ones, so we'll stick with the important case: lossless uniaxial materials.

Electromagnetic fields are *transverse*, which in a uniform medium means that for a plane wave, \mathbf{E} , \mathbf{H} , and \mathbf{k} are always mutually perpendicular, and that the Poynting vector \mathbf{S} always lies along \mathbf{k} . (The Poynting vector generally defines the direction the energy

[†]Landau and Lifshitz, *The Electrodynamics of Continuous Media*, has a lucid treatment of wave propagation in anisotropic media, which the following discussion draws from.

[‡]Less symmetric materials may be biaxial, that is, have three different indices, and in really messy crystal structures, these axes need not be constant with wavelength. Biaxial crystals exhibit some weird effects, such as *conical refraction* (see Born and Wolf).

goes in; that is, it's the propagation axis of the beam as measured with a white card and ruler).[†]

Neither of these things is true in birefringent materials, where we have only the weaker conditions that \mathbf{D} , \mathbf{B} , and \mathbf{k} are mutually perpendicular, as are \mathbf{E} , \mathbf{H} , and \mathbf{S} . For instance, the actual index seen by the e -ray changes with angle, unless the light propagates in the plane defined by the ordinary axes, for only then can \mathbf{E} lie exactly along the optic axis. The propagation vector \mathbf{k} defines an ellipsoid (where x , y , and z are the principal axes),

$$\frac{k_x^2}{\epsilon_x} + \frac{k_y^2}{\epsilon_y} + \frac{k_z^2}{\epsilon_z} = k_0^2. \quad (6.4)$$

The refractive index $n = k/k_0$ experienced by a given e -ray varies with its propagation direction. The extreme values of n_e are n_\perp (the o -ray index n_o) when \mathbf{k} is along the optic axis and n_\parallel when \mathbf{k} is normal to the optic axis. There is a lot of sloppiness in the literature, with n_\parallel often being referred to as n_e , whereas n_e really varies between n_\perp and n_\parallel . Light encountering the surface of a birefringent material is split apart into two linearly polarized components going in different directions. Phase matching dictates that \mathbf{k}_\perp is preserved across the boundary. The o -ray behaves just as if the material were isotropic with an index of n_\perp , so that's easy— \mathbf{S} is parallel to \mathbf{k} .

Determining \mathbf{k} and \mathbf{S} for the extraordinary ray is straightforward. The direction and magnitude of \mathbf{k}_e can be found from (6.4) and the phase matching condition. Once \mathbf{k}_e is known, the direction of \mathbf{S} can be found easily; it lies in the plane defined by \mathbf{k} and the optic axis, and the angles θ_k and θ_S separating the optic axis from \mathbf{k}_e and \mathbf{S} obey

$$\tan \theta_S = \frac{\epsilon_\perp}{\epsilon_\parallel} \tan \theta_k. \quad (6.5)$$

Remember, though, that the phase term is still $\exp(i\mathbf{k} \cdot \mathbf{x})$ —stick with this and don't get confused by trying to calculate propagation distance along \mathbf{S} and multiplying by $n_e k_0$ or something like that.

If light travels directly down the optic axis, \mathbf{E} has no component along it, so the material appears to be isotropic. This is useful where the birefringence is obnoxious, for example, “ c -axis normal” sapphire windows used for their strength and chemical inertness.

6.3.3 Retardation

Since the phase velocity of light in a material is c/n , the e - and o -rays propagate at different phase velocities, so the two linear polarization components with \mathbf{k} along z will be phase shifted with respect to each other by an amount δ , where

$$\delta = (n_e - n_o)k_0 z. \quad (6.6)$$

[†]Care is needed in identifying $\mathbf{E} \times \mathbf{H}$ with the local energy flux: Poynting's theorem applies to the integral of $\mathbf{S} \cdot d\mathbf{A}$ over a closed surface, or equivalently with the volume integral of $\nabla \cdot \mathbf{S}$. That means that \mathbf{S} is nonunique in much the same way as the magnetic vector potential—Poynting's theorem still holds if we add to \mathbf{S} the curl of any arbitrary vector field. It usually works.

Unless the incoming beam is a pure *e*- or *o*-ray, this will change the resulting polarization (as we saw in Section 1.2.8). This phenomenon is called *retardation* and is the basis for wave plates. Retardation is usually specified in nanometers, since it is a time delay Δt that causes a phase shift $\delta = \omega \Delta t$, in contrast to a reflection phase as in a Fresnel rhomb, which is almost wavelength independent. (In other words, retarders are wavelength dependent even when the material has no dispersion.)

6.3.4 Double Refraction

An oblique beam entering such a material from an isotropic medium splits into two beams, because the different refractive indices give different angles of refraction by Snell's law. This phenomenon is called *double refraction* (which is what birefringence means, of course).

6.3.5 Walkoff

Besides double refraction, birefringent materials exhibit *walkoff*, as shown in Figure 6.1. Although the \mathbf{k} vector behaves normally in a birefringent material, the Poynting vector does not; the energy propagation of the *e*-ray is not parallel to \mathbf{k} , but lies in the plane defined by \mathbf{k} and the optic axis, somewhere between them, so that the *e*-ray seems to walk off sideways. This weird effect arises because the Poynting vector is parallel to $\mathbf{E} \times \mathbf{H}$. The tensor character of ϵ prevents \mathbf{E} from being perpendicular to \mathbf{k} ,[†] and the




Figure 6.1. Polarizers based on beam walkoff: (a) simple walkoff plate or beam displacer and (b) the Savart plate, a symmetrical walkoff plate.

[†]Unless \mathbf{D} is an eigenvector of ϵ , that is, is a pure *o*-ray or lies along the optic axis.

cross-product relation then forces \mathbf{S} to be not along \mathbf{k} . This effect has nothing to do with double refraction; instead, it's a spatial analogue of the phase velocity/group velocity distinction for a pulse of light. A general beam normally incident on a planar slab of birefringent material will split apart into two beams going in different directions. Double refraction can't cause this directly, since at normal incidence no refraction occurs. Oblique beams walkoff, as well, but is less obvious then. Now you know why it's called the extraordinary ray.

Aside: Defuzzing Filters. Very thin walkoff plates, usually LiNbO_3 , are often used in CCD cameras to reduce the disturbing moiré patterns due to the way the pixels are arranged in color cameras (see Section 3.9.14). Two walkoff plates mounted at 90° to one another, with a $\lambda/4$ plate in between, split an image point into an array of four points, thus multiplying the OTF of the camera by $\cos(u dx) \cos(v dy)$, where dx and dy are the shift distances. This apodization rolls off the OTF to zero at frequencies where the moiré patterns are objectionable. (Sometimes quartz is used, but it has to be thicker, which causes more aberration.)

6.3.6 Optical Activity

A birefringent material has different refractive indices for different linear polarizations. A material that has some inherent helicity, such as a quartz crystal or a sugar solution, has different indices for different circular polarizations (helical antennas respond more strongly to one circular polarization than to the other). The different coupling between the bound electrons and the radiation field gives rise to a slightly different index of refraction for the two helicities. As a linearly polarized wave propagates in such a medium, \mathbf{E} changes direction; if there is no birefringence, \mathbf{E} describes a very slow helix as it propagates. This is called *optical activity* or *circular birefringence*.

Noncentrosymmetric crystals such as α quartz and tellurium dioxide may exhibit both optical activity and birefringence; this combination makes the polarization effects of a random hunk of crystal hard to predict.

If you put a mirror on one side of a piece of isotropic but optically active material (e.g., a cuvette of sugar water), the linear polarization that comes out is exactly the same as the one that went in; the rotation undoes itself. This is because the helicity reverses itself on reflection—each component crosses the material once as left circular and once as right circular, so that their total delays are identical, and the original polarization direction is restored.

The effects of optical activity are fairly weak but highly dispersive; for a 90° rotation in α quartz, you need 2.7 mm at 486 nm and 12 mm at 760 nm; this is around 100 times weaker than the effect of birefringence, so it dominates only when the light is propagating right down the optic axis. The dispersion is occasionally useful (e.g., in separating laser lines with low loss), but since it's an order of magnitude higher than a zero-order wave plate's, optical activity isn't much use with wideband light except for making pretty colors.

Due to the columnar morphology of coatings (see Section 5.4.7), it is possible to make artificial circular birefringent coatings by evaporation. The substrate is held at an angle to the source and slowly rotated about the source–substrate axis, producing helical columns that are highly optically active.

6.3.7 Faraday Effect

Another effect that leads to the slow rotation of \mathbf{E} is the Faraday or magneto-optic effect, which is often confused with optical activity because the effects are superficially very similar. Terbium-doped glasses and crystals such as terbium gallium garnet (TGG), immersed in a magnetic field, rotate the polarization of light propagating parallel to \mathbf{B} by

$$\Theta = VBl, \quad (6.7)$$

where Θ is the rotation angle, l is the path length, B is the axial magnetic field, and V is a material property called the Verdet constant. The difference here is that there is a special direction, defined by \mathbf{B} . Heuristically, if you imagine that the application of \mathbf{B} starts a bunch of bound currents going around in circles, then what matters is not the helicity but whether \mathbf{E} is rotating the same way as the currents or not, because the dielectric susceptibility will be different in the two cases.

The key point is that the rotation direction does not change on reflection. If we put our mirror at one side of the magneto-optic material, \mathbf{E} keeps going round the same way on both passes, so the helicity change on reflection does not make the delays equal; Faraday rotation doubles in two passes, instead of canceling—it is said to be *nonreciprocal*. This property allows us to build *optical isolators*, which allow light to propagate one way but not the other, and Faraday rotator mirrors, which help undo the polarization nastiness of optical fibers.

6.4 ABSORPTION POLARIZERS

Some materials exhibit polarization-selective absorption, as in Polaroid sunglasses. They do it by anisotropic conductivity, which is what you'd expect given the close relationship between conductivity and the imaginary part of n .

6.4.1 Film Polarizers

Film polarizers are made of anisotropically conductive polymer: stretched polyvinyl alcohol (PVA) doped with iodine. They work throughout the visible, but deteriorate in the infrared, and are useless in the ultraviolet since PVA is strongly absorbing there. There are several different kinds, for different wavelength intervals, but the good ones absorb about 20–40% of the allowed polarization and have open/shut ratios of 10^4 . The selectivity of older types used to degrade significantly in the blue, but the newer ones are much better. Their wavefront quality is relatively poor (about like window glass, $2\lambda/\text{inch}$), and they have a very low damage threshold, only 1 W/cm^2 or so.

In order not to be limited by the wavefront wiggles, put the polarizer near the image. An image formed in thermal light has very small phase correlations between points to begin with (since the phase of object points further than λ/NA apart is essentially uncorrelated), so phase wiggles are not much of a worry at an image.

6.4.2 Wire Grid Polarizers

Wire grids, which despite their name are arrays of very thin, closely spaced, parallel wires, function well in the mid- and far-infrared, but the difficulty of making the pitch

fine enough prevents their use in the visible—such a structure is obviously a diffraction grating, so the pitch has to be fine enough that the first diffracted order is evanescent (see Section 7.2). Their open/shut ratios are usually about 10^2 , and they absorb or reflect about 50% of the allowed polarization. They also reflect some of the rejected polarization, but how much depends on the metal and the geometry. Shorter-wavelength grids are usually lithographically deposited, so you have to worry about substrate absorption as well.

6.4.3 Polarizing Glass

A development of the wire grid idea is the dichroic[†] glass Polarcor, made by Corning. It is an optical glass with small metallic silver inclusions. During manufacture, the glass is stretched along one axis, which transforms the inclusions into small whiskers, aligned with the axis of the stretch. These whiskers function like little dipole antennas and are highly absorbing in a relatively narrow band. At present, Polarcor is best used in the near-infrared (out to $1.6\ \mu\text{m}$) but is available down to 600 nm in the visible (transmittance deteriorates somewhat toward short wavelengths). It has excellent transmission in one polarization (70–99%), and excellent extinction in the other ($\approx 10^{-5}$), so that its open/shut ratio is comparable to that of crystal polarizers. It has a wide ($\pm 30^\circ$) acceptance angle and good optical quality—though there are a few more striae than in ordinary optical glass, as one would expect. Polarcor's laser damage threshold is lower than calcite's—if you hit it too hard, the silver grains melt and lose their shape. The threshold is around $25\ \text{W}/\text{cm}^2$ CW, or $0.1\ \text{J}/\text{cm}^2$ pulsed.

6.5 BREWSTER POLARIZERS

At intermediate angles of incidence, reflections from dielectric surfaces are fairly strongly polarized. At Brewster's angle, R_p goes to 0, and for glass $R_s \approx 10\%$ per surface. The effect in transmission is not strong enough to qualify as a polarizer unless it's used intra-cavity, but can be enhanced by doing it many times. A pile of microscope slides at 55° or so makes a moderately effective polarizer, and (as we saw in Section 5.4.4) a $(HL)^m H$ stack of dielectric films can be highly effective.

6.5.1 Pile-of-Plates Polarizers

Assuming that the light is low enough in coherence that etalon fringes can be ignored, m glass plates stacked together and oriented at θ_B will attenuate the s -polarized light by a factor of 0.8^m , which for 31 plates amounts to 10^{-3} , ideally with no loss at all in the p polarization. This nice property is of course degraded as θ_i departs from θ_B , but it's useful over a reasonable angular range. The transmitted wavefront fidelity of a pile-of-plates polarizer is poor because of accumulated surface error and multiple reflections of the s -polarized beam between plates. The reflected beam is even worse; surface error affects reflected light more than transmitted, and the multiple reflections are not superimposed. The only real advantages are high power handling capability and ready availability of materials.

[†]The word *dichroic* has been given so many different meanings in optics that it's now next to useless.

6.5.2 Multilayer Polarizers

Alternating layers of high and low refractive index can be made into an effective polarizer, similar to the pile of plates but without the beam quality disadvantage. The similarity is not perfect, because interference effects cannot be ignored in thin films, even for white light.

We saw this trick in Section 5.4.4 with polarizing cubes, but such a film stack can also be deposited on a glass plate, forming a polarizing plate beamsplitter. Since there is no optical cement, the damage threshold of these devices is high, making them a good match for powerful pulsed lasers, such as ruby (694 nm) and Nd:YAG (1064 nm). They are rarely used elsewhere, for three main reasons: Brewster incidence is very oblique, so that the reflected light comes off at an odd angle; the angular alignment is critical (as in all Brewster polarizers), and there is no obvious cue for rough alignment as there is in polarizing cubes; and the large index discontinuity at the top surface of the film reflects an appreciable amount of *p*-polarized light, making the polarization purity of the reflected wave poor.

6.5.3 Polarizing Cubes

Next to film polarizers, the most common type of polarizer in lab drawers is the polarizing beamsplitter cube, which we discussed at length in Sections 4.7.2 and 5.4.5. These are superficially attractive devices that in the author's experience cause more flaky optical behavior than anything else, barring fiber.

6.6 BIREFRINGENT POLARIZERS

Birefringent materials can be used in various ways to make polarizers. The three main classes use (best to worst) double refraction, walkoff, and TIR.

Crystal polarizers are usually made of calcite because of its high birefringence, good optical properties, and reasonable cost. Quartz is sometimes used, but its optical activity causes quartz prisms to exhibit weird polarization shifts versus wavelength, field angle, and orientation—it matters which way round you use a quartz polarizer. None of these crystal devices is cheap, so use them only where you need really good performance. The CVI Laser catalog has an extensive discussion of polarizing prisms.

6.6.1 Walkoff Plates

A very simple polarizer or beam displacer based on beam walkoff (Section 6.3.5) can be made from a thick plate of birefringent material whose optic axis is not parallel to its faces, as in Figure 6.1. When a beam comes in near normal incidence, the *o*-ray passes through such a plate nearly undeviated, whereas the *e*-ray walks off sideways. The two are refracted parallel to the original incident beam when they leave the plate. Walkoff plates are inexpensive, because single plane-parallel plates are easy to make, and because no great precision is required in the orientation of the optic axis if only the *o*-ray is to be kept. This technique is frequently used in optical isolators for fiber applications, where cost is a problem and the angular acceptance is small. Note that the optical path length seen by the two beams is very different, so using walkoff plates as beamsplitters in white-light interferometers is difficult. The shift is a small fraction

of the length of the prism but works over a large range of incidence angles; thus the étendue is small if you need the beams to be spatially separated, but large if overlap is OK.

6.6.2 Savart Plates

The walkoff plate can be made more nearly symmetrical by putting two of them together to make a Savart plate. These consist of two identical flat, square plates of quartz, calcite, or LiNbO_3 whose optic axes are oriented at 45° to the surface normal. The plates are rotated 90° to each other and cemented together. (One plate's optic axis lies in the plane of the top edge and the other one's in the plane of the side edge.)

An *o*-ray in the first plate turns into an *e*-ray in the second, and vice versa, so that the two polarizations are offset by the same amount from the axis, in opposite directions, and emerge parallel to their initial propagation direction. At normal incidence, they have zero path difference and hence produce white-light fringes if they overlap.

Away from normal incidence, these are not zero path difference devices, since the *e*-ray is polarized at a large angle from the optic axis, the path difference changes linearly with angle, rather than quadratically as in a Wollaston prism; Section 19.1.1 has an example where this seemingly obscure point caused a disaster.

6.7 DOUBLE-REFRACTION POLARIZERS

Double-refraction polarizers exploit the different index discontinuity seen by *e*- and *o*-rays at an interface. Generally they have excellent performance, but like other refracting prisms their deflection angles change with λ , and they anamorphically distort the beam to some degree.

6.7.1 Wollaston Prisms

A Wollaston prism consists of two wedges of calcite, with their optic axes oriented as shown in Figure 6.2a (the diagram shows a Wollaston prism made from a positive uniaxial crystal such as quartz). A beam entering near normal incidence is undeviated until it encounters the buried interface. There, the *e*-ray will see the index go down at the surface and so will be refracted away from the normal, whereas the *o*-ray will see an index increase and be refracted toward the normal by nearly the same amount. (The bending goes the other way for negative uniaxial crystals.) Both beams hit the output facet at a large angle and so are refracted away from the axis.




Figure 6.2. Double-refraction polarizers have the best extinction and purity of any type. (a) Wollaston prisms have no etalon fringes. (b) Rochon prisms have one undeviated beam.

The result is a polarizing prism of unsurpassed performance: polarization ratios of 10^{-6} , nearly symmetrical beam deviation, and, crucially, no internal back-reflection to cause etalon fringes. The beams are not perfectly symmetrical because of the asymmetric incidence on the buried interface. The phase shift between beams is linear in y , independent of x , and varies only quadratically in the field angle, since the optic axis lies in the plane of the prism faces, making Wollastons good for interferometers. You can find a lateral position where the OPD between the two beams is zero, and moving the prism sideways makes a nice phase vernier. Quartz Wollastons have beam separations of 1° to 3.5° , while calcite ones are typically 10° to 20° , and special three-element calcite ones can achieve 30° . Wollastons have excellent étendue on account of their wide angular acceptance.

6.7.2 Rochon Prisms

Table 6.1 shows that the refractive indices of calcite are in the range of optical glass. You can thus make a Wollaston-like prism by using one calcite wedge and one glass one (with $n_{\text{glass}} \approx n_o$), as shown in Figure 6.2. The difference is that one beam is undeviated, as in a Glan–Thompson. Rochon prisms suffer from severe etalon fringes due to the undeviated path, but the polarization purity is similar to a Wollaston, and because of the larger angular acceptance of double refraction, you can use the Rochon tipped fairly far to avoid the back-reflection. The undeviated beam is an o -ray and so has no major chromatic problems.

Some Rochons (and the closely related Senarmont prisms) are made with the glass wedge replaced by a chunk of calcite with its optic axis normal to the output facet, so that both rays see nearly the o -ray index in the second wedge. It has the optical properties of the Rochon without its cost advantage, but is better color-corrected and less likely to delaminate due to differential thermal expansion.

Because of the variations of n_e with incidence angle, a prism (like the Rochon) that transmits the o -ray undeviated is probably superior in imaging applications, as it is easier to correct the resulting aberrations, particularly astigmatism, anamorphic distortion, and chromatic aberration.

Aside: Quartz Rochon Prisms. Occasionally a Rochon prism will surface that has both wedges made of α quartz; unfortunately, optical activity in the output wedge will

TABLE 6.1. Properties of Common Birefringent Materials

Material	n_o	(n_\perp)	n_\parallel	$n_\parallel - n_\perp$	$\lambda(\mu\text{m})$
Magnesium fluoride	1.3783	1.3906	0.0123	540	0.12–8
α Crystal quartz SiO_2	1.5462	1.5554	0.0092	546	0.15–2, 3.3–6
Sapphire Al_2O_3	1.768	1.760	–0.0080	589.3	0.2–6
Potassium dihydrogen phosphate (KDP)	1.5125	1.4705	–0.0420	540	0.2–1.5
Ammonium dihydrogen phosphate (ADP)	1.5274	1.4814	–0.0460	540	0.2–2.0
Barium titanate (BaTiO_3)	2.459	2.400	–0.0590	550	0.4–2
Lithium niobate (LiNbO_3)	2.3165	2.2285	–0.0880	550	0.3–8
Calcite (CaCO_3)	1.658	1.486	–0.1720	589.3	0.25–2.5
Rutile (TiO_2)	2.64	2.94	–0.3000	560	0.4–9




Figure 6.3. Wollaston prisms have a variety of uses: (a) beam splitting, (b) beam combining, (c) heterodyne interferometry, and (d) solid Fourier transform spectrometer.

severely degrade the polarization purity—not only rotating it, but rotating light at different locations and wavelengths by different amounts.

6.7.3 Cobbling Wollastons

Wollaston prisms are usually used as etalon-fringe-free polarizing beamsplitters, but that doesn't exhaust their usefulness. They can of course be run backwards as beam combiners, where two input beams are combined into a single output beam. More generally, a Wollaston can be used to make two beams of orthogonal polarization cross outside the prism, as shown in Figure 6.3.

Due to the slightly unequal angular deviations of the two beams, Wollaston prisms have a path difference that depends on position; an interferometer can have its operating point[†] adjusted by moving the prism back and forth slightly in the direction of the wedges. Note that if you're using the prism to combine beams of any significant NA, this asymmetry makes the plane of the fringes not exactly perpendicular to the axis, so that some in-out motion is needed to maintain the same average optical path in the two beams as you move the prism sideways.

Wollastons are quite dispersive, so they aren't especially useful for imaging in white light unless the beam separation is very small, as in *schlieren interferometry*, which takes advantage of the colored fringes formed by very small-angle Wollastons to make images of weak phase objects such as air currents.

Example 6.1: Solid Fourier Transform Spectrometer. Interferometers based on Wollaston prisms have been used fairly widely. One interesting approach is static Fourier transform interferometry, similar to an FTIR spectrometer (see Section 10.5.6) but with no moving parts. The wide angular acceptance of Wollastons makes high étendue interferometers easy to build. The limiting factor in such interferometers is the birefringence

[†]The operating point of an interferometer, amplifier, or what have you is the nominally stable point of a nonlinear response about which the (supposedly small) signal excursions occur.

of the plates, which makes the transmission phase of an off-axis ray a peculiar function of the angle and limits throughput. Using a positive uniaxial crystal (e.g., quartz) for the splitting Wollaston and a negative uniaxial one (e.g., ammonium dihydrogen phosphate) for the recombiner results in the birefringence canceling out, so that the full étendue is available[†].

6.7.4 Nomarski Wedges

A Nomarski wedge is a modified narrow-angle Wollaston prism. In the Wollaston, the beams diverge from the middle of the prism, so that building an interference microscope requires the microscope objective's exit pupil to be outside the lens, whereas it's usually just inside. The Nomarski prism has the optic axis of one of the wedges oriented out of the plane of the end face; thus the *e*-ray walks off sideways far enough that its exit pupil (where the *e*- and *o*-rays cross and the white-light fringes are located) is outside the prism. In a symmetrical system using two Nomarski wedges, the path difference between the two beams is zero, as is required for achromatic differential interference contrast (DIC) measurements.

6.7.5 Homemade Polarizing Prisms

The one serious drawback to birefringent prisms is that they're expensive, especially in large sizes. You can make an adjustable prism similar to a Wollaston of a few milliradians by bending a bar of plastic such as polycarbonate. Stress birefringence splits the polarizations, and the bending causes the bar to become trapezoidal in cross section (by the local strain times Poisson's ratio) so that the two polarizations are refracted in slightly different directions. This has been used in schlieren interferometers.[‡] The material creeps, so these prisms aren't very stable, and cast plastic isn't too uniform, so they have to be used near an image.

6.8 TIR POLARIZERS

The third major class of birefringent polarizers is based on TIR at a thin layer between two prisms of birefringent material. Because n_e and n_o are different, the *e*- and *o*-ray critical angles will be different as well, so that we can transmit one while totally reflecting the other. In a sufficiently highly birefringent material, the difference is large enough to be useful, although TIR polarizers always have a much narrower angular acceptance than double-refraction ones. The exit angles are set by *o*-ray reflections, so they are pretty well achromatic as long as the exit face is close to perpendicular to the beam (see Section 6.8.1).

The small angular acceptance leads to small étendue, and there are some familiar drawbacks such as poor polarization purity in the reflected beam.[§] A more subtle problem is

[†]D. Steers, B. A. Patterson, W. Sibbett, and M. J. Padgett, Wide field of view, ultracompact static Fourier transform spectrometer. *Rev. Sci. Instrum.* **68**(1), 30–33 (January 1997).

[‡]S. R. Sanderson, *Rev. Sci. Instrum.* **76**, 113703 (2005).

[§]It's a bit more complex than in a polarizing cube, because an oblique reflection at the TIR surface can mix *e*- and *o*-rays. Real prisms are carefully cut to avoid this effect.

that since calcite has a negative birefringence, it's the e -ray that is transmitted undeviated, and in imaging applications, its wavefronts are aberrated by the variation of n_e with angle (see the earlier Rochon prism discussion). All in all, TIR polarizers are inferior to double-refraction types for most uses.

6.8.1 Refraction and Reflection at Birefringent Surfaces

When calculating the behavior of obliquely incident light at the surface of a birefringent material, life gets somewhat exciting unless the optic axis is perpendicular to the plane of incidence. When doing this sort of problem, remember that phase is phase is phase—you calculate phase matching at the interface based on the \mathbf{k} vector of the incoming beam, period. The angle of reflection is a consequence of the fundamental physics involved, namely, the phase matching condition, which remains in force.

For example, imagine making the calcite prism of Figure 6.4a out of MgF_2 , so that the e -ray is now the one reflected. Light coming in s -polarized is a pure o -ray, but the p -polarized light starts out as a high index e -ray and winds up as a low index e -ray (if the wedge angle were 45° it would wind up as an o -ray). Thus the value of k changes, so the “law of reflection” is broken: $\theta_r \neq \theta_i$.

6.8.2 Glan–Taylor

As shown in Figure 6.4a, we can make a TIR polarizer from a simple triangular calcite prism, with the optic axis lying parallel to one edge of the entrance face, and with a wedge angle α whose complement lies between the e - and o -ray critical angles. It has calcite's wide transmission range (220–2300 nm), and because there are no cemented joints, its laser damage threshold is high, $100\text{W}/\text{cm}^2$ or so. This simple device has some serious disadvantages, too; as in the beamsplitter cube, the reflected polarization purity is poor, but there are more. The transmitted beam exits near grazing, because the two refractive indices are not very different; it is anamorphically compressed, which is usually




Figure 6.4. TIR polarizing prisms: (a) simple calcite prism ($n_e < n_o$); (b) quartz Glan–Taylor ($n_e > n_o$) adds an air-spaced second prism to straighten out the transmitted beam; and (c) calcite Glan–Thompson ($n_e < n_o$) uses cement with $n_e < n < n_o$.

undesirable, and it also exhibits chromatic dispersion. The angle of incidence can range only between the e - and o - ray critical angles, which limits the field severely.

If we put two such prisms together, as in Figure 6.4b, we have the Glan–Taylor prism. Here the first prism does the polarizing, and the second one undoes the beam deviation and almost all the chromatic dispersion. The Glan–Taylor keeps the wide transmission and high damage threshold of the calcite prism, but the stray light is quite a bit worse due to multiple bounces in the air gap, where the reflection coefficient is high due to the oblique incidence.

You don't usually try to use the reflected beam from a Glan–Taylor for anything, because its polarization is impure and changes with angle due to the birefringence of the entrance prism. For laser applications, you can cut the prisms at Brewster's angle for the e -ray, which lies above the o -ray critical angle. The resulting Glan-laser prism reduces the angular acceptance while improving the multiple reflections and stray light.

6.8.3 Glan–Thompson

The Glan–Thompson prism of Figure 6.4c is made by cementing two skinny calcite wedges together, hypotenuse to hypotenuse, using cement with an index of about 1.52. The superficial similarity to the Glan–Taylor is somewhat misleading; because n_{glue} is between the e - and o -ray indices, the e -ray cannot experience TIR, so the Glan–Thompson has a much wider angular acceptance than the Glan–Taylor, even though it is longer and skinnier.

As in the Glan–Taylor, the first prism does the polarizing, and the second straightens out the transmitted beam, so that the transmitted beam is undeviated. Because the indices of the calcite and the cement are not very different ($n_o = 1.655$, $n_{\text{glue}} \approx 1.52$), this requires near-grazing incidence, making the Glan–Thompson prism rather long for its aperture.

The o -ray makes two or three TIR bounces in the entrance prism, so that's a good place for a beam dump; Glan–Thompson prisms are usually embedded in hard black wax (a reasonable index match to the o -ray in calcite), so that only the undeviated beam emerges. With a four-sided prism on the entrance side, the reflected ray can be allowed to leave through its own facet, near normal incidence. This configuration is called a beamsplitting Thompson prism and is quite a good device; the closer index match at the interface makes the reflected polarization purer, and the reflected light doesn't see any serious birefringence since its polarization direction is unaltered. Nonetheless, Glan–Thompson prisms share most of the disadvantages of polarizing cubes, including strong etalon fringes and low damage threshold ($\approx 1 \text{ W/cm}^2$) due to the glue, and reflected light polarization purity inferior to that of double-refraction polarizers.

6.9 RETARDERS

The polarization of a monochromatic electromagnetic wave can be decomposed in terms of two arbitrary orthonormal complex basis vectors. This means that, for example, a linearly polarized light beam can be expressed as the sum of two orthogonal circularly polarized beams (of right and left helicity), and vice versa. The devices of this section all use this property to apply different phase delays to different polarization components.

6.9.1 Wave Plates

Retarders or *wave plates* are the simplest application of birefringence. A uniaxial plate of thickness d with its optic axis parallel to its faces will delay a normally incident o -ray by

$$\Delta t_o = n_o \frac{d}{c} \quad (6.8)$$

and similarly for the e -ray, so that the phases of the two are shifted by

$$\delta = (n_e - n_o) \frac{\omega d}{c}. \quad (6.9)$$

Retarders based on this principle are called *wave plates*. When $\Delta\phi$ is $\lambda/4$ for a given wavelength, you have a *quarter-wave plate*, and when it's $\lambda/2$, a *half-wave plate*; these are the two most useful kinds. Note that in the absence of material dispersion, this retardation is a pure time delay, so that the phase varies rapidly with λ . As we saw in Chapter 4, there are also retarders based on the phase shift upon TIR whose phase shift is nearly constant with λ . There are also achromatic wave plates made of multiple polymer layers, whose phase shift is reasonably independent of λ . As usual, there is a three-way trade-off between complexity (i.e., cost and yield), degree of achromatism, and bandwidth.

6.9.2 Quarter-Wave Plates

Quarter-wave plates are good for converting linear to circular polarization and back. Consider a $\lambda/4$ plate lying in the xy plane whose fast axis is \mathbf{x} , with a plane wave passing through it with \mathbf{k} parallel to \mathbf{z} . If \mathbf{E} is at $\pi/4$, the e -ray and o -ray amplitudes will be equal. These get out of phase by $\pi/2$ crossing the plate, so that the field exiting is

$$\mathbf{E}_{\text{out}} = \frac{E_{\text{in}}}{\sqrt{2}} (\hat{\mathbf{x}} \cos \omega t + \hat{\mathbf{y}} \sin \omega t), \quad (6.10)$$

which is left-circular polarization (goes like a right-handed screw). Putting the slow axis along x (or equivalently putting \mathbf{E} at $-\pi/4$) makes right-circular polarization instead.

6.9.3 Half-Wave Plates

A half-wave plate delays one Cartesian component by half a cycle with respect to the other, which reflects \mathbf{E} through the fast axis. This is very useful where both the initial and final polarization states are linear—you twist the wave plate until the polarization lines up just right. Linear polarization stays linear, but with circular or elliptical polarization, the helicity gets changed, so right and left circular are exchanged. Figure 6.5 shows how this works.

Combining differently oriented retarders is most easily done with rotation matrices; a retarder of $\beta\lambda$ whose slow axis is at θ_1 can be written

$$R(\beta, \theta_1) = \begin{bmatrix} \cos \theta_1 & \sin \theta_1 \\ -\sin \theta_1 & \cos \theta_1 \end{bmatrix} \begin{bmatrix} \exp(i2\pi\beta) & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \theta_1 & -\sin \theta_1 \\ \sin \theta_1 & \cos \theta_1 \end{bmatrix}, \quad (6.11)$$

which we'll come back to a bit later.




Figure 6.5. Retarders. Here two quarter-wave plates ($n_{\perp} < n_{\parallel}$) turn a linearly polarized beam first into circular, then into the orthogonal linear polarization, just as a half-wave plate would.

6.9.4 Full-Wave Plates

The retardation of a wave plate depends on its rotations about x and y , so that small errors can be removed by tilting it. For a uniaxial material, the retardation is increased by rotating it around the optic axis, since n_o and n_e are unaltered and the path length is increased; tipping the other way also increases it, because in almost all uniaxial materials, $|n_e - n_o|$ declines more slowly than the path length grows.

A full-wave plate nominally does nothing at all, but in fact the dependence of the retardation on the tipping of the plate makes it a sensitive polarization vernier control, able to provide small amounts ($< \lambda/8$) of retardation to balance out other sources of birefringence in the beam.

6.9.5 Multiorder Wave Plates

Much of the time, it is impractical to make the plate thin enough for such small retardations (a visible quarter-wave plate made of quartz would be only $20 \mu\text{m}$ thick). For narrowband light propagating axially, it is sufficient that the retardation be an odd multiple of $\lambda/2$ for a half-wave plate or $\lambda/4$ for a quarter-wave plate. It is more practical to make a 20.25λ plate (say), which has just the same effect.

Neglecting the change in n_e with angle, the retardation goes approximately as the secant of the incidence angle, so that a 20.25λ multiorder plate will have an étendue over three orders of magnitude smaller than a 0.25λ zero-order one for a given retardation tolerance, and chromatic and temperature shifts 80 times larger as well. This isn't usually a problem with laser beams used in the lab but is serious when you need to run a significant aperture or work in field conditions. One good thing about multiorder wave plates is that they can be made to work at more than one wavelength; for example, you can get quarter-wave plates that work at 633 nm and 1064 nm. It's easier to make them when the wavelengths are significantly different. That kind are often cut with the optic axis slightly out of the plane of the end faces, to make the ratios come out just right with a reasonable thickness. Small amounts of walkoff will result.

6.9.6 Zero-Order Wave Plates

The poor étendue of multiorder wave plates can be fixed by putting two of them together, oriented at 90° to each other so that the aperture-dependent phase shifts largely cancel.

We might laminate a 20.38λ plate with a 19.88λ plate to make a zero-order 0.50λ plate.

6.9.7 Film and Mica

The cheapest retarders are made from very thin layers of mildly birefringent materials such as PVA film or mica. These materials are easily prepared in very thin sheets, and since only a small retardation is desired, only a little is needed. Unfortunately their accuracy, uniformity of retardation across the field, and transmitted wavefront distortion have historically been poor. On the other hand, because even the cheap ones are usually zero-order devices, their dispersion and angular sensitivity are small. Polymer retarders are enjoying a bit of a renaissance, because when carefully made (usually in two-layer stacks) they can have unique properties, such as decent achromatism. In a low coherence imaging system, you can get around their wavefront distortion by putting them near a focus.

6.9.8 Circular Polarizers

The usual polarizer types resolve the incoming polarization state into two orthogonal linear polarizations. This is not of course the only choice; since an arbitrary linear polarization can be built up out of right and left circularly polarized light with a suitable phase shift between them, it follows that we can use these as basis vectors as well. Circular polarizers are quite useful for controlling back-reflections (e.g., glare on glass panels). They're made by laminating a film polarizer to a polymer film quarter-wave plate, with the fast axis of the plate at 45° to the polarization axis. One serious gotcha is that there's usually a retarder on only one side. If used in one direction, that is, with the polarizer turned toward an unpolarized light source, this does result in a circularly polarized output and will indeed attenuate reflected light returned back through it; but it won't pass a circularly polarized beam through unchanged, nor will it work if it's turned round. For that you need two retarders, one on each side.

6.10 POLARIZATION CONTROL

6.10.1 Basis Sets for Fully Polarized Light

We saw in Section 6.2.5 that light could be expressed in terms of Cartesian basis vectors, the *Jones vectors*:

$$\mathbf{E}_\perp = \begin{bmatrix} E_x \\ E_y \end{bmatrix} = E_x \hat{\mathbf{x}} + E_y \hat{\mathbf{y}}. \quad (6.12)$$

A similar decomposition in terms of circularly polarized eigenstates is useful in discussing optical activity and Faraday rotation. A plane wave propagating toward positive Z with complex electric field $\tilde{\mathbf{E}}_\perp$ can be decomposed as

$$\tilde{\mathbf{E}}_\perp \equiv \begin{bmatrix} \tilde{E}_x \\ \tilde{E}_y \end{bmatrix} = \frac{\tilde{E}_L}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix} + \frac{\tilde{E}_R}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}, \quad (6.13)$$

TABLE 6.2. Jones Matrix Operators for Common Operations

Coordinate rotation of θ	$\begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}$
δ Radian retarder, slow axis along \mathbf{x}	$\begin{bmatrix} e^{i\delta/2} & 0 \\ 0 & e^{-i\delta/2} \end{bmatrix}$
Analyzer along \mathbf{x}	$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$
Analyzer at θ to \mathbf{x}	$\begin{bmatrix} \cos^2 \theta & \sin \theta \cos \theta \\ \sin \theta \cos \theta & \cos^2 \theta \end{bmatrix}$

where left- and right-circular components \tilde{E}_L and \tilde{E}_R are given by

$$\tilde{E}_L = \mathbf{E} \cdot \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix} \quad \text{and} \quad \tilde{E}_R = \mathbf{E} \cdot \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix}. \quad (6.14)$$

Linearly polarized light has $\tilde{E}_L = e^{i\phi} \tilde{E}_R$, where ϕ is the azimuthal angle of \mathbf{E} measured from the x axis. Polarization gizmos can be represented as 2×2 matrix operators; in Cartesian coordinates, the result is the *Jones matrices* shown in Table 6.2.

Like *ABCD* matrices, these are not powerful enough to model everything without hand work; for example, treating reflections takes some care. We model optical activity and Faraday rotation as coordinate rotations, but since one adds and the other cancels on the way back, we have to be very careful about the bookkeeping; we'll do an example of this later when we consider the Faraday rotator mirror.

6.10.2 Partial Polarization and the Jones Matrix Calculus

Light can be partially polarized as well, most often by reflecting thermal light from a dielectric. This polarization is often fairly strong, especially when the reflection takes place near Brewster's angle, as we know from the effectiveness of Polaroid sunglasses in reducing glare from water, car paint, and glass. Rayleigh scattering also produces partially polarized light; try your sunglasses on a clear sky—when you look at 90° to the sun, the polarization is quite pronounced.

The vectors that are adequate for describing full polarization fail for partial polarization, which is an intrinsically more complicated situation. If the direction of \mathbf{k} is prespecified, we can express the polarization properties of a general narrowband beam as a four-dimensional vector (the *Stokes parameters*, see Born and Wolf) or by a 2×2 *coherency matrix*.[†] As Goodman explains,[‡] the coherency matrix formulation lets us follow the polarization state of our beam through the system by matrix multiplication of \mathbf{E}_\perp by the accumulated operator matrices, written in reverse order, just the way we did ray

[†]The two are very closely related; the elements of the coherency matrix are linear combinations of the Stokes parameters.

[‡]Joseph W. Goodman, *Statistical Optics*. Wiley, Hoboken, NJ, 1986.

TABLE 6.3. Coherency Matrices for Some Polarization States

Linear x	$E_0 \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$	Linear y	$E_0 \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$
Right circular	$E_0 \begin{bmatrix} 1 & i \\ -i & 1 \end{bmatrix}$	Left circular	$E_{0/2} \begin{bmatrix} 1 & -i \\ i & 1 \end{bmatrix}$
Unpolarized	$E_{0/2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$		

tracing with the $ABCD$ matrices, and it is easily connected to time-averaged polarization measurements. The coherency matrix $\underline{\mathbf{J}}$ is the time-averaged direct product $\underline{\mathbf{E}}\underline{\mathbf{E}}^{\text{T}*}$:

$$\underline{\mathbf{J}} = \begin{bmatrix} J_{xx} & J_{xy} \\ J_{xy}^* & J_{yy} \end{bmatrix} = \begin{bmatrix} \langle \tilde{E}_x \tilde{E}_x^* \rangle & \langle \tilde{E}_x \tilde{E}_y^* \rangle \\ \langle \tilde{E}_y \tilde{E}_x^* \rangle & \langle \tilde{E}_y \tilde{E}_y^* \rangle \end{bmatrix}. \tag{6.15}$$

It’s easy to see from the definition that (up to a constant factor) J_{xx} and J_{yy} are the real-valued flux densities you’d measure with an analyzer along $x(I_{0^\circ})$ and $y(I_{90^\circ})$, respectively. The complex-valued J_{xy} is related to the flux density I_{45° that you get with the analyzer at 45° , and the I'_{45° you get by adding a $\lambda/4$ plate with its slow axis along y before the analyzer,

$$J_{xy} = I_{45^\circ} - \frac{1}{2}(I_{0^\circ} + I_{90^\circ}) + i \left(I'_{45^\circ} - \frac{1}{2}(I'_{0^\circ} + I'_{90^\circ}) \right) \tag{6.16}$$

Table 6.3 has coherency matrices for some special cases.

6.10.3 Polarization States

It is commonly held that when you superpose two beams, their $\underline{\mathbf{J}}$ s add, but that assumes that they are mutually incoherent, which is far from universally true. You’re much safer sticking closer to the fields unless you know *a priori* that the waves you’re combining are mutually incoherent but still narrowband enough for the Jones matrix approach to work.

A couple of mathematical reminders: a lossless operator $\underline{\mathbf{L}}$ is unitary—all its eigenvalues are on the unit circle and $\underline{\mathbf{L}}\underline{\mathbf{L}}^\dagger = \underline{\mathbf{I}}$, that is, $\underline{\mathbf{L}}^\dagger = \underline{\mathbf{L}}^{-1}$, where the adjoint matrix $\underline{\mathbf{L}}^\dagger$ is the complex conjugate of the transpose, $\underline{\mathbf{L}}^\dagger = (\underline{\mathbf{L}}^{\text{T}})^*$. These lists can be extended straightforwardly by using the definition (6.15) and matrix multiplication. Remember that although the operators multiply the *fields* (6.13) directly, applying a transformation to $\underline{\mathbf{J}}$ or $\underline{\mathbf{E}}$ requires applying it from both sides; if $\mathbf{E}'_\perp = \underline{\mathbf{L}}\mathbf{E}_\perp$,

$$\langle \mathbf{E}'(\mathbf{E}')^{\text{T}*} \rangle = (\underline{\mathbf{L}}\mathbf{E})(\underline{\mathbf{L}}^*\mathbf{E}^{\text{T}*})^{\text{T}} = \underline{\mathbf{L}}\underline{\mathbf{J}}\underline{\mathbf{L}}^\dagger. \tag{6.17}$$

It’s worth writing it out with explicit indices and summations a few times, if you’re rusty—misplacing a dagger or commuting a couple of matrices somewhere will lead you down the garden path otherwise.

6.10.4 Polarization Compensators

Extending the zero-order quartz wave plate idea, we can make one plate variable in thickness by making it from two very narrow wedges rather than a single plate, yielding the *Soleil compensator*. Provided the two wedges are identical, the retardation is constant across the field and can be adjusted around exactly 0 by sliding one of them, which is a great help in wide-field and wideband applications, for example, looking at small amounts of stress birefringence with a nulling technique. Exactly zero retardation is important only in such cases; in narrowband, low NA systems, it's enough to have the retardation be $0 \bmod 2\pi$, and these compensators are no better than a pair of quarter-wave plates.

6.10.5 Circular Polarizing Film for Glare Control

Laminated circular polarizers of moderate performance can be made cheaply in very large sizes, which makes them attractive as a glare-reduction device for instrument displays; room light passing through the polarizer and then being reflected is absorbed on the return trip. Before CRT screens were AR coated efficiently, circular polarizers were very popular computer accessories, and they remain useful in similar situations.

6.10.6 Polarization Rotators

Optically active materials such as quartz or sugar solutions can be used for polarization rotators. Those made from amorphous materials (e.g., Karo corn syrup) have no birefringence, and so linearly polarized light stays linearly polarized regardless of wavelength, field angle, or what have you. They're inconvenient to make, hard to adjust, and ten times more dispersive than half-wave plates, so apart from special situations such as optical diodes, they have few compelling advantages.

6.10.7 Depolarizers

It is impossible to reproduce the extremely rapid polarization changes of thermal light when starting from a more coherent source such as a laser beam. Devices that purport to depolarize light never do it well enough for that; they just produce a speckle pattern varying more or less rapidly in time or space. If your experiment is slow enough, this may suffice, but in fact it rarely does. Polarized light is something we just have to live with.

There are two classes of depolarizers: wave plates whose retardation varies rapidly across their faces (e.g., Cornu depolarizers), and moving diffusers, such as a disc of ground glass spun rapidly on a shaft. A Cornu depolarizer can do a reasonable job on wideband light, providing the 2π period of the polarization change is sufficiently rapid and the spatial resolution sufficiently low.

Fixed depolarizers help to eliminate the mild polarization dependence of some optical instruments, for example, PMTs, grating spectrometers, and so on, when used with broadband light that may be partially polarized. They do a good enough job for that, certainly.

The spinning ground glass technique often tried with laser beams is much less successful: all you get is a rapidly rotating speckle pattern, which causes a whole lot of

noise. Unlike the situation in Section 2.5.3, the order-1 changes in instantaneous intensity at any point are not smeared out over a bandwidth of hundreds of terahertz, but concentrated in a few hundred kilohertz; the result is not pretty. The rotating speckle pattern also produces undesired correlations in space. These correlations can be reduced by using two discs rotating in opposite directions; if this is properly done, the speckles no longer appear to rotate. Doing it really properly is not trivial, however, and anyway the huge intensity noise remains. If you are forced to do your measurement this way, be prepared to integrate for a long time; it is normally preferable to use polarization flopping, where you do separate measurements in horizontal and vertical polarization and combine the two, or use a rotating $\lambda/2$ plate and integrate for a whole number of periods.

6.10.8 Faraday Rotators and Optical Isolators

Faraday rotators are straightforward applications of the Faraday effect, using a magneto-optic crystal such as YIG in a magnetically shielded enclosure with carefully placed, stable permanent magnets inside providing a purely axial field in the crystal. The main uses of Faraday rotators are in optical isolators and in transmit/receive duplexing when the returned light is not circularly polarized, for one reason or another.

These devices, shown in Figure 6.6, all rely on nonreciprocal polarization rotation. The simple isolator uses two walkoff plate polarizers, oriented at 45° to one another, and a 45° Faraday rotator. Light making it through the first rotator gets its \mathbf{E} rotated through 45° on the first pass, so that it is properly oriented to pass through the second polarizer without loss. Light coming the other way has to make it through the second polarizer and is then rotated 45° in the same direction, putting it at 90° to the first polarizer, so that none gets through. Ideally the isolation would be perfect, but it is more typically 30 dB per isolator, with a loss of about 1 dB.

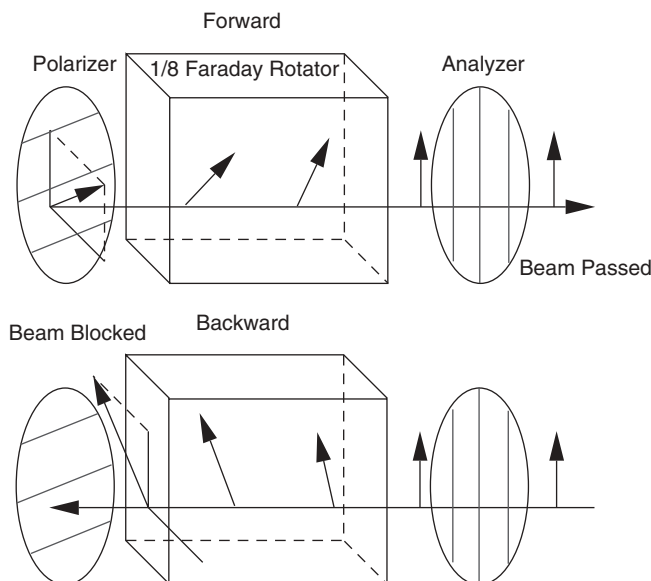


Figure 6.6. Two polarizers plus a 45° Faraday rotator make an optical isolator.

This simplest Faraday isolator requires fully polarized input light, but polarization-insensitive ones can also be made; since you have to use two different paths, it isn't trivial to preserve the input polarization in the process, unfortunately.

The most important uses of Faraday isolators are preventing feedback-induced instability in diode lasers and in preventing high-finesse Fabry–Perot cavities from pulling the frequency of sources being analyzed as the F-P scans.

It is possible to build a *circulator*, an M -port device where the input on port m goes out port $m + 1 \pmod{M}$. Circulators are common in microwave applications but rare in optics. A related device is the *optical diode*, a 45° Faraday rotator plus a -45° optically active cell, used in ring laser cavities; only one propagation direction is an eigenstate of polarization, so that the ring lases in one direction only (the other one gets killed by the Brewster windows).

6.10.9 Beam Separators

A polarizing beamsplitter such as a cube or Wollaston, plus a $\lambda/4$ plate, makes a beam separator, very useful for separating the transmit and receive beams in an interferometer or scanning system. The wave plate is aligned at 45° to the axes of the beamsplitter, as shown in Figure 6.7. On the transmit side, p -polarized light passes through the prism and gets turned into left-circular polarization. Specularly reflected light comes back right-circular, and so the $\lambda/4$ plate turns it into s -polarized light in the cube, which is reflected. If the components are lossless, the wave plate accurate, and the incoming light perfectly p -polarized, the beam suffers no loss whatever in its round trip.

6.10.10 Lossless Interferometers

In Section 4.8.1, we saw that sending light on two passes through a nonpolarizing beamsplitter costs you a minimum of 75% of your light. That's 93.75% of your detected electrical power, representing an SNR degradation of 6 dB in the shot noise limit and 12 dB in the Johnson noise limit—and that's in the best case, with a beamsplitter without excess loss.




Figure 6.7. A polarizing beamsplitter plus a quarter-wave plate make a beam separator, able to disentangle the transmit and receive beams of an interferometer or scanning system.

If we have a polarized source such as a laser, we can use a beam separator to split and recombine the light as shown in Figure 1.12. The only problem is that the two beams are orthogonally polarized and so don't interfere. The usual solution to this is to put an analyzer at 45° to the two beams, resulting in 100% interference but 6 dB detected signal loss. However, by using a polarizing beamsplitter oriented at 45° , detecting the two pairs of beams separately, and subtracting the resulting photocurrents, we get the equivalent of 100% interference with no signal loss, as in the ISICL sensor of Example 1.12.

6.10.11 Faraday Rotator Mirrors and Polarization Insensitivity

As an application of the Jones matrix calculus, let's look at the Faraday rotator mirror, which is widely used in fiber optics. It consists of a 45° Faraday rotator in front of a mirror, so that the light passes twice through the rotator, and so the total rotation is 90° . We have to complex-conjugate the beam to represent the mirror, because the helicity changes and in this model we've no way of expressing the propagation direction.

In operator representation, this is

$$\begin{aligned} \mathbf{E}' &= \underline{\mathbf{R}}_{\pi/4} (\underline{\mathbf{R}}_{\pi/4} \mathbf{E})^* = \underline{\mathbf{R}}_{\pi/2} \mathbf{E}^* \\ \Rightarrow \mathbf{E}'^* \cdot \mathbf{E} &= \begin{bmatrix} E_x & E_y \end{bmatrix} \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} E_x \\ E_y \end{bmatrix} = 0, \end{aligned} \quad (6.18)$$

that is, the light coming back is polarized orthogonally to the light coming in, *regardless of the incoming polarization*—it works for circular and elliptical, as well as linear. It's obviously orthogonal if the polarization exiting the fiber is linear (courtesy of the Faraday rotation) or circular (courtesy of the mirror). Elliptical polarizations have their helicity inverted by the mirror, and their major axis rotated 90° by the Faraday rotator.

The power of this is that the polarization funnies encountered by the beam, that is, birefringence and optical activity, are all unitary operations, so the incoming and outgoing polarizations remain orthogonal everywhere, as long as they traverse the same path. That means that our optical fiber can misbehave as much as it wants, in theory, and as long as we've got a Faraday mirror at the other end, the round-trip light comes out polarized orthogonally to the incoming light; if we send in vertically polarized light, it comes out horizontally polarized, no matter how many waves of birefringence it encountered. This doesn't work quite as well as we'd like, because the accuracy requirements are very high and it ignores scattering, multiple reflections, and transients. Nonetheless, we can build more-or-less polarization-insensitive fiber interferometers this way.

A slightly more subtle benefit is that the propagation phase is polarization insensitive. A lossless fiber has two orthogonal eigenmodes. If we decompose any incoming polarization into these modes, we find that the Faraday rotator mirror exchanges the components in the two eigenmodes, so that the total round-trip phase delay is the sum of the one-way delays of the eigenmodes. You do have to think about Pancharatnam's phase, though (see Section 6.2.4).

Exotic Optical Components

To know a thing well, know its limits. Only when pushed beyond its tolerance will its true nature be seen.

—Frank Herbert, *Dune*

7.1 INTRODUCTION

The mundane optical tasks of directing and transforming beams are mostly done with lenses and mirrors, with the occasional polarizer or wave plate thrown in, and this is enough for cameras and microscopes. They're the workhorses: steady, dependable, and reasonably docile, but not very fast. Building advanced instruments is more like racing. The more exotic components like fibers, gratings, scanners, and modulators are thoroughbreds—they're good at what they're good at, and aren't much use otherwise. (Watch your step in this part of the paddock, by the way.) Fibers are left until Chapter 8, but there are lots of connections between there and here.

7.2 GRATINGS

A diffraction grating is an optical surface with grooves in it, shaped and spaced so as to disperse incident polychromatic light into a sharply defined spectrum. There are lots of variations, but they're all basically holograms—the grooves reproduce the interference pattern between an incident monochromatic beam and the desired diffracted beam, and so form an optical element that transforms the one into the other. Some gratings are made holographically, and some are ruled mechanically, but the operating principle is the same.

The most common type is the *classical plane grating*, a flat rectangular surface with equally spaced grooves running parallel to one edge. Phase matching at the surface governs their operation; as with a planar dielectric interface, this condition can be satisfied over a broad continuous range of incident \mathbf{k} vectors.

There are also Bragg gratings, where the grating structure runs throughout some volume, a more complex structure that is important in fiber devices, holograms, acousto-optic cells, and some diode lasers, as we'll see later in this chapter.




Figure 7.1. Plane diffraction grating.

7.2.1 Diffraction Orders

We can make this a bit easier to understand by looking at the special case of Figure 7.1. A metallic plane grating lies in the xy plane, with G sinusoidal grooves per unit length,

$$h(x, y) = a \sin(2\pi Gx), \quad (7.1)$$

where $a \ll \lambda$ and G , k_x , and k_y are small compared to k (i.e., a weak grating of low spatial frequency). A plane wave $\exp(i\mathbf{k}_{\text{inc}} \cdot \mathbf{x})$ hits the surface and gets its phase modulated by the changing height of the surface. If the medium is isotropic, linear, and time invariant, the modulus of k can't change,[†] so we've got a pure spatial phase modulation, as in Section 13.3.7. Thus the phase modulation produces mixing products (see Chapter 13) with wave vectors \mathbf{k}_{Dm} ,

$$\mathbf{k}_{Dm} = \mathbf{k}_i + m\mathbf{k}_G, \quad (7.2)$$

where $k_G = 2\pi G\hat{\mathbf{x}}$ and $m = \dots, -1, 0, 1, 2, \dots$. Equation (7.2) can also be derived immediately from the phase matching condition: because the surface is periodic, the fields have to be invariant (apart from an overall phase) if we shift it by an integer number of cycles. Only a finite range of m produces propagating waves at the output—only those whose $|k_{xm}| < k$; that is,

$$\frac{-\sqrt{k^2 - k_y^2} - k_x}{k_G} < m < \frac{\sqrt{k^2 - k_y^2} - k_x}{k_G}. \quad (7.3)$$

Although we've kept k_y in this formula, gratings are nearly always oriented to make k_y as small as possible, so that it enters only quadratically in θ_d . Since G is wavelength independent, we can solve (7.2) (in the special case $k_y = 0$) for λ , yielding the *grating equation*[‡]

$$\lambda_m = \frac{\sin \beta - \sin \alpha}{mG}. \quad (7.4)$$

[†]When we get to acousto-optic modulators, this won't be true anymore, and frequency shifts will occur.

[‡]You sometimes see it used with the other sign convention, so that there is a plus sign in $\Delta u = \sin \theta_d - \sin \theta_i$; in any event, specular reflection ($m = 0$) has $\Delta u = 0$.

The illuminated patch on the grating is the same width for both beams, but because of obliquity, the diffracted beam undergoes an anamorphic magnification in x of

$$M = \frac{\cos \beta}{\cos \alpha}. \quad (7.5)$$

In general, the spectrum gets spread out in a cone, but in the $k_y = 0$ case, it gets spread out in a line, with only a slight curvature due to the inescapable finite range of k_y in real apparatus. If we send broadband light in at θ_i , we can select a narrow wavelength band centered on λ by spatial filtering.

The nonlinear relation of θ_d to θ_i for a given wavelength means that classical plane gratings cause aberrations if the light hitting them is not collimated in the x direction. These aberrations reduce the resolution of the spectrum and cause spectral artifacts, so we normally allow only a single k_x and a limited range of k_y .

Example 7.1: Czerny–Turner Monochromator. A monochromator is a narrowband tunable optical filter, based on the Fourier optics of gratings and slits. The classical design is the Czerny–Turner, shown in Figure 7.3. Polychromatic light from the entrance slit is (spatially) Fourier transformed by spherical mirror M1, so that every point in the slit produces its own plane wave at each wavelength. These are then dispersed by the (rotatable) grating and transformed back by M2 to produce a set of images of the entrance slit on the plane of the exit slit, of which only one is allowed to escape. Rotating the grating moves the dispersed spectrum across the exit slit, so that a different λ emerges.

The spherical mirrors are used well off-axis, so there is a significant amount of coma and astigmatism as well as spherical aberration, which must be accounted for in the design. Note the anamorphic magnification and pupil shift in the figure; normally we use mirrors that are somewhat larger than the grating to combat this.

In designing monochromators and spectrometers, we have to remember that most of the light doesn't make it out the exit slit, but bounces around inside the box, so we need good baffles. Otherwise this stray light would bounce all over the place inside, and some of it would eventually escape through the exit slit and contaminate the output spectrum. There's no place in Figure 7.2 to put a baffle that won't obscure the optical path, so real Czerny–Turners don't have planar layouts. The mirrors are canted down a bit (into the




Figure 7.2. Czerny–Turner monochromator.

page), which depresses the grating position enough to fit a good baffle in over top, which helps a lot.

Another thing to watch out for in grating instruments, especially adjustable ones, is temperature-compensating the slit opening. Carelessness here, such as using brass slits on steel screws and a long mechanical path, typically leads to throughput changes of *several percent per °C*.

The other main problem with spectrometers is that they're all polarization sensitive. The *p-to-s* diffraction efficiency ratio of a transmission grating is massively wavelength dependent, and mirrors used off-axis can easily contribute several percent polarization (see Example 5.2).

7.3 GRATING PATHOLOGIES

So far, a grating is a reasonable spatial analogy to a heterodyne mixer (see Section 13.7.1). The analogy can be pressed further, because the grating also has analogues of LO phase noise (scattered light), LO spurs (periodic errors in the grating phase, giving rise to *ghosts*), and spurs due to LO harmonics (multiple orders, leading to overlap). It starts to break down when the 3D character of the intermodulation starts entering in; spatial frequency differences can arise from shifts in ω or θ_i , but the resulting fields are quite different.

7.3.1 Order Overlap

For any grating and any θ_i , the function $\lambda(\theta_d)$ is multivalued, so that more than one wavelength will make it through a monochromator at any given setting. Simple grating spectrometers are limited to a 1-octave range in the first order, as shown in Figure 7.3, and the limitation gets tighter at higher orders. The best way to reject light that would be aliased is to use *cross-dispersion*: just put a second grating or prism at 90° to separate out




Figure 7.3. M th order wavelength as a function of $\Delta \sin \theta$.

the orders; the second grating's dispersion limits the allowable slit length. Cross-dispersed gratings are a good match to 2D detector arrays such as CCDs.

7.3.2 Ghosts and Stray Light

If we consider a grating as a frequency mixer, it isn't surprising that irregularities in the fringe spacing get transformed into artifacts in the diffracted spectrum. Small-scale irregularities give rise to a smoothly varying diffuse background of scattered light, which sets the maximum rejection ratio of a spectrometer. Low frequency variations in the grating pitch, caused, for example, by diurnal temperature variation in the ruling engine during a run, produce close-in sidebands on the diffraction spectrum just the way they would in an RF mixer; these close-in artifacts are called *Rowland ghosts*. Ghosts occurring further away are called *Lyman ghosts* and are even more objectionable, since it's much harder to connect them to their true source. As we saw, baffles help the stray light a lot, but ghosts are more of a problem, since they follow the normal light path through the exit slit. Both can be dramatically reduced by using another grating and a third slit, to make a *double monochromator*, and for some very delicate measurements such as Raman spectroscopy, people even use triple monochromators. It's amazing that any light makes it through all that, but it does if you do it properly, and you get double or triple the linear dispersion, too.

7.4 TYPES OF GRATINGS

Classical plane gratings are wonderful at dispersing different wavelengths, but bad at almost everything else—they cost a lot, aberrate converging beams, treat *s* and *p* polarizations differently enough to be a bother but not enough to be useful, require throwing away most of our light to get high spectral resolution, the list goes on and on. Lots of different kinds of gratings have been developed to try to deal with some of these difficulties.

Nearly all gratings sold are replicas, made by casting a thin polymer layer (e.g., epoxy) between the master grating and a glass blank (using a release agent to make sure it sticks to the glass and not the master). Reflection gratings (the usual kind) are then metallized.

Aside: Grating Specifications. Since everything about gratings depends on $\mathbf{k}_{||}$, their properties are usually specified for use in the *Littrow* configuration, where $\mathbf{k}_d = -\mathbf{k}_i$ (i.e., the diffracted light retraces its path). This implies that $k_y = 0$, and that there is no anamorphic magnification of the beam, which simplifies things a lot, but isn't necessarily representative of what you should expect far from Littrow.

7.4.1 Reflection and Transmission Gratings

The essential function of a grating is to apply a very precise spatial phase modulation to an incoming beam. This can be done by reflecting from a corrugated surface, or by transmission through different thicknesses of material. Transmission gratings have the advantages of lenses: compact multielement optical systems, lower sensitivity to flatness errors, and less tendency for the elements to get in one another's way. On the other hand, multiple reflections inside the transmission grating structure give rise to strong artifacts,

making them unsuitable for high resolution measurements in general. With the exception of hologon scanners, you should use reflection gratings for nearly everything.

Reflection gratings are usually supplied with very delicate bare aluminum coatings. Touching a grating surface will round off the corners of the grooves and ruin the grating, and anyone who tries to clean one with lens paper will only do it once. Pellicles can be used to protect gratings from dust. Gold-coated gratings are useful, especially around 800 nm where the efficiency of aluminum is poor.

7.4.2 Ruled Gratings

Ruled gratings have nice triangular grooves, which allow high diffraction efficiency, but since the process of cutting the grooves produces minute hummocks and irregularities in the surface, they also have relatively high scatter. The increasing scatter limits ruled gratings to a practical maximum of 1800 lines/mm.

Ruled gratings can be *blazed* by tipping the cutting point so that the grooves are asymmetrical; by the array theorem of Fourier transforms,[†] a regular grating illuminated with a plane wave will have an angular spectrum equal to the product of the angular spectrum of a single grating line (the envelope) times the line spectrum from the III function. Blazing is a way of making the peak of the envelope coincide with the diffracted order, by tipping each grating line so that the specular reflection from that line is in the diffracted direction. The gratings of Figure 7.4(a) and (b) are blazed.

A grating blazed at λ_B works well from about λ_B to $1.5\lambda_B$, but falls off badly at shorter wavelengths. Like most other grating parameters, the blaze wavelength is quoted for Littrow incidence.

7.4.3 Holographic Gratings

Holographic gratings are frozen interference fringes, made by shining two really, really good quality laser beams on a photoresist-coated substrate. The grating period can be adjusted by changing the angle between the beams, or a transparent substrate can be suspended in an aerial interference pattern with fringe frequency $\Delta\mathbf{k}$, for example, from a laser bouncing off a mirror, and tipped to change the spatial frequency $\Delta\mathbf{k}_{\parallel}$ at the surface. If the resist layer is thin and is developed after exposure (i.e., unexposed resist is washed away), a surface grating results. This grating is then transferred to a durable metal surface by plating it, attaching the plated surface to a stable support (e.g., a glass blank), and then chemically or mechanically stripping the resist, leaving its image in the metal. This metal submaster is then used to make replicas.

The grooves in a holographic grating are ideally sinusoidal in shape (although they can be blazed by ion milling or evaporation of metal from oblique incidence, or by special lithography techniques). The peak-to-peak phase modulation in the reflected wavefront is then roughly

$$\Delta\phi = 2k_z d, \quad (7.6)$$

where d is the peak-to-valley groove height, and shadowing has been neglected.

Best overall efficiency with a sinusoidal modulation is obtained when the specular order goes to 0, which happens with $\Delta\phi = 4.8$ radians (in RF terms, the first Bessel

[†]We'll see it in Example 13.8 and Section 17.4.3 as convolution with a III function.




Figure 7.4. Diffraction gratings: (a) ruled, (b) replicated, and (c) holographic.

null at a modulation index of 2.405—see Section 13.3.7). Deep gratings are therefore used for long wavelengths, and shallower ones for short. Holographic gratings have less small-scale nonuniformity than ruled ones, so they exhibit less scatter and possibly fewer ghosts. The diffraction efficiency of holographic gratings is less strongly peaked than blazed ruled gratings, so they may be a better choice for weird uses.

7.4.4 Concave Gratings

From a diffraction point of view, there's nothing special about a flat surface; since the grooves embody the interference pattern between the incident and diffracted light, we can sample the pattern anywhere we like. Concave gratings combine the collimating, focusing, and diffracting functions in a single element. They are very expensive but are worth it in the deep UV, where mirror coatings are very poor ($R \approx 0.20\text{--}0.35$), so extra bounces cost a lot of photons.

The trade-off is bad aberrations; the focusing mirror is used off-axis, and the diffracted beam fails to obey the law of reflection on which mirror design is based. Concave mirror spectrographs are therefore usually highly astigmatic. This is not as bad as it sounds, since the image of a point in an astigmatic system is a line in the tangential plane (at the sagittal focus) or in the sagittal plane (at the tangential focus) (see Section 9.4.3). Providing the slits are at the tangential focus, the errors caused by astigmatism can be kept small, since the astigmatism's main effect is then to smear out the image along the slit. Of course, in a real situation the tangential image is not a perfectly straight line, and its curvature does degrade the spectral resolution somewhat.

There exist aberration-reduced holographic concave gratings, where the groove spacing is intentionally made nonuniform in such a way as to nearly cancel the leading aberrations over some relatively narrow range of θ_i and λ , which are great if you have the budget for a custom master, or a catalog item happens to fit your needs.

7.4.5 Echelles

An echelle grating (shown in Figure 7.5) is a coarse-ruled grating used in a high order, near grazing incidence; the scattering surface used is the short side of the groove instead




Figure 7.5. Echelle grating.

of the long side (like the risers of the stairs, rather than the treads). It is not unusual for an echelle to be used in the 100th order. Echelles usually have between 30 and 500 lines per millimeter and a really big one (400 mm) can achieve a resolving power of $2W/\lambda \approx 10^6$, a stunning performance. Because of the angular restriction, echelles are usually used near Littrow.

Problems with echelles include expense and severe grating order overlap. The expense is due to the difficulty in maintaining precision while ruling the coarse, deep grooves required, and of course to the low manufacturing volumes; the overlap requires cross-dispersion or a very well understood line spectrum.

7.5 RESOLUTION OF GRATING INSTRUMENTS

7.5.1 Spectral Selectivity and Slits

The usefulness of grating instruments lies in their selectivity—their ability to treat the different wavelengths differently. It's no use dispersing the light at one θ_i if it instantly remixes with other components at different θ_i , different ω , but the same θ_d (take a grating outside on a cloudy day, and try seeing colors). Just how we want to treat each wavelength varies; in a monochromator, we use a spherical mirror to image the angular spectrum on a slit, in order to select only one component, whereas in a pulse compressor, we just add a wavelength-dependent delay before recombining.

7.5.2 Angular Dispersion Factor

In computing the resolution of a grating instrument, we need to know the scale factor between λ and θ_d , the *angular dispersion* D :

$$D = \frac{\partial \beta}{\partial \lambda} = GM \sec \beta = \frac{(\sin \beta - \sin \alpha)}{\lambda \cos \beta}. \quad (7.7)$$

7.5.3 Diffraction Limit

The wavelength resolution of an ideal grating is limited by size, order, and operating wavelength. A uniformly illuminated square grating of side W has a sinc function lineshape as a function of \mathbf{k}_{\parallel} ,[†]

$$E(u, v) = E_{0m} \frac{W}{\lambda} \operatorname{sinc} \left((u - u_i - mG) \frac{W}{\lambda} \right) \operatorname{sinc} \left((v - v_i) \frac{W}{\lambda} \right) \quad (7.8)$$

[†]The xy projection of \mathbf{k} , that is, (k_x, k_y) .

(with $u = k_x/k$ and $v = k_y/k$ as usual, and assuming that \mathbf{k} is along \mathbf{x}). By smearing out the angular spectrum, this effect sets a lower limit to the useful slit width in a spectrometer. As we'll see in great detail in Section 17.6, this sinc function is a nuisance, having slowly decaying sidelobes in its transform. (Don't worry about the sinc functions due to the slits—the exit slit is at an image of the entrance slit, whereas those sidelobes are in the pupil.)

In applications needing the cleanest peak shapes, it is sometimes worth apodizing the incoming light so that its diffraction pattern decays before hitting the edge of the grating. A carefully aligned linear fiber bundle can do a good job of this, provided the aperture of the spectrometer is a bit larger than the fibers'.

The equivalent width of this sinc function is $\Delta \sin \theta_d = \Delta u = \lambda/W$, or in angular terms,

$$\Delta\beta = \frac{\lambda}{W \cos \beta}. \quad (7.9)$$

The theoretical *resolving power* R of a grating is the reciprocal of the diffraction-limited fractional bandwidth:

$$R = \frac{k}{\Delta k} = \frac{\lambda}{\Delta\lambda} = mN = \frac{(\sin \beta - \sin \alpha) W}{\lambda}, \quad (7.10)$$

where N is the number of grating lines illuminated. This gives a useful upper limit on R ,

$$R_{\max} = \frac{2W}{\lambda}, \quad (7.11)$$

with the limit achieved when $\sin \theta_d = -\sin \theta_i = 1$ (i.e., coming in at grazing incidence and coming back along the incident path). The resolving power is a somewhat squishy number, of course, because the definition of resolution is arbitrary, you never use slits as narrow as the diffraction limit, the grating is never perfectly flat nor perfectly uniformly ruled, and the effects of coma are a limiting factor anyway. Typical values of R_{\max} range from 10^4 to as high as 10^6 for very large UV gratings.

7.5.4 Slit-Limited Resolution

We normally don't try to achieve the diffraction-limited resolution, because it requires extremely narrow slits, which are very fiddly and let through very little light. Thus in most cases, it's sensible to use ray optics to discuss spectrometer resolution.

A slit of width w used with a mirror of focal length f produces a beam with an angular range of w/f radians. The grating and the second mirror will reimage the entrance slit on the exit slit, with an angular magnification of $1/M = \cos \theta_i / \cos \theta_d$. The spatial pattern that results is the product of the two slit patterns, so the angular smear is the convolution of that of the two slits, taking account of magnification,

$$A(\beta) = \text{rect}\left(\frac{f(\beta - \beta_0)}{w_{\text{exit}}}\right) \star M \text{rect}\left(\frac{Mf(\beta - \beta_0)}{w_{\text{ent}}}\right), \quad (7.12)$$

which has equivalent width

$$\Delta\theta_{d\text{slit}} = (w_{\text{exit}} + w_{\text{ent}}/M)/f. \quad (7.13)$$

This can be translated into spectral resolution by dividing by $\partial\theta_d/\partial\lambda$,

$$\left. \frac{\Delta\lambda}{\lambda} \right|_{\text{slit}} = \frac{\Delta\beta}{\lambda D} = \frac{(w_{\text{exit}} \cos \beta + w_{\text{ent}} \cos \alpha)}{\sin \beta - \sin \alpha}. \quad (7.14)$$

7.5.5 Étendue

Neglecting diffraction, the étendue $n^2 A\Omega'$ of a grating spectrometer is the product of the entrance slit area wL and the projected solid angle of the grating as seen from the entrance slit, which is approximately W^2/f^2 at normal incidence. The oblique projection of the grating area goes down by $\cos \theta_i$, and the anamorphic magnification $M = \cos \theta_d / \cos \theta_i$ changes the effective size of the exit slit, but those are both effects of order 1, so we'll sweep them under the rug and say

$$A\Omega' = \frac{wLW^2}{f^2}, \quad (7.15)$$

which is tiny; if $f = 250$ mm, a 25 mm grating with a $5 \text{ mm} \times 20 \text{ }\mu\text{m}$ slit ($4\times$ the diffraction limit at $\lambda = 500$ nm, about typical), $n^2 A\Omega = 10^{-5} \text{ cm}^2 \cdot \text{sr}$. Unfortunately, the resolution goes as $1/w$, so we are faced with a 1:1 trade-off of resolution versus photon efficiency for a fixed grating size. The diffraction efficiency of the grating isn't that great, about 0.8 if we're lucky, and we lose another 5% or so at each mirror. Furthermore, it is only at the centre wavelength that the entrance slit is imaged at the exit slit; at other wavelengths it is offset more and more until the overlap goes to 0, so the total efficiency is reduced by almost half. A good ballpark figure is that the efficiency of a monochromator is around 30%.

7.6 FINE POINTS OF GRATINGS

7.6.1 Order Strengths

In general, accurately predicting the strengths of the various diffraction orders requires a vector-field calculation using the true boundary conditions, but there are some general features we can pull out by waving our arms.

Right away, we can see by conservation of energy that there are liable to be sharp changes in the grating efficiency whenever some m enters or leaves the range (7.3), whether by changing θ_i or λ . These are the so-called Wood's anomalies, which are sharp peaks and troughs in the diffraction efficiency curves. With some grating profiles, these anomalies are extremely large and sharp. Figure 7.6 shows the calculated diffraction efficiency of a ruled plane grating with a 9° blaze angle, used in Littrow, which displays strong polarization sensitivity and Wood's anomalies.

7.6.2 Polarization Dependence

The diffraction efficiency, of a grating is usually a strong function of polarization, being highest for light polarized across the grooves (i.e., p -polarized, when $k_y = 0$). This is intuitively reasonable when we consider the behavior of wire antennas—light polarized




Figure 7.6. Theoretical diffraction efficiency in Littrow of a 9° blazed grating. Note the strong Wood’s anomalies and polarization dependence. (Courtesy of The Richardson Grating Laboratory. Note that the Richardson book interchanges p and s.)

along the wires (*s*) causes strong currents to flow in the wires, which therefore reflect it (or absorb, if there’s a load attached). The current is not interrupted by a change in the direction of the surface, because it’s flowing along the grooves; thus there is little light scattered. On the other hand, light polarized across the grooves (*p*) causes currents that have to go over the top of the grooves, leading to strong scatter. By and large, this is how gratings behave, although there are lots of subtleties and exceptions.

7.6.3 Bragg Gratings

The detailed theory of free-space Bragg gratings is fairly involved because of multiple scattering, but weak ones with constant fringe spacing can be treated by coupled-mode theory, since there are only a few orders to worry about.

The main feature of Bragg gratings is that the phase matching condition has to apply throughout the volume, rather than just at one plane, leading to the Bragg condition, which for a sinusoidal grating is

$$\mathbf{k}_d - \mathbf{k}_i = \pm \mathbf{k}_G, \tag{7.16}$$

which is known as the *Bragg condition*.

It’s a very stiff condition, since as Figure 7.7 shows, \mathbf{k}_G is the base of the isosceles triangle made up of \mathbf{k}_i and \mathbf{k}_d ; for a given wavelength and an infinite grating, there’s only a single choice of \mathbf{k}_i that works. This \mathbf{k}_i selectivity is smeared out by the finite depth of the grating, just as the angular spectrum is by the finite width of the grating. Nonetheless, a deep Bragg grating with many fringes is a high-*Q* filter in \mathbf{k} -space.

Bragg gratings can have diffraction efficiencies approaching unity, and in fact since the diffracted wave meets the Bragg condition for being diffracted back into the incident wave, Bragg diffraction is a coupled-modes problem of the sort we’ll encounter in Section 8.3.3.

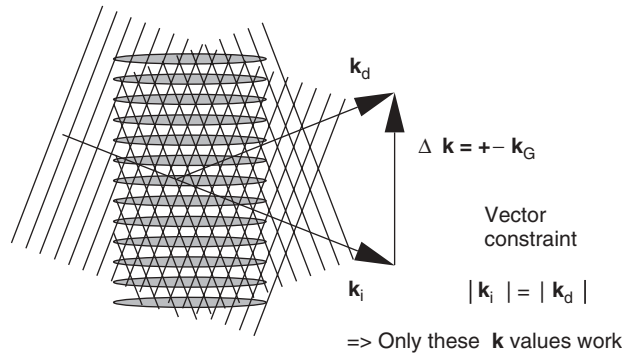


Figure 7.7. Bragg grating.

7.7 HOLOGRAPHIC OPTICAL ELEMENTS

We noted that all gratings are holograms, in the broad sense of embodying the interference pattern of the incident light with the diffracted light. The class of holograms is of course much more general than just diffraction gratings. Fresnel zone plates are basically holograms of lenses, for example, and with the development of computer-generated holograms we can conveniently make more general patterns, for example, holographic null correctors for testing aspheric optics with spherical reference surfaces, and even beam-splitters producing several diffracted beams in different directions with widely differing strengths.

It is somewhat ironic, but a consequence of their very generality, that holographic elements tend to be application specific. You probably won't find an off-the-shelf item that does what you want, so holograms are used mostly in high volume applications such as bar code scanners.

One exception is the holographic diffuser, whose range of scattering angles can range from 1° to 60° . These work poorly with lasers due to speckle but are just the ticket for situations where an ugly illumination function has to be spread out, for example, LEDs, fiber bundles, and liquid light pipes.

Holograms function differently depending on the number of distinct phase or amplitude levels available. Simple holograms have only two levels: zone plates made from annular holes in chrome masks or from a single layer of photoresist. This presents some problems. For one thing, a two-level zone plate functions as both a positive and negative lens, since with only two levels, $\exp(i\mathbf{k} \cdot \mathbf{x})$ is indistinguishable from $\exp(-i\mathbf{k} \cdot \mathbf{x})$. Just as adding bits helps a DAC to mimic an analog signal more accurately, so more levels of phase or amplitude improves holographic optics. The most popular way to implement this idea is *binary optics*: add more levels of lithography, with binary-weighted phase shifts, just as in a DAC.

7.7.1 Combining Dispersing Elements

It is often convenient to combine two dispersive elements in one system, such that their angular dispersions add or subtract. The usual rule is that if the diffraction direction tends to straighten out, the dispersion effects tend to subtract, whereas if the beam is being sent round and round in a circle, the effects add. This of course applies only for elements of




Figure 7.8. Cobbling gratings to add or subtract the diffractive effects, while possibly changing the beam shape anamorphically. Leaving the second grating near grazing multiplies the tuning sensitivity of both diffractions.

the same sort; gratings bend shorter wavelengths less, whereas prisms bend them more. A combination grating and prism could have additive effects even with a nearly straight optical path, like an Amici prism (Section 4.9.3).

Gratings and prisms can be used to change the beam diameter anamorphically (i.e., with magnification different in x and y). The gratings in Figure 7.8 first expand and then contract the beam. The width out of the page is of course unaltered. This anamorphic property is often useful with diode lasers, and Figure 7.8 shows how it can be accomplished with and without tuning sensitivity of the angular deflection.

7.8 RETROREFLECTIVE MATERIALS

We're all familiar with retroreflecting materials, used for bicycle reflectors, traffic signs, and safety clothing. They have lots of possibilities in instrument design too, and so everyone should know something about them. The basic idea is to use large numbers of small, poor quality retroreflectors, so as to return incident light generally back toward its source, with an angular spread of 0.5° to 5° in half-angle and (unlike larger corner cubes) no significant lateral shift.

There are two main kinds: glass beads and corner cube arrays embossed in plastic (see Section 4.9.8). A sphere in air makes a pretty good retroreflector if its index is chosen so that incident light is focused on the opposite surface of the sphere. A ray at height h , parallel to the axis, has an angle of incidence $\sin \theta_i = h/R$, and to focus on the back surface, $\sin \theta_d = h/(2R)$, so by Snell's law,

$$\frac{n_2}{n_1} = \frac{\sin \theta_i}{\sin \theta_r} = \frac{h/R}{h/(2R)} = 2, \quad (7.17)$$

which can just about be done in glass. The angular spread can be adjusted via defocus, by varying n . To prevent most of the light getting out the other side, aluminum-coated beads are used, with the aluminum etched off the top surface after the coating is applied. The embossed plastic stuff has a narrower acceptance angle than the spheres, because away from the symmetry axis of the corner, more and more of the light fails to make the three TIR bounces required for retroreflection. It's also a bit of a puzzle to mount, because you can't touch the TIR surface or you'll deactivate all the little cubes. The

standard solution is to use a quilted foam backing that touches only a few percent of the area, which makes the retroreflection spatially nonuniform.

The figure of merit for a retroreflective material is its *specific brightness*, which is measured in inverse steradians, although it's quoted as lux per steradian per lux or other cute units that amount to the same thing: if you put in 1 W/m^2 , how much flux is radiated per steradian at specified angular offsets from exact back-reflection? For a retroreflector with an RMS beam spread half-angle of $\Delta\theta$ and photon efficiency η , the specific brightness is

$$B \approx \frac{\eta}{\pi(\Delta\theta)^2}. \quad (7.18)$$

For a $\Delta\theta$ of 0.5° , this is around 4000—a factor of 13,000 over a Lambertian ($\Omega' = \pi$) reflector, assuming that $\eta = 1$. The best real materials are more like 900 or 1000, a factor of 3000 or so over Lambertian. This is an enormous effect—a strip of tape in the right place can get you a 70 dB (electrical) signal level increase, which is well worth the labor and the compromises it takes to use these materials (they're very speckly when used with lasers, for example).

The stuff made for use in signs and clothing isn't that great; the specific brightness is usually below 300, and 60 is not uncommon, but those numbers still represent big signal increases in systems where most of the illumination is not collected. There does exist material that gets into the 10^3 range, but it isn't easy to find. You can also buy just the beads,[†] made usually from barium titanate glass with an index of 1.9 or a bit higher. They're used for spraying onto traffic paint before it dries and may be helpful for special situations where it's inconvenient to use the made-up sheet material. Other considerations include rain—glass bead retroreflector relies on refraction at the air–glass boundary, and so doesn't work well when it's wet. Assuming the TIR surfaces remain dry, the corner cube stuff is almost unaffected by rain.

The best retroreflective materials in the 3M catalog for instrument use are Scotchlite 2000X Very High Gain Sheeting (corner cubes) and Scotchlite 7610 High Gain Reflective Sheeting (spheres). Both are available in tape rolls and can really make a difference to your SNR, especially in a fiber-coupled instrument. The other main manufacturer of this stuff, Reflexite, also has cube material tailored to produce a nearly fixed offset angle (not 180°).

TIR film will mess up the polarization of your beam, as in Section 4.9.8. Reflexite makes metallized corner cube material, which reduces this problem. Plastic retroreflector is naturally of no use in the UV, but the high brightness glass bead stuff is quite good, as it has no plastic overcoat on top of the spheres.

7.9 SCANNERS

Scanning systems are one of the thorniest areas of electro-optical instrument building. The whole point of scanning is to interrogate a huge $A\Omega$ volume sequentially with a low- $A\Omega$ system. None of the available methods is as good as we'd like, and the cost of commercial solutions is enough to make you gasp (how about \$4000 for a 25 mm, two-axis galvo scanner with analog driver cards and no power supply?).

[†]Suppliers include Potters Industries and Cataphote.

The difficult design issues, together with the very high cost of playing it safe, make scanner design worth a considerable amount of attention. The main points to worry about are scanner and encoder accuracy, rectilinearity, range, speed, jitter, aberration buildup, field flatness, temperature drift, and power consumption (other than that, it's easy). We'll start with mechanical scanners.

Before we begin, it is useful to have a bit of vocabulary to describe different scanner vices. Nonrepeatable motions, caused, for example, by out-of-round bearings, are called *jitter* if they're in the scan direction and *wobble* otherwise. Repeatable errors are *scan nonlinearity* along the scan direction, and *scan curvature* out of it. Temperature drift, which is sometimes very serious, comprises offset (zero) drift and gain drift. Scanning is often combined with signal averaging to reject noise and spurious signals; see Section 10.9.2.

7.9.1 Galvos

Galvanometer scanners are electric motors that don't turn very far ($\pm 45^\circ$, maximum) but can accelerate very fast and move very accurately. They usually run on ball bearings, but some use flexure bearings. Single- and double-axis galvos are available that provide 12 bit angular accuracy with settling times of several milliseconds. Small ones are much faster than large ones, because the moment of inertia I goes as mr^2 , which tends to grow as r^5 . The angular accuracy of the encoders isn't all that great over temperature, typically 10 arc seconds drift and 0.05% gain/ $^\circ\text{C}$, with some being much worse. If you really need that 12 bits, you have to compensate for those in some way. Jitter typically runs 10–20 arc seconds, and wobble more like 5–10. Those are really what limit the resolution. The torque from any motor tends to go as the volume—due to field and current limitations, the available force per unit area is approximately constant, so the torque goes as the surface area times the diameter. The moment of inertia grows faster than that, so big galvos tend to be slow.

Resonant galvos shrink the mirror and the motor, and add a big torsion spring to get their restoring force, which enormously reduces their moment of inertia. This makes them quite a bit faster (500 Hz), but being high- Q resonant systems, they cannot be controlled within a cycle; only the amplitude and phase of their sinusoidal oscillation can be changed, and many cycles are required for settling afterwards. Thus resonant galvos are good fast-scan devices, where they compete with rotating polygons and hologons; the trade-off is adjustable angular range and sinusoidal scan versus uniform scan speed through a fixed angular range.

7.9.2 Rotating Scanners

All reciprocating scanners have to slow down and turn around at the end of their travel, which makes them relatively slow. What's worse, their varying scan speed makes it relatively difficult to take data points equally spaced in angle—it requires a continuously varying clock frequency. This can be implemented with a lookup table in hardware, or done by resampling the data afterwards (see Section 17.8). Since the dwell time on each pixel is different, more lookup tables may be needed to take out the resulting gain error and offset. All these lookups, whose contents depend on unit-to-unit differences such as the details of loop damping and rotor imbalance, require an onerous calibration procedure for high accuracy applications.

Nonconstant scan speed is particularly obnoxious when you're using a continuous frame scan, since it leads to hooks at the ends of the scan line. Large amounts of overscan are necessary to combat it. A scanner with good acceleration can make sharper turns, so less overscan is needed with high torque motors, low moment of inertia (i.e., small mirrors), and slower scan rates. This is all manageable, but still a constant, fast-scan speed (constant in m/s or rad/s depending on the application) would be useful for raster applications.

One partial solution is a continuously rotating scanner, such as a polygon mirror or holographic scanner.

7.9.3 Polygon Scanners

A polygon scanner is a spinning wheel with flat mirror facets on its periphery (Figure 7.9). These may be oriented normal to the radius vector, so that the wheel is a polygonal cylinder, or angled like a cone. In order to get a straight-line scan with a cylindrical polygon, the beam has to come in normal to the rotation axis, although other configurations exist with tilted facets.

With the beam so aligned, rotating a mirror through $\theta/2$ deviates the reflected light by θ , so an n -facet polygon deflects light through an angular range $\Delta\theta$ of

$$\Delta\theta = \frac{4\pi}{n}, \quad (7.19)$$

although you can't use all that range, since at some point the edge of the facet has to cross your beam, leading to a dead time on each facet. A polygon rotating at constant speed naturally produces a constant angular velocity (CAV) scan, which is great for some things (e.g., lidar) but a problem for others (e.g., document scanning, where a constant linear velocity is much more convenient). Polygons can go very fast; speeds over 50,000 rpm can be achieved with a small solid beryllium scanner in a partial vacuum, though not easily or cheaply. The ultimate limit is set by deformation and then failure of the polygon itself. High end printers (\$1M) use polygons with 10–12 facets running at up to 40,000 rpm on air bearings (with 10 beams, that's a 70 kHz line rate), but in more pedestrian applications, keep below 7000 rpm, and remember that below 3000




Figure 7.9. Polygon scanner.

things get much easier. Polygons cause no aberration of the scanned beam. They have a unidirectional scan pattern, with a retrace interval as the edge between two facets crosses the beam. Well-made polygons can have a facet-to-facet angular accuracy of a few arc seconds, which is good enough for most applications. Cheap polygons are closer to an arc minute, but cost only a few dollars.

7.9.4 Polygon Drawbacks

Polygons have a number of drawbacks. Their moment of inertia is high, so that the scan rate cannot be changed rapidly. We usually use them to get high speed, so that the kinetic energy is high, which compounds the speed problem and adds wind resistance, turbulence, and noise, ultimately forcing us to use a vacuum chamber.

The constant angular velocity of the beam means that it scans a flat surface at a nonuniform rate, unless we do something about it. A subtler difficulty is that since the rotation axis does not pass through the surface of the mirror, as it does with a galvanometer, the scanned beam translates as well as pivoting during a line. Thus a polygon-scanned system lacks a true pupil. You can get (very expensive) f - θ lenses, which have just the right amount of built-in barrel distortion to convert a constant angular velocity scanned beam into a constant linear velocity spot, and a flat field; they're big chunks of glass, used at very low aperture (e.g., a "250 mm $f/16$ " lens whose front element is 90 mm across). It is somewhat galling to have to use a \$1200 lens with a \$25 polygon.

The remaining trouble is their very high sensitivity to shaft wobble. A polygon accurate to a few arc seconds is very much easier to get than a motor whose shaft is straight and stable to that accuracy, especially at the end of its life. Air bearings are a good choice for high speed polygons, but they don't usually work as well at low speed.

7.9.5 Butterfly Scanners

In Section 4.9.4, we encountered the pentaprism, which is a constant deviation 90° prism that works by having two reflections; tipping the prism increases one angle while decreasing the other, making the total constant. The same principle can be applied to scanning, resulting in the butterfly scanner of Figure 7.10, which is a nearly complete solution to the shaft-wobble problem; drawbacks are complexity, expense, probably worse fixed facet-to-facet errors, and much higher air resistance, noise, and turbulence.

7.9.6 Correcting Rasters

Once the shaft wobble has been corrected, the scan is still not perfect. To get really good rasters from any multisegment device, you really have to have software that knows which segment you're on, and dials in the appropriate offsets. While this requires calibration, it's not a difficult problem since it's only line-to-line variation and can be done in the instrument itself using a vertical bar test pattern. Furthermore, the dimensional stability of the hologon or polygon means that it can be done once and forgotten about.

In-plane errors cause only timing trouble, which can be eliminated completely. Out-of-plane errors are more obnoxious visually, causing obvious nonuniformity in raster line spacing, and are more difficult to handle, requiring an additional deflection element such as a Bragg cell.




Figure 7.10. The butterfly scanner works on the pentaprism principle.

7.9.7 Descanning

In order to reduce the required $n^2 A \Omega'$ of the detection system (which makes it smaller, faster, cheaper, and dramatically more resistant to background light), we usually need to *descan* the received light. There's an easy and a hard way to do this.

The easy way is to work in backscatter, and use the same scanner on the transmit and receive sides of the instrument, for example, to interrogate a sample one point at a time. If the scanner is at the pupil of the objective lens, light backscattered from the sample will be recollimated, so that it comes back antiparallel to the transmit beam. The mirror won't have had time to move far during the time of flight, so by symmetry, both the scan angle and any angular jitter are removed, leaving only a bit of Doppler shift.

The hard way is to use a second scanner, synchronized to the first. You have to do this sometimes, for example, in a long path sensor where you're sweeping a focused beam back and forth without steering it, and can't work in backscatter for some reason. This really is a miserable way to make a living, so work hard to avoid it; a corner cube or some retroreflecting tape will often let you off this hook.

Aside: Preobjective and Postobjective Scanning. Before scanning, our beam is usually collimated, so its NA is very low. It seems a bit perverse to take something like that, which we can focus very well with inexpensive lenses, scan it through perhaps 45° , and then try to focus it with a very expensive f - θ lens. Couldn't we focus first and scan afterwards?

If the NA is low enough and the working distance long enough, this is a good possibility. The major drawback is that the field flatness and nonuniform scan speed (in m/s on the sample surface) are uncorrected unless you do something fancy yourself. This may not matter in your application, in which case this *postobjective* scan strategy will work well. Just don't try it with hologons (see Section 7.9.9). A hybrid scheme, where the line scan is preobjective and the frame scan is postobjective, is also widely used.

7.9.8 Constant Linear Scan Speed

It isn't that easy to achieve constant scan speed with a mechanical scanner. Rotating a mirror at a constant speed $\dot{\theta}/2$ produces a reflected beam whose angular speed $\dot{\theta}$ is

constant; if the scanner is a distance h above a planar surface, the scan position x on the surface is

$$x = h \tan \theta, \quad (7.20)$$

which is stretched out at large angles, just like late-afternoon shadows. Reciprocating scanners such as galvanometers and voice coils slow down and stop at the ends of their angular travel, so they tend to scan more slowly at the edges; if the beam scans sinusoidally through $\pm\theta_{\text{pk}}$ at radian frequency ω , then the scan speed v is

$$v(\theta) = h\omega\sqrt{\theta_{\text{pk}}^2 - \theta^2} \sec^2 \theta \quad (\dot{\theta} > 0). \quad (7.21)$$

The slowdown of $\dot{\theta}$ at large θ approximately compensates the stretching out of x , so that the resulting curves look like the Chebyshev filters of Section 15.8.3; choosing $\theta_{\text{pk}} = 40.6^\circ$ produces a maximally flat response. Table 7.1 shows optimal scan parameters for an equiripple error from $\pm 0.01\%$ to $\pm 5\%$: tolerance, linear range θ_L , peak range θ_{pk} , duty cycle (or scan efficiency), and the corresponding error with a CAV scan of $\pm\theta_L$. Note how the duty cycle improves as the error band is relaxed, and how much lower the maximum error is for the galvo versus the polygon, at least when used unaided. Usually we have to accept more scan speed variation and compensate for it with slightly nonsinusoidal motion (easiest), nonuniform pixel clock speed, resampling digitally afterwards, or (as an expensive last resort) an f - θ lens.

If we need to focus the beam on the surface, it's usually enough to put the galvo at the pupil of a flat field lens, with due attention to the lens's geometric distortion.

7.9.9 Hologons

A holographic scanner consists of a spinning surface containing holographic elements (Figure 7.11). The most common type is the *hologon*, short for holographic polygon. A hologon is a spinning glass disc containing radially segmented transmission gratings (like orange segments), with the grating lines oriented tangentially.

Hologon scanners are best operated near minimum deflection, that is, when the incoming and outgoing beam make equal angles with the surface normal. Small amounts of wobble in the shaft then cause only second-order angular deviations in the frame direction, which is an important advantage of holographic scanners over simple polygon mirrors, though butterfly scanners do even better. Beiser shows that for a scanner producing 90°

TABLE 7.1. Approximating a Constant Linear Velocity Scan with a Sinusoidal Galvo

Speed Tolerance ($\pm\%$)	θ_L (\pm°)	θ_{pk} (\pm°)	Duty Cycle (%)	Constant AV Error ($\pm\%$)
0.01	10.25	41.03	16.1	1.6
0.05	15.0	41.68	23.4	3.5
0.1	17.7	42.17	27.6	4.9
0.5	26.2	44.2	40.5	10.8
1	31.0	45.7	47.5	15
5	52.8	56.8	76	46




Figure 7.11. A hologon scanner is the diffractive analogue of a polygon. It isn't as efficient but has much lower jitter, weight, and wind resistance.

deviation ($\theta_i = \theta_o = 45^\circ$), a large shaft wobble of 0.1° (360 arc sec, or 1.75 mrad) produces a scan wobble in the frame direction of only 1.3 arc sec, an improvement of nearly 600:1 over a polygon, with even better performance for smaller wobbles.

The scan line produced by a hologon scanner is in general curved, because k_y can't always be 0 when we're rotating the grating in azimuth. By choosing $\theta_i = \theta_d = 45^\circ$, the scan can be made almost perfectly straight, which is how they are usually used. The deflection is easily found by applying phase matching; if the grating is rotated through an angle ϕ_{shaft} from the center of the facet, the change in k_x of the diffracted beam is equal to that of the grating, so in the middle of the scan line,

$$\frac{\partial \theta_{\text{az}}}{\partial \theta_{\text{shaft}}} \approx \frac{k_G}{k}, \quad (7.22)$$

which is equal to $\sqrt{2}$ for the 45° – 45° scanner. The angular scan speed $\dot{\theta}$ is also mildly nonuniform, being slightly compressed at the ends of the travel (we saw that this is actually an advantage in scanning flat surfaces). The effect is nowhere near as strong as with galvos.

As a practical matter, hologons are restricted to collimated beams. A focused beam used with a collimated-beam hologon produces an astounding amount of astigmatism—dozens of waves over a 45° scan angle, even with an NA of only 0.01. Since they are holograms, it is possible to make a scanner that focuses as well as deflects the light. It might seem that the resultant saving of a lens would be very valuable, but in fact doing this throws away the major advantage of hologons, the insensitivity to shaft wobble. Resist the temptation to be too fancy here, unless your performance specs are modest (e.g., in hand-held bar code scanners). One possible exception would be adding a few waves of optical power to correct for aberrations in a simplified scan lens, because the wobble effect would then still be small. The angular accuracy of a hologon can be as good as 10 arc seconds facet to facet, although 30 is easier and hence cheaper.

If the facets are made with slightly different values of G they will deflect the beam at slightly different angles, so that an N -facet hologon by itself can perform an N line raster scan, which allows a useful trade-off between scan speed and alignment accuracy. (Doing this with a polygon would make it dynamically unbalanced.)

The diffraction efficiency of hologons is normally quite good—80–90%, but that isn't as good as a properly chosen mirror coating, so you'll pay a decibel or two in detected signal for using a hologon.

7.9.10 Fast and Cheap Scanners

If your scan range and accuracy requirements are modest, don't forget the obvious candidates, for example, mounting a laser diode on a piezoelectric translator or a bimorph, and letting the collimating lens do the work. Life doesn't get much easier than that.

7.9.11 Dispersive Scanning

It is also possible to scan over a small range by using a tunable source (e.g., a diode laser) with a dispersive element, such as the second compound grating device in Figure 7.8. This is a very good technique for some purposes, because it is extremely fast (~ 20 resolvable spots in 3 ns), and although you do have to avoid mode jumps and cope with power variations, it presents few problems otherwise.

7.9.12 Raster Scanning

Raster scanning requires a 2D scanner or two 1D scanners in cascade. You can get two-axis voice coil type scanners, which tip a single mirror about two axes; they behave a bit like galvos but have only a few degrees' scan range and aren't as stable or repeatable, because the mirror mount usually relies on a single wire in the middle for its location, and the orthogonality of the tilts is not well controlled.

If we need to use two scanners, we must either accept a pupil that moves around a good deal (several centimeters with most galvos), or use a relay lens to image the center of one scanner on the center of the other. The usual approach is to use a small fast scanner first, to do the line scan, and a large, slow one afterwards for the frame scan, although the relay lens approach allows both scanners to be the same size. The moving pupil problem is worst with a pure preobjective scan approach, but if you can put the scan lens between the line and frame scanners, it gets a lot easier; in the ideal case of pure postobjective scanning, you can use a simple lens for focusing, with perhaps a weak toric field-flattening lens to correct for the different object distances at different scan positions.

7.9.13 Mechanical Scanning

Another approach is to keep the optical system very still and move the sample, as is done in some confocal microscopes and step-and-repeat photolithography tools. This is slow and prone to mechanical jitter, due to the requirement to start and stop the motion of a large mass quickly, and to the instabilities of translation stages. On the other hand, your point-spread function is really constant with position, and there is no limit on the number of resolvable spots. Another mechanical scanning method is to rotate or translate the entire optical system, as in a periscope or an astronomical telescope, which scans slowly to correct for the Earth's rotation.

7.10 MODULATORS

Diode lasers are unique among optical sources in being highly directional and easily modulated at high speed. Unfortunately, most sources of interest are not like that, so we need external modulators. Under this rubric lie a fairly motley collection of out-of-the-way physical effects, all of which have serious drawbacks, not all widely known. Modulators in general are troublesome devices if you need nice pure beams with uniform polarization, no etalon fringes, and smooth amplitude profiles.

Most modulators are based on bilinear interactions[†] between two fields in some material, for example, the electro-optic effect, where the applied electrostatic field causes the dielectric tensor to change, or the acousto-optic effect, where the optical wave diffracts from a sinusoidal phase grating produced by the traveling acoustic wave.

7.10.1 Pockels and Kerr Cells

Optical modulators based on the linear (Pockels) or quadratic (Kerr[‡]) electro-optic effects are shown in Figure 7.12. These are among the fastest modulators of all, but they are a genuine pain to use. Think of them as voltage-controlled wave plates.




Figure 7.12. Pockels and Kerr cells.

[†]A bilinear interaction is one that is linear in each of two independent variables; that is, it is expressible as $f(x, u) = g(x)h(u)$. An electronically variable attenuator is an example of a bilinear device if its gain is linear in its control voltage.

[‡]The electro-optic Kerr effect is not the same as the magneto-optic Kerr effect, which leads to polarization rotation in linearly polarized light reflected from a magnetized medium.

Kerr cells are based on a quadratic electro-optic effect in isotropic materials (usually nasty inflammable organic liquids such as carbon disulfide or nitrobenzene). Their quadratic characteristic makes them harder to use in applications needing linearity, of course, but since the static birefringence is 0, they are pretty predictable. Kerr cells are excited transversely by dunking capacitor plates into the liquid cell. They are normally used singly, with bias voltages around 10 kV. The organic liquids are strongly absorbing in the UV, so Kerr cells are generally limited to the visible and near IR. Getting decent uniformity requires limiting the fringing fields, which (as we'll see in Section 16.2.5) means making the plate dimensions several times their separation.

The variable retardation of electro-optic cells can be used to modulate the polarization and phase of a beam, as shown in Figure 7.13. Pockels cells are built from crystals such as potassium dihydrogen phosphate (KDP) or lithium niobate, nasty birefringent things whose dielectric tensor $\underline{\epsilon}$ depends on the applied \mathbf{E} . The dependence of $\underline{\epsilon}$ is generally complicated; a given applied \mathbf{E} can change all the coefficients of $\underline{\epsilon}$. Since the material is already anisotropic, the leading-order change is linear in applied field.

Pockels and Kerr cells are usually used as amplitude modulators, by following the polarization modulator with an analyzer. They can also be used as phase modulators, by aligning the polarization of the beam with one of the crystal axes, so that the polarization remains constant but n varies. It's hard to get this really right in a multielement cell,




Figure 7.13. E-O modulators: (a) polarization, (b) phase, and (c) amplitude.

because the optic axes of the elements may not be aligned well enough. Fancier things can be done, for example, frequency shifters made by putting a rotating \mathbf{E} field on a crystal between crossed circular polarizers, but they tend to be rare compared with phase, amplitude, and polarization modulation applications.

There are two main designs for Pockels cells. Because the applied \mathbf{E} changes more than one element of the dielectric tensor, the field can be applied longitudinally (parallel to \mathbf{k}) or transversely. Longitudinal cells usually have a ring electrode around the outside of the face, which produces some polarization nonuniformity that limits their ultimate extinction ratios to a few hundred, even with long thin cells. Transparent electrodes such as indium–tin oxide are also used, but they’re pretty lossy, and not conductive enough for the highest speed applications; for really fast stuff, the champion electrode is a low pressure neon plasma, which improves the extinction to a few thousand, even while improving the étendue.[†]

Transverse cells are simply metallized. The trade-off between the two styles is in étendue and capacitance versus voltage; a longitudinal Pockels cell has a reasonably large étendue (especially the ITO and neon plasma types) but requires high voltage, whereas a transverse one has high capacitance and (because it is long and narrow) very low étendue.

Since the effect is linear, the number of waves of birefringence in a longitudinal cell is proportional to the line integral of $\mathbf{E} \cdot d\mathbf{s}$ along the ray path, that is, to the voltage drop across the crystal. Since going from off to fully on requires a change in retardation of one-half wave (why?), the figure of merit for a given electro-optic material is the *half-wave voltage* V_π , which—most inconveniently—is usually several thousand volts.[‡] Because both polarizations get phase shifted in the same direction, the retardation is less than the phase modulation, so phase modulators can work at somewhat lower voltages.

The Pockels effect is fast; a properly designed transverse cell such as a 40 Gb/s telecom modulator can switch in 10 ps, and the intrinsic speed is even higher. The problem in bulk optics is that with a longitudinal cell whose half-wave voltage is (say) 4 kV, to get a 1 ns edge we have to make the bias voltage change at a rate of 4,000,000 V/ μ s, which is an interesting challenge—if the 50 Ω connecting cable is 10 cm long, it’ll take 80 A for 1 ns just to charge the cable. You can do that with a spark gap or a series string of avalanche transistors, but only just; the usual method is a big thyatron producing a 10 ns edge, running into a ferrite-loaded coaxial pulse forming network.[§] All these are limited to shuttering applications since you can’t stop an avalanche once it starts. (These techniques are discussed in Section 15.14.1.) Accordingly, people have worked out ways to ease the voltage requirement. The main way is to take many thin plates of electro-optic material dunked in index oil and drive them in parallel as in an interdigitated capacitor, as shown in Figure 7.12. You can get these down into sub-400 V territory, which is a lot easier although still not trivial.

The optical path in a Pockels cell contains a great big chunk of birefringent material, so it has a huge static retardation (many waves), and thus has a few percent nonuniformity

[†]Pockels cell people enjoy suffering almost as much as femtosecond dye laser people used to.

[‡]This is an example of the extreme linearity of optical materials—even if we get to pick an arbitrarily horrible material, and orient it for maximum effect, we still need thousands of volts to make an order-unity difference in the transmitted field.

[§]Saturation in the ferrite causes the back end of the pulse to move faster than the front end, which leads to shock formation, like the breaking of an ocean wave on a beach. You can get below 100 ps rise time for a 20 kV pulse in 50 Ω , but you have to really want to, and the EMI problems are, *ahem*, interesting.

of retardation across its face, fairly poor wavefront fidelity, and a significant amount of temperature drift. Many of the crystals used for Pockels cells are piezoelectric, and so they may exhibit low frequency resonances; those made with biaxial crystals have particularly nasty temperature dependence, since (unlike uniaxial crystals) the optic axes can move around with temperature. For a device with as much retardation as a Pockels cell, this can be a serious drawback. Low voltage devices have lots of etalon fringes too. For high accuracy applications, longitudinal Pockels cells need a lot of babysitting, and that confines them pretty much to lab applications.

7.10.2 House-Trained Pockels Cells: Resonant and Longitudinal

Many applications of modulators are relatively narrowband, so that we can stick the cell in an electrical resonator to reduce the required voltage by a factor of Q . Cells like that are available from 100 kHz up past 30 GHz, with operating voltages of 6–30 V.

Transverse cells tend to be long and thin because we win by a factor of L/d , which allows operating voltages down to the tens-of-volts range, a much better match to ordinary circuitry. This leads to higher capacitance, but since the electrical power goes as CV^2 , we win anyway, and traveling-wave structures can be used to make the device look resistive when that becomes a problem.[†] The most serious limitation is their very small étendue. Even for light going right down the axis, the beam often has to be so small in diameter that diffraction limits the length of the cell. This is an especially serious limitation in the infrared, where diffraction is worse and more retardation is required to get a half-wave shift. Transverse modulators are most commonly found in integrated optics and fiber-coupled applications, where they are entirely practical; a single-mode waveguide made of electro-optic material needs very little étendue, the field confinement of the waveguide avoids any length limitation due to diffraction, and nonuniformity causes fewer problems since only one mode is involved. The really fast traveling-wave integrated-optic Pockels cells used for telecom usually need about 100–200 mW of RF power and have rise times as short as 12 ps or so. Telecom modulators are usually *zero chirp*, that is, they produce little or no phase modulation, which otherwise shows up as spurious FM sidebands. Chirp is one of the main limitations of directly modulated diode lasers, so this matters. If you really need fast beam modulation, consider using one of these and expanding the beam later.

7.10.3 Liquid Crystal

Another class of electro-optic devices is based on liquid crystals (LCs). These odd materials are liquids made of large molecules, which show some large scale orientation effects even in the liquid state. The physics of liquid crystals is very rich (read complicated). A very slightly grooved surface (e.g., glass that has been wiped in one direction with a cloth pad) can determine the orientation for fixed applications such as wave plates; an applied voltage can change their alignment, which changes the resulting birefringence. Because they rely on the physical motion of molecules, rather than electrons, all liquid crystal modulators are slow (1 μ s to 1 ms). You use them like big, slow, low voltage Pockels cells, to modulate polarization, phase, or amplitude.

[†]Think of coaxial cable, which is 100 pF/m, but can handle gigahertz signals over many meters because of its traveling-wave character.

They come in two basic types: the extremely slow, continuously variable *nematic* ones, and the somewhat-faster, binary *ferroelectric* ones. One of the best things about LC devices is their huge étendue; you can get 100 mm diameters with $\Omega \approx 0.5$ sr. They are also nearly indestructible—their damage thresholds are so high they're not easy to measure.[†] Being liquids, they make intimate contact with the electrodes; because their birefringence is so high, they can be very thin. This makes it easy to build spatially addressable LC *spatial light modulators* (SLMs). Besides the familiar LCD displays, SLMs are used to make shutters, masks, and low resolution computer-generated holograms, among other things.

Example 7.2: Phase Flopping Interferometers. One especially good use of LC modulators is in making zero-background imaging measurements by inverting the phase of the signal but not the background, and frame subtracting. For example, many years ago a colleague of the author's, T. G. Van Kessel, built a very successful Nomarski interference system for measuring the latent image in photoresist. (The image is latent between exposure and development.) It was a normal Nomarski-type metallurgical microscope (see Example 10.2) with the addition of an early liquid crystal variable wave plate before the analyzer, oriented parallel to the Nomarski axis (45° to the analyzer axis). Changing the retardation from 0 to $\lambda/2$ on alternate video frames caused a π relative phase shift in the combined beams; this inverted the Nomarski contrast but preserved the background amplitude. Under frame subtraction, the weak phase contrast signals added and the strong background canceled out, making an electronically zero background measurement (see Section 10.8).

7.10.4 Acousto-optic Cells

The most common Bragg grating in bulk optics is the acousto-optic Bragg cell. We encounter the piezo-optic effect in Section 8.5.6, where it causes stress birefringence in lenses and prisms. Launching a strong acoustic plane wave in a material with a big piezo-optic coefficient makes a moving Bragg grating. Typical frequencies are 40 MHz to 2 GHz, which produce acoustic wavelengths of 2–100 μm . If the interaction zone is too skinny, phase matching perpendicular to \mathbf{k}_A is no longer a constraint, so we get many weak diffraction orders, spaced at multiples of \mathbf{k}_A . This is the *Raman–Nath* regime, shown in Figure 7.14.

That grating has some unique properties: the diffracted light gets frequency-shifted by $\pm f_A$ depending on which direction it was diffracted. Also, the diffraction efficiency can be smoothly varied from 0% to 80% or so merely by changing the RF power from 0 to a watt or two (and more than that for some materials, e.g., glass).

The phase matching condition can be modified (and sometimes considerably relaxed) by using a birefringent material. By a suitable choice of cut, the change in k_{diff} with incidence angle or grating period can be compensated by that due to the change of n . This trick is used all the time in acousto-optic devices.

Acoustic waves in solids are tensor waves, which include scalar (longitudinal) and vector (transverse) waves, but more general shear waves can propagate too. Predicting the effects of a given order and type of diffraction can be done by classical field theory, but it is far easier and less blunder-prone to take a simplistic quantum view. We know

[†]That doesn't apply to the film polarizers on LC shutters, however.




Figure 7.14. Acousto-optic cells: Raman–Nath and Bragg regimes.

that a photon has energy, momentum, and angular momentum; well, so do phonons, and they are all conserved during the interaction, on a quantum-by-quantum basis. A photon that absorbs a phonon (the $+$ or *anti-Stokes* branch) gets its frequency upshifted ($E = \hbar\omega$), and is bent along the acoustic propagation direction ($\mathbf{p} = \hbar\mathbf{k}$)—the energies and momenta add. If instead it emits one (by stimulated emission, the $-$ or *Stokes* branch), it's downshifted and bent away from $\mathbf{k}_{\text{acoustic}}$. Similarly, conservation of angular momentum means that a linearly polarized photon that absorbs or emits a shear phonon has its polarization shifted—it goes from s to p or p to s . A second-order diffraction gets twice the shift in each, because it involves emitting or absorbing two phonons, and so on.

Acousto-optic cells are usually used with laser beams, because their aperture is so small; the major use is as medium-speed amplitude modulators (DC to 10 MHz easily, DC to 100 MHz if you really work at it—focused beams, fast materials, small crystals). One exception is the acousto-optic tunable filter (AOTF), which achieves a wide field angle (and hence a decent étendue) by *noncritical phase matching*, where the curve of λ versus the phase-matched θ_i has a maximum, so the phase matching error is quadratic in $\Delta\theta_i$. These are available in both collinear and noncollinear designs. Narrowness of the passband requires more waves of interaction zone, as we saw, so the angular acceptance goes down as the selectivity goes up; a typical cell with $\Delta\nu/\nu = 0.2\%$ works over $\pm 5^\circ$ ($\Omega = 0.023$ sr), a much wider range than a grating instrument with the same resolution.

You can image through AOTFs, but it doesn't work very well unless you're careful. The images are corrupted by ghosts and prism aberrations, and the ultimate spectral selectivity is limited by a high background light level. Both of these problems are caused by sinc function sidelobes due to a finite interaction length and the random phase matching of all the light and acoustic energy bouncing around inside the crystal. Putting two in a row is an effective but pricey solution to the ghost problem, and putting the AOTF at an image instead of a pupil pretty well solves the aberration problem too.[†]

7.10.5 AO Deflectors

The same crystal-cutting tricks allow a reasonable range of output angles (± 2 – 3°) for a 1-octave frequency shift, making a narrow-range but very fast scanner with high

[†]Dennis R. Suhre et al., Telecentric confocal optics for aberration correction of acousto-optical tunable filters. *Appl. Optics* **43**(6), 1255–1260 (February 20, 2004).

diffraction efficiency, the acousto-optic deflector (AOD). In this case, we want θ_i to be constant over a wide range of acoustic frequency, a condition called *tangential phase matching* that can be met in TeO₂.

There are interesting tricks you can do with AODs. If you gently focus a beam at the center of an AOD, it will pivot about its focus with f . Recollimating produces a beam that moves from side to side without steering, which is very useful for long path sensors such as the extinction system of Section 10.8.3. Galvos work for this too, of course, but since an AOD has no moving parts, you can get a much higher scan frequency, with zero jitter and wobble.

Getting high resolution out of an AOD requires lots of fringes, just like any other Bragg grating; the number of resolvable spots is equal to the transit time–bandwidth product. Like Rayleigh and Sparrow resolution criteria, there's a factor of order 1 in front that we can argue about, depending on your beam profile and how much overlap you allow between distinct spots.

Aside: Acoustic Phase Delay. Bragg cells are often used in heterodyne interferometers, and it is sometimes important to remember that the acoustic propagation delay translates into a huge phase shift. This acoustic phase shift gets impressed on the optical phase, and so appears in the phase data. It is often far from negligible; if you're using two passes through an 80 MHz cell that 200λ delay has a phase sensitivity of 31 rad/MHz. This is a nuisance in wide-range AOD measurements, or where it makes the oscillator spurs show up in the data, but in fixed-frequency applications it can be useful—you can stabilize the operating point of your system by using feedback to adjust the acoustic frequency. Shear-wave TeO₂ devices are the best overall in the visible. Optics people used to birefringent materials with $(\delta n)/n$ of a percent or less are usually surprised that the slow shear wave in TeO₂ goes at 600 m/s while the longitudinal wave goes at 4200. A really big TeO₂ cell can have 1000 resolvable spots in a single pass, though several hundred is more typical.

While they're the best of the fast bulk-optics modulators, AO cells have some major drawbacks. Cells with small Bragg angles (longitudinal devices at lowish frequency) have severe etalon fringes. AODs are less prone to these, because of the high angles, high diffraction efficiency, and polarization shift. There is also usually some beam apodization due to acoustic nonuniformity, and ripples in the diffraction efficiency in space and frequency due to acoustic standing waves. The standing wave effect is eliminated in shear wave devices by cutting the bottom of the cell at 5°; because of the huge Δv , this totally destroys the phase matching between the reflected acoustic wave and the light.

Some people say that AO cells have poor wavefront fidelity, but the author has never had a moment's problem with it. Scanning rapidly does introduce aberrations however. It takes some time for the wave to cross the beam, so a frequency ramp produces astigmatism by sending rays at different x in different directions; a frequency staircase with time allowed for settling avoids this problem. The polarization eigenstates of a shear wave cell are also very slightly elliptical, which one occasionally needs to remember. Overall, a slow shear wave AOD is a pretty trouble-free device.

7.10.6 Photoelastic Modulators

Besides a change in refractive index, the acousto-optic effect also induces stress birefringence in the crystal. Normally we don't worry too much about this, especially with TeO₂

devices, where the incoming beam has to be polarized in the right direction to get good performance. Photoelastic modulators are acousto-optic devices that exploit this effect. The basic idea is to run an AO cell at a very low frequency, 20–100 kHz, so that the acoustic transducer basically just shakes the entire crystal back and forth, and tune the frequency to the lowest order longitudinal vibration mode of the crystal—just like an organ pipe. The acousto-optic effect leads to a more or less uniform phase modulation, but the stress birefringence (the photoelastic effect) basically turns the crystal into an oscillating wave plate, whose retardation can reach $\pm\frac{1}{2}$ wave. Photoelastic modulators thus act a bit like acoustic Pockels cells, only much slower. Their big advantage is greater uniformity of the birefringence across their field.

7.10.7 Acousto-optic Laser Isolators

The acousto-optic effect is symmetrical, so reflected light propagating backwards along the beam will be diffracted back into the laser. The returned first-order beam enters the cell at $f \pm f_A$ from its first pass, and winds up at $f \pm 2f_A$, because the sign of the optical \mathbf{k} vector has changed. The frequency shift is so small that the difference in deflection angle is negligible, and the light goes straight back into the laser.

The laser is a Fabry–Perot resonator, though, and provided f_{acoustic} has been properly chosen, and the cavity finesse is high (as in gas lasers), virtually none of that light will make it back into the laser cavity to cause feedback problems. Even with diode lasers, where the finesse is low, and a lot of the light does make it into the cavity, the beat frequency $2f_A$ is so much higher than the ~ 100 kHz of mode hops that its effect is much reduced. (Why doesn't this work with the zero-order beam?)

Fiber Optics

All mankind in our age have split up into units, they all keep apart, each in his own groove; each one holds aloof, hides himself, hides what he has from the rest, and he ends by being repelled by others and repelling them.

—Elder Zosima, in *The Brothers Karamazov* by Fyodor Dostoevsky

8.1 INTRODUCTION

We all know that fiber optics and lasers have revolutionized long-haul and high bandwidth terrestrial communications. Fibers are unequaled at that job, there's no doubt about it. They're also good for instruments, though the situation there is less clear.

Fiber optics really constitutes a different domain, a bit like \mathbf{k} -space, where different things are easier or harder than in bulk optics. There is some new physics available in fibers, and a lot of old physics under a new aspect and in new operating regimes. The trade-offs are different enough that they're worth a considerable amount of study, so we'll concentrate on the characteristics of fibers in instruments. There are a lot of books available on fiber communications and fiber sensors, but not a lot on why to choose fiber or bulk for a particular application.

The thing to remember about using fiber is that it's far harder than it looks. It's seductive; a fiber sensor can easily be made to work at some level, which makes it easy to suppose that making that sensor robust enough for field use is straightforward—and it isn't straightforward at all.

8.2 FIBER CHARACTERISTICS

An optical fiber is a thin cylinder of doped fused silica, with nonuniform n near the center, covered with a protective plastic or metal jacket. It works as a very low loss dielectric waveguide. The most common type is *step-index fiber*, where a central *core* has a uniform index n_1 , which abruptly changes to n_2 where the annular *cladding* begins. The cladding index is also constant out to the edge of the jacket. Most fibers use a germania (GeO_2) doped silica core and a pure silica cladding. (See Table 8.1.)

TABLE 8.1. Fiber Parameters

Core radius	a	2.5–5 μm (single mode) 25–900 μm (multimode)
Core index	n_1	1.47–1.52
Cladding index	n_2	1.464
Normalized index difference	$\Delta = \frac{n_1^2 - n_2^2}{2n_1^2}$	0.004–0.04 (single mode) Δ usually small
Mode volume	$V = an_1k_0\sqrt{2\Delta}$	1.8–2.4 (single mode)
Fracture strength	S	0.01–4 GPa
NA (multimode)	$\text{NA} = \sqrt{n_1^2 - n_2^2} = n_1\sqrt{2\Delta}$	0.1–0.4
Étendue	$n^2A\Omega$	10^{-8} to 10^{-5} $\text{cm}^2\cdot\text{sr}$

8.2.1 Fiber Virtues

Fibers have lots of virtues, even besides those that all optical components share, such as wide temporal bandwidths.

Cheapness. It's nearly free (10 cents per meter in large quantity for single-mode communications fiber), and components are getting cheaper. Fiber connectors are amazingly cheap for such precise artifacts; they're only a few dollars, as long as you're using 125 μm OD fiber. You can thus bolt stuff together with fibers very quickly, which in principle would allow checking out a measurement idea faster than with bulk optics, where everything has to be aligned manually. And components such as directional couplers are mass produced, which makes them far cheaper than previously.

By putting the optical complexity at the far end, the business end of a fiber sensor can be made very small and cheap, even cheap enough to be disposable, as in fiber catheters, or sacrificial, as in fiber strain gauges cast into concrete structures or shrapnel velocity sensors for ordnance development.

Good Transmission Properties. For long distance transmission, fibers are tops, as already mentioned. They have very low loss (0.2–10 dB/km), exhibit no crosstalk, and transmit over a reasonably wide wavelength range—250–2000 nm at least, and special fibers go somewhat beyond both extremes.

EMI Immunity. Not being inductive or capacitive at ordinary frequencies, fibers are totally immune to electromagnetic interference (EMI), ground loops, and pickup; you can send a very weak signal a long way in a hostile environment. This is a key virtue for distributed sensing and smart structures.

Versatility. Most things that you can do with bulk optics, you can do with fiber, at least in principle. It provides some altogether new capabilities and more convenient ways of doing the same things, or is lower in cost because of economies of scale (driven by fiber transmission), simplified mounting, and reduced size and mass.

Ease of Use. Fibers are mechanically flexible and can bend light around in a pretzel shape if you want. They couple well to many kinds of lasers. Once you've got the light going in and out the ends, no more alignment is necessary. In connectorized systems,

the alignment is done by the manufacturer of the connector. Convenient collimators are available inexpensively.

Environmental Robustness. Fibers, being made almost entirely of quartz, survive high and low temperatures extremely well. When properly jacketed, they are immune to dirt and most fluids, too, though they become much more vulnerable to breakage when wet.

New Physics. The long interaction lengths possible with fiber allow us to base good measurements on intrinsically very weak effects, such as the Faraday effect in silica. We can make all-fiber sensors of completely new types, for example, multiplexed or distributed sensors like OTDR and fiber Bragg grating types.

8.2.2 Ideal Properties of Fiber

Besides these comparative advantages, fiber systems have a few *perfect* characteristics, which can be used to good effect.

Pointing Stability. Single-mode fibers have essentially perfect pointing stability, because with due care, we can ensure that really only one mode reaches the end facet. The mode pattern depends only on wavelength and material properties, so it is extremely stable, merely changing very slightly in diameter and NA with temperature, due to $\partial n/\partial T$ (see Example 9.7).

There-and-Back Phase Symmetry. There is a broad class of cases where the one-way phase delay is exactly the same in both directions in a fiber. The mode structure doesn't change when we replace k_z with $-k_z$, so the propagation velocity of a single mode is unaffected (Figure 8.1). In the face of mode coupling, we have to be a bit more careful. For the mode coupling to undo itself, the fiber has to be lossless and the




Figure 8.1. There-and-back polarization stability with Faraday rotator mirrors.

polarization has to come back in complex conjugated (i.e., with the same major axis and helicity) from the way it went out (see Section 8.3.3 below).

Scatter will degrade the matching stochastically, and there are also three effects that modify it deterministically: Pancharatnam's topological phase (see Section 6.2.4) in non-planar systems, the Sagnac effect in accelerated or rotating ones, and the Faraday effect in magnetically excited ones.

Transient effects will of course arrive at different times in the two directions, depending on where in the fiber they occur. Although silica is a very linear material, in a long fiber path you may see nonlinear effects, for example, the optical Kerr effect that causes phase offsets in Sagnac interferometers.

Two-Pass Polarization Stability. Polarization has two degrees of freedom per mode. Light in a lossless single-mode fiber can thus be resolved into two orthogonal polarization states, and given any one state ϕ_1 , we can always construct a ϕ_2 that is orthogonal to it.

Light launched into the fiber gets its polarization modified wholesale, due to bend birefringence, topological phase, twist-induced optical activity, and other effects, but throughout the process, ϕ_1 and ϕ_2 remain orthogonal, and so maintain their identity if not their appearance. This is a consequence of the unitarity of the Jones matrix of a lossless polarization element (see Section 6.10.2); an infinitesimal length of fiber can be represented by a differential Jones matrix, and the whole fiber's Jones matrix follows by passing to the limit of many segments of very small length.

A Faraday rotator mirror (or FRM, see Section 6.10.11) turns any monochromatic polarization state into its orthogonal partner, which retraces its path through the fiber but remains orthogonal to the first pass. Thus the light reemerging from the input facet has a stable polarization, which is orthogonal to the input light, and so can be separated from it with a polarizing prism. Of course, really high polarization selectivity (e.g., 50 dB optical) requires unfeasibly accurate fiber couplers, end faces, and FRMs.

8.2.3 Fiber Vices

Fiber has its own peculiar vices as well, some of which are crippling, and few of which are obvious. Most stem from having a too accurately aligned system with a long path length in glass that's hard to get light into. Attentive readers may note a certain similarity between the lists of virtues and vices—that's where our discussion of fibers in instruments centers.

Fiber Isn't So Cheap in Real Life. That 10 cents per meter pricing requires you to buy many kilometers of fiber, and it only applies to garden-variety single-mode telecom fiber—not PM or multimode fiber. Fiber components and connectors are also not cheap except at 1310 and 1550 nm and 125 μm OD.

Mechanical Fragility. It's made of glass, after all, and is vulnerable to snagging, kinking, and abrasion. Bending also causes loss.

Really Tiny Étendue. It's so low that fiber doesn't couple well to anything but focused laser beams. The same effect is responsible for fiber's horrible etalon fringe problem—there's nowhere for reflections to go except exactly back the way they came. (Angle polishing helps a lot, but not enough.) Its power handling capacity is limited, too.

Mode Instability. Multimode fiber, especially if it has only a few modes, exhibits large amounts of noise and drift in its illumination pattern due to mode coupling.

Etalon Fringes and FM–AM Conversion. Every optical system is an interferometer, and fiber optic systems are several interferometers at once, each one demodulating phase noise like a delay line discriminator. The too accurate alignment leads to severe etalon fringes and laser feedback and (as we saw in Section 2.5.3) removes the spatial averaging that helps so much to eliminate the FM–AM conversion from interference with delayed light.

Besides the obvious facet reflections, there are distributed effects that are often even larger; multiple Rayleigh scattering within a long fiber causes serious FM–AM conversion too (see Section 13.5.6). This sounds like a trivial effect, but unfortunately it is not; it is usually by far the dominant broadband AM noise source in a single-mode fiber sensor when the fiber length is a significant fraction of both a coherence length and an absorption length.[†] The physics is similar to the coherence fluctuation noise of Sections 2.5.3 and 19.1.1.

Phase and Polarization Instability. The long path in glass makes fibers vulnerable to piezo-optic shifts, temperature changes, vibration, and bending, so that every fiber device has serious instabilities of phase and polarization with time.

Sensitivity to Everything Except EMI. Use of fibers for DC measurements is extremely difficult, because apart from pure intensity instruments, every fiber sensor is also a thermometer. For example, fiber strain sensors are plagued by temperature sensitivity, because 1 °C of temperature change gives the same signal as a fairly large strain (10 $\mu\epsilon$), and separating the two effects well enough to get high sensitivity thus requires very accurate, independent temperature measurements, and an accurate knowledge of the response of the system to temperature alone.

8.3 FIBER THEORY

In order to be able to use the good properties of fiber, and avoid being blown up by the bad ones, we need to spend some time on their theoretical properties.

8.3.1 Modes

A step-index fiber can be modeled as a right-circular cylinder of radius a and index n_1 , surrounded by a practically infinite thickness of n_2 . Most of us probably remember how to solve the wave equation in a uniform medium with cylindrically symmetric boundary conditions: the Laplacian operator separates, so we get a one-dimensional wave equation in the axial direction, Bessel's equation in the radial coordinate, and an eigenvalue equation for the azimuthal coordinate, leading to solutions of the form

$$\begin{aligned}\psi_m^{(1)}(r, \phi, z) &= J_m(\kappa r)e^{i(m\phi+k_z z)}, \\ \psi_m^{(2)}(r, \phi, z) &= N_m(\kappa r)e^{i(m\phi+k_z z)},\end{aligned}\tag{8.1}$$

[†]A. Yariv, et al., *IEEE Trans. Lightwave Technol.* **10**(7), 978–981 (1992).

where J_n and N_n are the Bessel functions of the first and second kind, respectively, and

$$k_z^2 + \kappa^2 = k^2 = n_1^2 k_0^2. \quad (8.2)$$

From experience, we expect to apply boundary conditions and get a denumerable set of modes for any given k . For a finite radius, only a finite number can propagate, as usual.

For step-index fibers, we have two dielectrics, so we have to patch the fields together at the boundary, using the phase matching condition in z (because of translational invariance) and continuity in r and ϕ of tangential E and H and of perpendicular D and B , leading to an eigenvalue equation for the allowed values of κ . This requires a significant amount of bookkeeping and results in four families of modes— TE , TM , EH , and HE . For large diameters and small index differences, these can be combined into nearly transverse electromagnetic (TEM) forms, and the scalar approximation is valid; thus we use the scalar patching conditions, that is, phase matching in z and continuity of ψ and $\partial\psi/\partial r$. With a little bit of hand work, this yields the linearly polarized LP modes. While these solutions are approximate, it's much better in practical work to have an intuitive approximation than an unintelligible exact result.

We normally confine our attention to those modes that are guided, so that the field decays with r outside the core of the fiber. Such modes occur when k_z is large enough so that (8.2) forces κ to be imaginary in the cladding while being real in the core, which occurs when

$$n_2 k_0 < k_z < n_1 k_0. \quad (8.3)$$

When κ is imaginary, the solutions are more conveniently written in terms of the modified Bessel functions,

$$\begin{aligned} \psi_m^{(3)}(r, \phi, z) &= I_m(\gamma r) e^{i(m\phi + k_z z)}, \\ \psi_m^{(4)}(r, \phi, z) &= K_m(\gamma r) e^{i(m\phi + k_z z)}, \end{aligned} \quad (8.4)$$

where $k_z^2 - \gamma^2 = n_2^2 k_0^2$. These functions are both monotonic: $I_m(r)$ grows exponentially and is regular at 0, while $K_m(r)$ decays exponentially but has a pole at 0. Because the cladding is very thick and doesn't extend to 0, only the $K_m(\gamma r)$ solution is physical, and the patching conditions lead to an eigenvalue equation for κ , which sets the allowed modes.

This may seem a bit *recherché*, but it has a very important consequence: because we have to patch J_m and $d/dr(J_m)$ together with a monotonically decaying function, $|J_m(\kappa r)|$ must be decreasing at $r = a$. Thus the m th radial mode can only propagate when $d/dr(|J_m(\kappa a)|) \leq 0$. Below the first zero of $J_1'(\kappa a)$, only a single mode can propagate, so the fiber is single mode when

$$\lambda_c > \frac{2\pi}{2.405} a n_1 \sqrt{2\Delta}. \quad (8.5)$$

The light is not confined entirely to the core, as it is in a metal waveguide, but spreads a considerable way into the cladding; the patching condition can always be met for the

lowest mode, so there is no long-wavelength mode cutoff of a fiber the way there is for a metal waveguide. On the other hand, the guiding becomes weaker and weaker as more and more of the light is in the cladding, so there is a practical minimum for the core diameter. Guiding behavior is captured by the *mode volume* or *normalized frequency* V , which from Table 8.1 is

$$V = an_1k_0\sqrt{2\Delta}. \quad (8.6)$$

Good guiding and reliable single-mode operation occurs for V from about 1.8 to 2.4. The lowest order mode, LP_{01} , has a round top and exponentially decaying sides, so it looks roughly Gaussian. The *mode field radius* w_0 is defined similarly, as the $1/e^2$ radius of the best-fit Gaussian, and for single-mode fiber is approximately

$$w_0 \approx a \left(\frac{1}{3} + \sqrt{\frac{2.6^3}{V}} \right). \quad (8.7)$$

As with Gaussian beams, we normally care about the details of only the lowest order mode, because any device that relies on a multimode fiber having exactly predictable mode properties is doomed; a multimode fiber has such strong coupling between modes that the mode amplitudes and phases change rapidly with propagation distance and environmental perturbations. One exception is when we know we have exactly two well-behaved modes, for example, in a polarization preserving fiber (see Section 8.4.4).

8.3.2 Degeneracy

When we talk about a fiber being *single mode*, we are of course talking about the scalar approximation, whereas light in the fiber has two degrees of freedom in polarization. For perfect cylindrical symmetry and homogeneous, isotropic materials, the fiber has rotational symmetry, which means that the two polarization states have exactly the same k_z ; that is, they are *degenerate*. Degenerate states don't get out of phase with each other, so even a very small coupling causes complete energy redistribution between them, as we'll see.

8.3.3 Mode Coupling

Many fiber devices are based on weak lossless (*adiabatic*) coupling between two waveguide modes, either the LP_{01} modes in two fibers brought together, as in a directional coupler, or between two modes in a single fiber. By *weak*, we mean that the mode amplitudes change on a scale much longer than a wavelength, not that the energy transfer isn't large eventually. Under these conditions, we can treat mode coupling by first-order perturbation theory, which is a good break because for that we need only the zero-order waveguide modes. Consider the situation of Figure 8.2: two modes, ψ_1 and ψ_2 , that are orthonormal[†] except for a small coupling, which is constant in z and t . We write the

[†]That is, orthogonal and with the same normalization, so that unit amplitude in one mode corresponds to unit power.

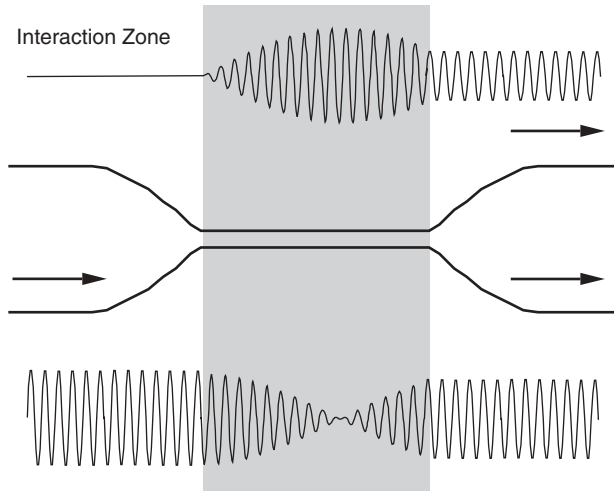


Figure 8.2. Coupled modes.

mode amplitudes as $(A_1, A_2)^T$. The first-order perturbation equation for \mathbf{A} is then

$$\frac{d}{dz}\mathbf{A} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{bmatrix} \begin{bmatrix} A_1 \\ A_2 \end{bmatrix}, \quad (8.8)$$

with $|c_{12}|, |c_{21}| \ll |k_1|, |k_2|$. If the perturbation is lossless, then $d/dz|\mathbf{A}|^2 = 0$. Rearranging a few indices leads to

$$\underline{\mathbf{c}} + (\underline{\mathbf{c}}^T)^* = 0 \Rightarrow \underline{\mathbf{c}} = \begin{bmatrix} ik_{z1} & c_{12} \\ -c_{12}^* & ik_{z2} \end{bmatrix}, \quad (8.9)$$

that is, $\underline{\mathbf{c}}$ is skew-hermitian,[†] and all its eigenvalues are therefore imaginary. We'd expect the pure imaginary elements on the main diagonal, because that's the behavior of the uncoupled modes; the skew-hermitian symmetry of the off-diagonal elements makes sure that the power leaving one mode winds up in the other.

This constant-coefficient system is solved by finding the eigenvalues $i\beta_j$ and eigenvectors ϕ_j of $\underline{\mathbf{c}}$, because in the eigenvector basis (the principal components basis), the equations decouple, leading to solutions $\psi_j(z)$ with a simple exponential z dependence, $\phi_j = \exp(i\beta_j z)$,

$$\beta_j = \frac{1}{2} \left(k_{z1} + k_{z2} \pm \sqrt{(k_{z1} - k_{z2})^2 + 4|c_{12}|^2} \right). \quad (8.10)$$

We'll use primed quantities in this basis (e.g., \mathbf{A}' is the principal component representation of \mathbf{A} , but it's the same vector really). If the two waveguide modes have the

[†]A skew-hermitian or antihermitian matrix is i times a hermitian one, so with that modification we can apply all the usual theorems about hermitian matrices—orthogonal eigenvectors, real eigenvalues, and so on.

same k_z (e.g., the fiber coupler case, or polarization coupling in an ordinary single-mode fiber),

$$\beta_i = k_z \pm |c_{12}| \quad (8.11)$$

and ϕ_j becomes

$$\phi_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix}, \quad \phi_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}. \quad (8.12)$$

This isn't too spectacular looking—just two quadrature pairs with slightly different propagation constants. Now, however, let's consider the usual case, where all the light is initially in $\psi_1 = 1/\sqrt{2}(\phi_1 + \phi_2)$, so $\mathbf{A}(0) = (1, 0)^T$ and $\mathbf{A}'(0) = (\beta_1/\sqrt{2}, \beta_2/\sqrt{2})^T$. In that case, we get a time dependence

$$\mathbf{A}(z) = e^{ik_z z} \begin{bmatrix} \cos(c_{12}z) \\ i \sin(c_{12}z) \end{bmatrix}, \quad (8.13)$$

so that after a distance $\pi/(2c_{12})$, all the energy has moved from ψ_1 to ψ_2 , and it then transfers back and forth sinusoidally, forever. Thus regardless of the magnitude of c_{12} , we can pick a length for our coupler that will give us any desired output power ratio from 0 to ∞ .

If $k_{z1} \neq k_{z2}$, we still get sines and cosines, but with offsets due to the maximum coupling being less than 1; the power doesn't finish moving from ψ_1 to ψ_2 before the phase shift between modes exceeds π , and it starts moving back again. It's useful to define the *beat length* λ_b , which is the period of the spatial beat between modes,

$$\lambda_b = \frac{2\pi}{|k_{z1} - k_{z2}|}. \quad (8.14)$$

Looking at the square root in (8.10), there are two distinct limits,

$$|c_{12}|^2 \gg (k_{z1} - k_{z2})^2/4 \quad (8.15)$$

(long beat length) and

$$|c_{12}|^2 \ll (k_{z1} - k_{z2})^2/4 \quad (8.16)$$

(short beat length), which govern how we do the binomial expansion of the square root (see the problems in the Supplementary Material). In the long beat length limit, the situation is very similar to the degenerate case, except that the maximum coupling is reduced slightly. In the short beat length limit, though, the rapidly growing phase shift between modes prevents much power from coupling at all, so the modes remain more or less uncoupled.

8.3.4 Space-Variant Coupling

The situation is quite different when the coupling coefficient c_{12} is a function of z , especially if we make it periodic near λ_b . If we do so, for example, by launching an acoustic wave of that wavelength, making a photorefractive Bragg grating, or by pressing a corrugated plate against the fiber, we can recover the strong coupling that we had lost. If the coupling has spatial frequency $2k_z$, light going to positive Z will get backscattered strongly, as in a fiber Bragg grating. The coupled-mode theory becomes a bit more complicated in this case, but is still quite manageable.

8.3.5 Dispersion

Because different order modes have different k_z , they have different phase velocities,

$$v_p = \frac{\omega}{k_z} \quad (8.17)$$

and because of the nonlinearity of the expression for k_z in terms of k and κ , their group velocities

$$v_g = \frac{\partial \omega}{\partial k_z} \quad (8.18)$$

are different as well. This *intermodal dispersion*[†] severely limits the transmission bandwidth of step-index multimode fibers, as is well known. A simple ray optics model of a ray bouncing off the core–cladding boundary at exactly θ_C gives an estimate of it,

$$\Delta v_m = \frac{c}{n_1} \left[1 - \frac{1}{\sqrt{1 + (\text{NA})^2/n_1^2}} \right] \approx \frac{c(\text{NA})^2}{2n_1^3}, \quad (8.19)$$

which for an NA of 0.2 predicts a velocity spread of over 1%—around 45 ns/km.

The exact modulation behavior of a given fiber depends on its length, manufacturing variations, and the details of its mode spectrum. We can feel pretty safe as long as the phase spread between modes is below a radian or so at the highest modulation frequency, but we expect to be in big trouble when it gets to half a cycle. That leads to a bandwidth–distance trade-off,

$$L \times \text{BW} < \frac{cn_1}{\pi(\text{NA})^2}, \quad (8.20)$$

which is pretty sharp; that 45 ns/km limits us to 3.5 MHz in a 1 km length. Bending the fiber into a coil reduces this effect and in rectangular cross-section waveguides can almost eliminate it; bending redistributes light between low- and high-angle modes, so that most of the light will spend some time in each, reducing the transit-time spread. (This is rather like the quenched mode coupling of Section 8.12.1.)

Gradient-index fiber works like a big long GRIN lens, ideally making rays at all angles see the same propagation delay. Waveguide effects and manufacturing errors limit

[†]This is of course a different use of the word *dispersion* than elsewhere in optics, where it refers to variation of n with λ .

its intermodal dispersion to 0.5–2 ns/km, which is much better, but still pretty big. Single-mode fibers don't have this problem, which (along with much lower loss and lower fiber cost) is why they are popular.

A single-mode fiber suffers *waveguide dispersion*[†] as well, because the fiber structure influences κ , so that k_z is a nonlinear function of k , and v_p slows down at long wavelengths. However, the quartz also has normal material dispersion, with n increasing at short wavelengths; for ordinary silica fiber, the two cancel out in the neighborhood of 1.3 μm , the zero-dispersion wavelength. You can send a fast signal a long way at 1.3 μm .

The loss minimum is at 1.55 μm , not quite the same as the dispersion minimum. In order to save money on repeaters, people have designed *dispersion-shifted fiber*, whose index profile is modified to reduce the waveguide dispersion enough to shift the zero-dispersion wavelength to 1.55 μm . There is also *dispersion-flattened fiber*, which is a compromise between the two and has low but not zero dispersion between 1.3 and 1.55 μm .

8.4 FIBER TYPES

8.4.1 Single-Mode Optical Fibers

Single-mode fiber is pretty well behaved stuff on balance; it will carry light long distances without doing serious violence to it, except for messing up its polarization, and it can be coupled into and out of good quality laser beams without great difficulty (once you get used to it). Once you get the light in, you can pipe it around corners, split it, combine it, modulate it, attenuate it, and detect it much the way you would with RF signals in coax. Due to the fiber's polarization degeneracy, the polarization of the transmitted beam is unstable, and because of its perfect spatial coherence, there are a lot of etalon fringes that have to be controlled with angled fiber ends, which is very awkward. Most of the single-mode fiber sensors we'll encounter later are limited by these twin problems.

The étendue of a single-mode fiber is the smallest of any optical component's—it's about the same as the $A\Omega$ product of a perfectly focused Gaussian beam, which is $\lambda^2/2$, on the order of 10^{-8} $\text{cm}^2\cdot\text{sr}$ at 1.5 μm . A laser is the only source that can cope with that.

8.4.2 Multimode Optical Fibers

Multimode fibers have a large V , which means that many modes can propagate; the number goes as V^2 , and is on the order of 200 for a step-index fiber. For the same value of V , a graded-index fiber will have half the number of modes of a step-index fiber. (Why?) These modes couple into one another at the slightest provocation, and as we saw, in step-index fibers they have very different propagation velocities, leading to severe bandwidth degradation.

Graded-index multimode fibers are the most common type; making the index of the core decrease smoothly from a broad peak on the axis brings the propagation velocities of the modes very close to one another, and so reduces the intermodal dispersion to a

[†]The term “mode dispersion” should be avoided, because different writers use it to mean both intermodal dispersion and waveguide dispersion.

few nanoseconds per kilometer. With laser light, the output of the fiber is a speckly mess that wiggles around like sunlight filtering through the trees on a windy day. It's even worse if the whole mode volume hasn't been excited, as often happens when we couple a laser into it with a low-NA lens.

On the other hand, multimode fiber is pretty stable with broadband illumination, if we carefully excite all the modes, for example, by using a diffuse source such as an integrating sphere, or homogenizing them by mandrel wrapping.

Multimode fiber is inherently lossier than single mode. High angle modes are weakly guided and so are vulnerable to scatter from irregularities in the fiber, or to bending and pressure. You see this as a gradual decrease of the measured NA of the fiber as its length increases. This hits a steady state eventually, since low angle light gets coupled into high angle modes, eventually spreading the loss evenly throughout the modes. The high loss encountered by high angle modes makes the far field pattern of a long piece of multimode fiber approximately Gaussian.

We often want to couple thermal light into multimode fibers, and so need to know something about their étendue. It's pretty easy to calculate for a step-index multimode fiber, because the acceptance angle is more or less constant across the facet, so $n^2 A \Omega' \approx (\pi a^2)(\pi(\text{NA})^2)$. Calculating the étendue of graded-index fiber is not so easy, because high angle rays near the edge of the core aren't well guided, but go into leaky modes. We can intuitively see that the étendue will be significantly less, because high angle rays can only enter near the fiber axis. A pure ray optics approach could calculate the NA as a function of radius and integrate up to $r = a$, but we'll content ourselves with arm waving here; it turns out that GI fiber has half the modes of SI.

At any given wavelength, each mode is fully coherent, and so we can imagine making a couple of multiple-exposure holograms that would transform each individual mode into a different focused spot. Thus each mode should contribute about the same étendue as a focused spot, that is, about $\lambda^2/2$. We thus expect the étendue of fiber with a given NA to go as

$$n^2 A \Omega' \approx \frac{\lambda^2 N}{2}. \quad (8.21)$$

The wisdom here is that to pipe low spatial coherence light around, step-index multimode is really the way to go unless you need the fast temporal response.

8.4.3 Few-Mode Fiber

The worst case is fiber with only a few modes, for example, telecom fiber used in the visible. There aren't enough modes for their spatial pattern to average out, and their coupling is a really serious source of noise because their propagation speeds are very different; thus the relative phase delay grows rapidly with z , leading to FM-AM conversion. Because the two modes will not in general have the same polarization, phase drift leads to severe polarization instability. Furthermore, the pointing instability of few-mode fiber is appallingly bad, with the beam wandering around by a degree or more depending on the phase delay.

8.4.4 Polarization-Maintaining (PM) Fiber

Polarization-maintaining (PM) fiber is SMF intentionally made asymmetric in profile to break the polarization degeneracy of circular-core fiber, and is seriously misnamed.

A built-in stress birefringence or (less often nowadays) an elliptical core makes the two polarizations travel at slightly different speeds. The speed difference is fairly slight; the beat length is a few centimeters in ordinary PM fiber to 1–2 mm in high birefringence (HiBi) fiber. As we've just seen, a short beat length means that the coupling of the two modes is incomplete. If we launch light into the fiber in one polarization state and look at what comes out, we find that the polarization is reasonably stable; regardless of mild abuse, bending, and so on, the polarization ratio at the output is typically 20 dB. That's enough to help, but far from enough to fix the polarization noise and drift problem; a 20 dB polarization ratio means that the polarization axis can wander $\pm 6^\circ$, and you aren't going to do a dark field measurement with that.

Furthermore, the short beat length means that the delay discriminator action is working moderately hard at demodulating the FM noise of the laser; a PM fiber has lots of noise if the orthogonal output polarizations are made to interfere, for example, by a misaligned polarizer or an angled fiber end, and the mode coupling eventually does the same thing to us.

The real reason that PM fiber is misnamed is that it doesn't preserve any old input polarization, just the two eigenmodes. Its birefringence drifts with temperature, and the drift rate is proportional to the total retardation, so high birefringence PM fiber (1 mm beat length) with equal amounts of light in both modes has around *three orders of magnitude more* polarization drift with temperature than ordinary single-mode fiber.

We therefore use PM fiber with as pure a polarization mode as we can give it. Since the output polarization purity is typically 20 dB, using PM fiber that way makes the small-scale instability worse, but restricts it to a limited range of elliptical polarizations about the desired mode.

The great virtue of PM fiber is thus to stabilize the operating point of our polarization-sensitive sensors well enough that a properly normalized AC measurement can pull good data out without needing active stabilization, and that's truly a major point gained—we don't have to ship a graduate student with each instrument.

8.4.5 Chalcogenide Fibers for IR Power Transmission

A silica fiber will transmit at least some light from 250 to 2000 nm, and fluoride fibers reach 3000 nm, but the loss gets pretty large near the extremes. In the mid-IR, good fibers can be made from chalcogenide glasses. The main types are sulfide glasses, $\text{As}_{40}\text{S}_x\text{Se}_{60-x}$, which work from 1 to 6 μm , and telluride glasses, $\text{Ge}_a\text{As}_b\text{Se}_c\text{Te}_d$, for 3 to 11 μm . These are flexible enough, and low enough in loss, to be an attractive way of piping mid-IR power around, for example, CO_2 light for heating or surgery. Even at their best, chalcogenide fibers' losses are between 0.2 and 1 dB per meter, thousands of times worse than single-mode silica at 1.5 μm , but they can handle CW power levels of 100 kW/cm² and pulsed levels of 200–400 MW/cm².

8.4.6 Fiber Bundles

The relatively low cost of fiber lends itself to some creative brute force solutions. For example, say we want to be able to pipe light from a lamp and condenser around easily. The étendue of a multimode fiber may be 10^{-5} cm²·sr, but we can use 10^3 of them in

a bundle, which starts to be quite respectable (and no longer so cheap). Bundles work best with broadband illumination, since otherwise speckles due to mode shifting from bending and stress will drive you crazy. They come in all sizes from two fibers to 20 mm diameter or larger; it depends on how much you want to pay for material, and for all that laying, cleaving, and polishing. They are always multimode because of the ratio of the core and cladding areas, but cladding modes are important and may need to be removed by potting in black wax.

The basic distinction is between image bundles, where each point on the output surface connects to the corresponding point on the input, so that an image can be transmitted, and random bundles, where no such care has been taken. Random bundles aren't really very random; they won't homogenize a space-varying illumination function, for example. The fibers just go any old way they happened to land. Now that LEDs are so good, random-bundle illumination is much less useful than it once was—we might as well mount the LEDs at the business end and run wires instead.

Imaging bundles are further divided into fused and flexible types. A flexible bundle has the fibers glued together only at the ends, leaving it free to bend. Fused bundles are used mainly for manipulating images; a short bundle with one concave and one flat surface makes a fiber-optic field flattener (also called a fiber-optic faceplate), which gives image tubes a flat image plane. Two image tubes can be cascaded with a biconcave fiber plate, which works better than any lens system and is dramatically smaller too. Tapered bundles, made by softening the bundle and stretching one end, magnify the image by the ratio of the face diameters. You can even twist the softened bundle, to make a fiber image rotator.

Of course, the phase gets completely screwed up in the process of transmission—the fibers are multimode, and no two are identical anyway. None of these image bundles can produce an aerial image—if the image is out of focus on the entrance face, it's out of focus forever, just like a ground glass screen. That's why the common term *coherent bundle* is a misnomer. On the other hand, the bulk optics alternatives are either a long succession of relay lenses or a long, specially designed GRIN rod. These are no picnic to get right, on account of the buildup of field curvature in the lenses and the specialized nature of the GRIN rod, but most of all because they can't bend without fancy and delicate articulated mirror mounts at each hinge.

8.4.7 Split Bundles

You can get split bundles that allow you to couple both a source and detector to a remote head. This can be pretty useful sometimes; with a tungsten bulb on one end and a spectrometer on the other, it allows decent low resolution spectroscopic measurements in tight spaces, for example. Combining a split bundle with a light bulb, band-limiting filters, and a monochromator is easy and very successful. If the monochromator is actually an integrated grating spectrometer/OMA design that plugs into a PC, you can put a nice film thickness measurement device together inexpensively that will do measurements under water, in awkward places, and in the face of lots of vibration and shock.

Another application of bundles is in fiber optic slip rings. A modulated laser on a rotating part is focused onto a surface having an annular ring of fibers. All the fibers are gathered together into one photodiode at the other end. The focused spot radius is about one fiber diameter, so that communications don't drop out periodically as the spot rotates. Communication the other way can be done with a more much powerful laser

feeding the entire bundle, with the photodiode on the rotating part, but this is seldom done because of the laser power required.

8.4.8 Coupling into Bundles

We usually use a condenser to launch light into a bundle, as we saw in Example 2.1. The output end is a bit dicier. Well-homogenized multimode fibers illuminated with white light have nearly uniform output brightness across the core, but it narrows and dims for long lengths and short wavelengths, due to the excess loss in the high angle modes.

The fibers themselves do not pack all that tightly (see the problems), so that the illumination function is inherently nonuniform. The cleaved fiber ends have some nonzero wedge angle, which is usually nonuniform among the fibers in a bundle, so the far field pattern is a bunch of more-or-less overlapping cones rather than a single cone. This can lead to artifacts in grating spectrometers, for example.

Many-fiber bundles are more predictable than few-fiber ones, because the holes in the illumination pattern are smaller by comparison with the whole bundle. A fiber bundle used to receive light from the apparatus subjects it to a very odd spatial filtering function, which is usually poorly understood, as shown in Figure 8.3. Don't do it if you don't have to, and if you do, think about putting the bundle far enough out of focus that the pattern doesn't cause artifacts. A $125\ \mu\text{m}$ fiber has to be surprisingly far out of focus for the modulation of the intensity to be below 1%, say.

Fiber bundles are more or less telecentric, since the angular pattern from each fiber is more or less the same. You can get fiber ring lights, which go around the outside of a lens, and provide a nice uniform illumination with an odd annular pupil function.

8.4.9 Liquid Light Guides

A variation on the fiber bundle is the liquid light guide, which is just a huge flexible waveguide made from high index liquid in a low index plastic tube. They're flexible but have poor illumination patterns. These have been pretty well superseded by white LEDs mounted where the light is needed.




Figure 8.3. Pupil function of a fiber bundle.

8.5 OTHER FIBER PROPERTIES

8.5.1 Leaky Modes

Modes that are just allowed or just cut off are generically called *leaky modes*. They get excited any time there's a mode mismatch. Their losses are very high, of course, which makes the overall fiber loss look anomalously high for the first bit of fiber after a launcher, splice, or connector.

8.5.2 Cladding Modes

Even leakier are *cladding modes*, in which the jacket takes over the role of the cladding and the whole fiber that of the core. You'll get these nearly always, but they're especially troublesome with single-mode fiber because the étendue of the highly multimode-fiber/jacket system is far larger than that of the core/cladding system. Cladding modes can be eliminated in a short distance by stripping the fiber jacket and putting a blob of index-matched black wax on the cladding, or by mandrel wrapping.

8.5.3 Bending

Fiber guiding occurs because there is no propagating mode in the cladding whose k_z phase matches with the light in the core. As soon as we bend the fiber into an arc, though, that is no longer true; light going round the arc at a larger radius sees a longer path length, so it can move faster and still stay phase matched. For any finite bend radius, there is a distance into the (notionally infinite) cladding at which a propagating mode is phase matched.

For long bend radii, that isn't a problem, because the guided mode has died off more or less completely before reaching that distance. As the bend gets tighter and tighter, though, the phase matched distance gets closer and closer, until it starts overlapping the propagating mode significantly and light starts getting lost. At that point, the loss increases exponentially as the bend radius declines. You can easily see it by putting a HeNe beam through a fiber and bending it with your thumb and forefinger.

For a given fiber, the guiding gets weaker as λ increases. The *bend edge* is the longest wavelength for which guiding occurs, as a function of the radius of the bend. You normally won't get very close to it.

8.5.4 Bending and Mandrel Wrapping

Bending loss has its uses, especially in stripping unwanted cladding modes and filling up the mode volume in a multimode fiber. It's pretty safe to bend a fiber into a coil 600 fiber diameters across, and you can go as small as $200\times$ briefly. For lab use, where broken fibers are easily replaced, you can be much more daring; a $125\ \mu\text{m}$ OD fiber can be wrapped 10 times around a pencil to strip leaky modes and cladding modes, for example. In an unattended system, stick to the $600\times$ rule and use lots more turns, or use the black wax trick for cladding mode stripping instead.

8.5.5 Bend Birefringence and Polarization Compensators

Wrapping a fiber around a mandrel makes it birefringent, with the fast axis normal to the axis of symmetry. How birefringent a given piece of fiber becomes at a given wavelength

is usually a pretty well kept secret—you have to measure it. A rule of thumb is that the retardation in meters of a single-mode fiber wrapped into one turn of diameter D is

$$\beta_b = K_\lambda \frac{d^2}{D^2}, \quad (8.22)$$

where $k \approx 0.13$ m at 633 nm for silica fiber; a $5/80$ μm HeNe fiber looped into one turn of 72 mm diameter makes a nice quarter-wave plate, and that in turn allows us to make all-fiber polarization compensators.

Just as two quarter-wave plates can turn any fully polarized state into any other, so two loops of fiber attached to flat discs, hinged so as to allow them to rotate through 180° , will allow us to control the polarization fannies of our fiber, at least until it starts moving again. Three-paddle compensators (two $\lambda/4$ and one $\lambda/2$) are often used to provide some tolerance for errors and wavelength shifts, because if your wave plates are not exactly $\lambda/4$, there are some places you can't get to—just the way you can touch your nose to your shoulder and to your wrist, but not to your elbow. Since we're usually interested in linear polarization, after tweaking the two $\lambda/4$ paddles to get rid of ellipticity, we can get any linear polarization by turning the $\lambda/2$ paddle.

8.5.6 Piezo-optical Effect and Pressure Birefringence

Squashing a fiber produces a piezo-optical shift in n , as we saw in Section 8.5.6. Quartz isn't very piezo-optic, but there's a lot of it in your average fiber system—and even a random effect grows with \sqrt{L} .

8.5.7 Twisting and Optical Activity

Analogously, if we twist a fiber about its own axis, it becomes optically active. The total polarization rotation of a fiber twisted through ξ radians is

$$\Delta\theta = g\xi, \quad (8.23)$$

where g is about 0.15 for silica fiber at 633 nm.

Aside: Normal Propagation. When we discuss the propagation of light in circular core, single-mode fibers under peculiar conditions (e.g., bending, twisting, and so on), we are looking at the deviation of the light from the normal situation, that is, \mathbf{k} and \mathbf{E} being constant in the lab frame. Twisting a fiber drags the polarization direction along with the twist, but that's solely a material property. Maxwell's equations and the waveguide properties *per se* aren't altered, so without the material change, the light would continue with the same polarization in the lab frame, completely oblivious to the change in the guide. Things are a bit more complicated with PM fiber, because we have to look at the twist or bend as a coupled-modes problem.

8.5.8 Fiber Loss Mechanisms

Silica fibers are amazingly low loss devices, far more transparent than air in some spectral regions. Loss is dominated by Rayleigh scattering and microbending at short wavelengths, although there is also strong electronic absorption deeper into the UV.

The Rayleigh scatter goes as λ^{-4} , so it's very strong below 400 nm and almost nonexistent beyond 1.2 μm . There are overtones of the OH vibrational resonance at 2.73 μm , which occur in bands near 950, 875, 825, and 725 nm, and molecular vibrational absorption beyond $\sim 2 \mu\text{m}$. The sweet spot is in the 1–1.6 μm range, where the absorption, Rayleigh scatter, and microbending are all weak, and the fiber can achieve losses of 0.15–0.5 dB/km. By comparison, a 1 m thick piece of window glass looks as green as a beer bottle (try looking through a bathroom mirror edgewise sometime). The few meters of silica fiber in the typical instrument will transmit light from 250 to 2000 nm or thereabouts.

At 254 nm, silica fiber's absorption goes up to around 1000 dB/km, and permanent darkening or *solarization* may start to occur due to UV damage. You can get solarization-resistant fiber that works stably down to 180 nm (1.5 dB/m). Multimode fiber is especially flaky in the UV, because the mode coupling is very strong. Low angle modes get coupled into high angle modes, which are attenuated much more strongly. Wiggling the fiber makes the amplitude go up and down, by as much as 1 dB/m at 254 nm.

Minor manufacturing errors also contribute to fiber loss; small-scale roughness at the core–cladding interface (a form of microbending) is the major one, but there is also the occasional bubble or inclusion. Microbending is worst at high Δ , just as high index lenses are more vulnerable to surface errors. These intrinsic losses dominate at long lengths, but in instruments we're more often fighting splices, connectors, and larger scale bending.

8.5.9 Mechanical Properties

Glass is a brittle material, which is to say that cracks that are not held shut will eventually propagate through it, causing failure. The time to failure goes as a very high negative power of the tensile stress. At the crack tip, chemical bonds are being broken, which takes energy. It takes a lot less energy if there's water around to bond to the newly created surface, so the time to failure is enormously longer for dry fiber versus wet.

It's odd at first sight, but the strength of a fiber depends on its length. For a uniformly loaded fiber, it's the largest flaw in the entire length that causes failure, just like the weak link in a chain. The statistics of how big the largest flaw is are sensitive to how long it is, and of course very sensitive to the history of that piece of fiber.

The fracture strength of fibers thus varies all over the place, with some published curves showing 0.1% probability of fracture per kilometer at stresses of only 50 kPa, which is less than the stress on your spine from holding up your head. On the other hand, a really good piece of fiber breaks at around 4 GPa, which is three times stronger than the best steel.

8.5.10 Fabry–Perot Effects

In Section 11.9.2, we'll see that minor misalignment protects us from etalon fringes to a considerable degree. There's no way to misalign a single-mode fiber, so we would seem to be in trouble; along with the polarization instability, etalon fringes are the leading cause of trouble in laser-based fiber instruments. One thing that helps is to cut the ends of the fibers at an angle, so that the reflection is outside the fiber NA, and so goes off into the cladding to be absorbed. Fibers tend to cleave normal to the axis, but twisting them while cleaving will produce an angled facet. Angled cleavers exist that work reasonably

repeatedly, and you can get angled physical contact (APC) connectors and matching fiber polishing kits. These will get you down to around 10^{-4} reflectance, which is still poor in the scheme of things, but is pretty good for fiber. We're often forced to choose fragile APC connectors and expensive Faraday isolators, which significantly reduces the attraction of fiber.

One gotcha is that the etalon fringes in fiber are polarization dependent, because different polarizations see different delays. In a broadband system, this will decorrelate the noise in the two polarization eigenstates.

8.5.11 Strain Effects

If you imagine stretching a hollow metal waveguide to $1 + \epsilon$ times its original length, it's easy to see that the two main effects are a propagation distance increased by the same factor, and a decreased propagation speed due to the slight narrowing of the guide due to the stretch.[†]

Fiber behaves the same way, with the addition of the change in n due to strain. The rate of change of phase with respect to stretching in a fiber is

$$\frac{d\phi}{dL} = nk_0\xi, \quad (8.24)$$

where ξ is given by

$$\xi = 1 - \frac{n^2}{2}(P_{12} - \mu(P_{11} + P_{12})) \approx 0.78 \quad (8.25)$$

so that

$$\frac{d\phi}{dL} \approx 0.78nk_0. \quad (8.26)$$

8.5.12 Temperature Coefficients

We saw in Section 4.2.2 that the temperature component of the optical path length is


$$\frac{\text{TC}_{\text{OPL}}}{\text{OPL}} = \frac{1}{n} \frac{\partial n}{\partial T} + \text{CTE}. \quad (8.27)$$

The CTE of fused quartz is very small, about $5 \times 10^{-7}/^\circ\text{C}$, but its TCN is around $+9 \times 10^{-6}/^\circ\text{C}$, so that its TC_{OPL} is $7 \times 10^{-6}/^\circ\text{C}$.

Although the exact value is slightly modified by waveguide effects, this is the dominant effect on the phase change in fiber devices with temperature, and also in the temperature tuning of fiber Bragg gratings.

If the fiber has a strong plastic cladding (e.g., 1 mm diameter nylon), the temperature sensitivity will go up by an order of magnitude due to the huge CTE of the plastic straining the fiber.

[†]Stretching things makes their dimensions perpendicular to the strain shrink. Poisson's ratio μ tells how big the relative shrinkage is. For something like rubber or chewing gum, where the total volume stays constant, $(1 + \epsilon)(1 - \mu\epsilon)^2 = 1 + O(\epsilon^2)$, so $\mu = 0.5$. Glass has $\mu \approx 0.4$, and most metals, 0.33.



$$E = \sum_{n=0} t_1 t_2 (r_1 r_2)^n e^{i2n\Delta\phi}$$

Figure 8.4. Demodulation of laser noise by etalon fringes and double Rayleigh scatter: this is generally what limits the SNR of fiber measurements.

8.5.13 Bad Company: Fibers and Laser Noise

We saw in Section 2.5.3 that multiple reflections can cause serious intensity noise. As shown in Figure 8.4, even a short length of fiber can do it, and so can multiple Rayleigh scattering even in a perfectly terminated fiber. This coherence effect is one of the main SNR limitations of fiber sensors.[†]

If the scatterers are localized (e.g., facet reflections), you get fringes in the spectrum from localized mirrors. Even tiny frequency fluctuations get turned into huge amplitude variations due to the very steep slope of the fringes with frequency.

In the double Rayleigh case, the phase of all the little scatterers is random, so the scattering contributions add in power. Either way, for fibers longer than $1/(n\Delta\nu)$, you basically get the whole interference term of the scattered light turning into noise with a frequency spectrum that is the autocorrelation of the source spectrum. There's the usual 3 dB reduction in intensity noise, as in Section 13.6.9, since some of the interference term remains phase noise instead of all of it becoming amplitude noise. As we saw in Section 2.5.3, this effect can become dominant with surprisingly short path differences.

The fringe phase is strongly dependent on fiber length and stress birefringence, so two fibers fed from the same source will produce a lot of uncorrelated noise. This effect makes double-beam noise reduction systems such as laser noise cancelers (see Section 10.8.6) far less effective with fibers.

8.6 WORKING WITH FIBERS

8.6.1 Getting Light In and Out

Getting light out of a fiber is relatively simple, because the light is already coming out on its own; you just stick the facet at the back focus of a collimating lens, and you're done. The NA of the fiber makes $10\times$ to $40\times$ microscope objectives a good choice; if you want good beam quality, use single-mode fiber, a good lens, and a good solid

[†]Amnon Yariv et al., Signal to noise considerations in fiber links with periodic or distributed optical amplification. *Opt. Lett.* **15**, 1064 (1990).




Figure 8.5. Coupling to fibers.

mount—not some big six-axis stage. With multimode fiber, there’s no avoiding an ugly beam, so don’t worry about the lens quality so much.

If you just want to dump the light, or shove it into a photodiode, you can dunk the fiber end into some index matching material (oil, gel, or wax). If you get the index right, the back reflection will be very small, and virtually none of the escaped light will bounce back into the fiber—the low étendue works in our favor here.

Getting light into the fiber in the first place is much more difficult, because you don’t have the light itself to guide you (see Figure 8.5.). The first thing you need is to believe in reciprocity. Reciprocity says that an antenna used for receiving has the same spatial pattern and the same gain as when it’s transmitting, or that a beam coupling into a fiber has the same loss that you’d see if it were coming out of the fiber and being spatial-filtered to look just like the incoming beam.

Thus it is necessary to mode-match to get the best coupling efficiency. Because the beam coming out of the fiber is nearly Gaussian, with an NA specified by the manufacturer, you can just use the same NA on the launching side, and be sure of the best coupling.

If your coupling efficiency is below 80%, you haven’t got it quite right; measure the NA of the output of the fiber (remember we’re talking about the $1/e^2$ diameter), and use that on the input side. Your coupling efficiency should improve to the 80–90% range. The coupling efficiency as a function of mode field diameter mismatch is roughly

$$\eta_c \approx \frac{2w_1w_2}{w_1^2 + w_2^2}. \quad (8.28)$$

Focus error and aberrations do to coupling efficiency just what they’d do to an interferometer, and for the same reasons. There are formulas available to calculate how close they should be, but basically you need to be within, say, 0.1 Rayleigh ranges of perfect focus over temperature. This isn’t particularly onerous, since the Rayleigh range is tens of microns.

The major aberration we have to deal with is astigmatism in diode lasers, and even this isn’t normally such a big problem since the diode’s beam is oblong and we’re only using the central bit, where the aberration hasn’t had time to get too large; the coupling efficiency is dominated by the shape mismatch unless you circularize the beam. There are formulas for calculating all these things, but nobody ever uses them in real life.

Aside: A Slightly More Rigorous Look at Reciprocity. An electromagnetic definition of a reciprocal component is that its behavior is unchanged when we replace t with $-t$, \mathbf{E} with \mathbf{E}^* , and \mathbf{H} with $-\mathbf{H}^*$. Consider shining a laser beam through a component, and putting a lossless phase-conjugate mirror after the component, so as to send the beam back through it the other way. The component is reciprocal if and only if the two-pass beam is a phase-conjugated replica of the incoming beam, independent of polarization. Lossless lenses, mirrors, wave plates, beamsplitters, and fibers are reciprocal; attenuators, polarizers, and Faraday rotators are not. Loss can be put in by hand, so we often loosely think of attenuators and polarizers as sort-of reciprocal, though a specific term for that would be useful. Hand-waving reciprocity is a very useful working concept.

8.6.2 Launching into Fibers in the Lab

Those nice fiber collimators that Thor Labs and others sell are great for coming out of fiber, but much less good for going in, because you can't see what you're doing during alignment very well. If you don't mind searching blindly for some time, you can put one of these in a two-axis tilt mount and twiddle till you see something. Alternatively, if yours are connectorized, you can start with a multimode fiber patch cord running into an optical power meter. Find the points where the intensity falls off, and then adjust so you're halfway in between. Then when you put your single-mode patch cord on, you'll see some light right away. For visible light, you don't need the power meter—provided the total power is low enough to be eye-safe, just look at the end of the fiber.

Sometimes nobody makes collimators suitable for our wavelength or beam diameter. The wavelength problem is due to chromatic aberration in the coupler lens, which we can fix, and mistuned coatings, which we can't. The power of lenses generally increases toward shorter wavelength, so if your coupler's center wavelength is too short (e.g., using a 1310 nm collimator at 1500 nm), try a weak convex lens in front of it, and if it's too long, try a weak concave one.

For the fully manual approach, the first thing to start with is a 20 \times , 0.4 NA microscope objective. That's far higher than the NA of most fibers, but then most laser beams are far smaller than the pupil of the microscope lens, so the actual NA you wind up with is closer to 0.2 in most cases anyway. Using a 10 \times , 0.2 NA will require filling the whole pupil—5 mm in diameter or more—and doesn't leave any elbow room for the wings of the Gaussian. Here are the steps (refer to Figure 8.6).

1. If the laser is visible, measure its power and adjust it with an ND filter until it's well below the Class II limit (200 μW is a good number); the most sensitive way to do the early aiming is by looking into the fiber, and it's nice to not hurt ourselves in the process. (This applies in fault conditions too; for example, you bump the power control knob or the monitor photodiode lead comes loose. If you don't have a really trustworthy power meter, don't look into the fiber.)
2. Get the laser beam running horizontally, parallel to one edge of the optical table, and at a convenient height (about 150 mm). Mark the wall where it hits.
3. Strip and cleave the fiber to get a clean end with 1 or 2 cm of stripped length, and have a good look at it with a fiber microscope. Put it in a holder, either a ceramic V-groove collet or a milled inside corner, and clamp or wax it in place.
4. Put the fiber on a three-axis mount with locking screws, and check with a ruler that the axis of the fiber chuck is horizontal and that both it and the z translator




Figure 8.6. Launching light into a fiber.

- are running parallel to the table edge ($\pm 1^\circ$ or so is OK). Check that the light beam hits the fiber end.
5. Put a $20\times$ microscope objective on a sturdy mount, and move it until the center of the output beam coincides with the spot on the wall. Look at the shadow of the fiber in the light beam when it's backed off—the shadow should stay centered and not walk off sideways as z is adjusted.
 6. Put an optical flat (or even a microscope slide, in a pinch—AR coated, preferably) on an ordinary two-tilt mirror mount between the laser and the objective. (This is a vernier adjustment, because the amount of tilt you can get isn't too big, and since it only tilts the output beam and doesn't translate it, it should be orthogonal to the xyz adjustments.)
 7. Watch the way the light reflects from the fiber facet as you adjust z to get it centered, a millimeter or so past focus. Light should be visible coming from the fiber.
 8. Twiddle the xyz knobs to maximize the light coming out of the fiber end. It probably won't be too impressive yet because we haven't adjusted the angle carefully.
 9. Adjust the optical flat tilts to maximize the transmitted power. Iterate the xyz and angular adjustments until the coupling is maximized.

If you're using ordinary SM fiber, you're done. It's a bit tougher with PM fiber, because we have to keep the polarization pure going in, and we don't in general know the orientation of either end of the fiber. For PM fiber, add these steps.

PM Fiber. Before adjusting,

- 4b. Make sure the polarization going in is really linear.
- 4c. Put a rotatable analyzer at the output of the fiber.

- 5b. Use a zero-order quartz $\lambda/2$ plate instead of the flat; it's thick enough to do both jobs if we're careful. You can use both if you like, or if your wave plate is too skinny to displace the beam much without a huge tilt.
9. Iterate rotating the analyzer to minimize the transmitted light and rotating the wave plate to minimize it some more.
10. Jiggle the fiber and look at how much the intensity jiggles. Iterate some more until the jiggle is as small as it's going to get.
11. If you're using a diode laser, you can use minimum detected noise on a photodiode as your criterion for polarization twiddling instead of minimum transmission.
12. Mark the axis orientation of the fiber at both ends for next time.

8.6.3 Waveguide-to-Waveguide Coupling

The easy way to couple two fibers is to use a connector splice. To avoid the resulting broad etalon fringes, you can use microspheres, little spherical lenses that you can put between two fibers. This isn't too easy to do manually. GRIN rods are perhaps a bit more useful; a GRIN rod of $\frac{1}{4}$ period or a bit more will couple two fibers together very nicely, and you can increase the NA by going closer to half a period, at the expense of working distance. Stick the GRIN rod in place, position the first fiber so that light comes out with the same NA that it went in (easy to see on a white card a little way away), and then go through the steps in the previous section.

Connecting fibers to diode lasers and integrated optics devices is usually done in one of four ways: microspheres, lensed (i.e., domed) fiber ends, GRIN rods, or proximity (butting the fiber right up to the device). GRIN rods and proximity are the two easiest ways if you're rolling your own.

8.6.4 Connecting Single-Mode to Multimode Fiber

You can easily send light from a single-mode fiber into a multimode one, but not the other way. In general, trying to get light from multimode into single mode will cost you about 20 dB in signal, together with very large (order-1) variations due to bending and temperature changes. This is because your average multimode fiber has about 50–100 modes, all of which are orthogonal. Because of their orthogonality, they can't all couple well into your fiber mode—in fact, if all N modes are illuminated equally with incoherent phases, you can get at most $1/N$ of the light into your single-mode fiber.

8.6.5 Fibers and Pulses

One situation where fiber really does behave a lot like coaxial cable is when you're using short pulses with a very low repetition rate. A single pulse entering a complicated fiber system will generate all sorts of smaller pulses due to reflections, which will rattle round the system for some time. On the other hand, if the pulse is short compared with the delays between different reflections, and the rep rate is low enough for them all to die away before the next pulse, things can be pretty well behaved. The author has a 20 Hz, 20 ps tunable laser system, which works fine with fibers as long as the reflections are suppressed by time gating. Of course, the polarization instability problems remain.

8.6.6 Mounting Fibers in Instruments

In instruments, we have another set of problems, caused by dirt, vibration, shock, and temperature drift, and mechanical instabilities due to lubricant flow and stick–slip energy storage (you know, like the San Andreas Fault). It isn't a problem inside the fiber, obviously, but it is for launching and collimating. The easiest and most expensive way round this is to start with pigtailed components and just stick them together with connectors. You can get armored fiber patch cords, whose steel gooseneck jackets enforce a minimum bend radius and can be bolted down, which helps a lot. The corrugated interior causes microbending, though, so they aren't as useful with multimode fiber, especially in the UV.

One thing to remember is that connectors for weird fiber may be impossible to find (and the price may wish you hadn't found them if you do), so try to stick with 125 μm cladding diameters which is the standard for communications fiber.

The other thing about using fibers in instruments is that you have to be very careful in choosing a jacket material. Teflon and polyimide are the most common choices at present, and they're both good as long as they're not thick enough to stretch the fiber. We've already talked about a thick plastic jacket straining the fiber with temperature, but there are other possible problems. Nylon and some other plastics are hygroscopic, so that they swell in humid conditions—you don't want that source of instability either.

8.6.7 Connectors

Fiber connectors used to be expensive and very miserable, but now they're cheaper and more mildly miserable. (At least if your fiber has a 125 μm cladding OD—special sizes are much more expensive, (for example, \$40 for an 80 μm ST connector, versus \$8 for 125 μm .) This is of course a jaundiced view, but how else do you describe a splice that is guaranteed to send a few tenths of a percent of the light right back into the source (or 0.01% for angled ones)? Even that 100 ppm is more than enough to make a diode laser misbehave and will cause percent-level etalon fringes.

8.6.8 Splices

The lowest-reflection way to connect two fibers is fusion splicing, where two cleaved fibers are butted and then melted together, but that isn't too convenient. We more often use epoxy splices, which come in kits and are easy to use. Fusion splices have reflections of 1 ppm or so, which is a big help in reducing instability. There are also temporary splices, mostly based on very accurate V-grooves in ceramic materials, with a variety of clamping strategies. You stick the fibers in, usually with a little index gel, and engage the clamp. Expect 0.1 dB coupling loss for a fusion splice and 0.3–0.5 dB for an epoxy splice.

Remember the egregious polarizing beamsplitter cube in Example 4.1, whose reflections were 1%, and whose length was only 25 mm—it had a temperature coefficient of transmission of 12%/°C. A fiber patch cord is even worse, at least with highly coherent sources, and you even get strong etalon fringes between the two fiber ends in the connector itself. As usual, you get strong fringes with single-mode fiber because there's no misalignment to help suppress them.

You'll probably use FC or ST connectors, in either normal or physical contact (FC/PC, ST/PC) varieties. The PC types use fibers with a slightly convex polish on the ends; they

have lower loss and back-reflections, but they are easily destroyed by dust on an end facet. You can also get angled physical contact (APC) connectors, with a nominal 8° facet angle, and these help quite a bit with back-reflections, as long as the reflected wave is all beyond the core/cladding critical angle, so that it all leaks out. APC connectors are unforgiving of angular misalignment, so they come with keys to keep them straight. These keys are the same kind used in PM fiber connectors, which can make life interesting if you're not careful.

You can't mate an APC connector to another kind, because the light comes out at a 4° angle to the fiber axis, and because you can't get the ends close enough together on account of the angle. That 4° angle makes trouble in coupling light in and out, too.

Connector manufacturers sell inexpensive kits that help a lot with getting the end facet polish right, so by all means buy one. You can clean fiber facets using a Q-tip and some very clean ethanol, or (for quick and dirty use) some scotch tape.

8.6.9 Expanded-Beam Connectors

At the price of a couple of decibels' loss per pair, you can use connectors with integral collimating lenses. They are tolerant of dirt and other sorts of abuse, and you can shoot the beam some distance in mid-air if you need to, which is a plus in instruments. The lenses aren't good enough to make a decent beam, unfortunately.

8.6.10 Cleaving Fibers

Cleaving fibers is easy but takes a bit of practice; most of us are much too rough at first. Semiautomatic fiber cleavers exist, but careful hand work is just as good. Strip the fiber back a couple of inches, wet it, and nick it ever so gently with a V-point diamond scriber in the middle of the stripped section (don't buy the cheap carbide scribers). Bend the fiber gently at the nick until it breaks.

Another technique is to strip it back a bit further, wet and nick it as before, and bend it back in a hairpin loop between thumb and forefinger. Pull gently on one end so that the loop slowly gets smaller until it breaks.

Bad cleaves are usually due to too big a nick or to pulling the loop in too fast; hackles[†] form when the fracture propagates faster than the speed of sound in the glass. Big fibers are much harder to cleave than little ones. Twisting the fiber makes an angled cut.

8.7 FIBER DEVICES

It's worth emphasizing the distinction between all-fiber devices, which use some intrinsic property of the fiber, and fiber-coupled devices, where the fiber is grafted onto a bulk optic or integrated optic device. Directional couplers, Bragg gratings, and polarization compensators are all-fiber, but Pockels cells, isolators, and detectors are merely fiber coupled. All-fiber devices often have a cost advantage, whereas fiber-coupled ones are more expensive than bulk. Bulk devices are covered elsewhere, so here we'll stick with the all-fiber ones.

[†]A hackle is a bump or dip where a bit of glass has remained attached to the wrong side of the cleave.

8.7.1 Fiber Couplers

By bringing the cores of two single-mode fibers close together, the mode fields overlap, leading to the coupled-mode behavior we saw in Section 8.3.3. Phase matching makes such *evanescent wave couplers* highly directional. If the beat length is longer than the interaction zone, nearly 100% coupling can be achieved, but we more often want 3 dB, which gives maximum fringe visibility in interferometers. Couplers are usually made by fusing two fibers together and stretching the fused region until its diameter is slightly less than that of a single fiber, as shown in Figure 8.7a, which shows a 2×2 tapered coupler. By conservation of étendue, the NA must increase as the core diameters drop, and the high angle light is no longer guided. It therefore goes off into cladding modes. The reverse process occurs as the core broadens out again, corralling most of the light into the cores again. Because most of the light is in cladding modes at some point, the outside of the cladding must be clean and untouched; thus tapered couplers are usually suspended in air inside a case.

All these devices are inherently $2N$ -port, so in combining two beams with a 3 dB coupler, you're going to lose half of each one out the other coupled port. This is easy to see thermodynamically, because otherwise you could feed 100 fibers from a low temperature source, combine them, and use the result to heat a higher temperature sink. The four-port character is not necessarily a disadvantage, but you do have to be aware of it.

Coupler figures of merit are the coupling ratio, directivity, and excess loss. For a four-port (2×2 , i.e., two fibers in, two fibers out) coupler, if we feed P_1 into port 1, the coupling ratio is $P_4/(P_4 + P_2)$, the directivity is P_3/P_1 , and the excess loss is $P_1 - (P_4 + P_2)$.

Provided the taper is sufficiently gentle, the reflections are small, and the excess loss is well below 1 dB; large-core multimode splitters can have 1.5 dB of excess loss, and some single-mode ones reach 0.1 dB. The directivity is usually -50 to -60 dB (optical), and back-reflections are relatively low too, around -50 dB, because the fiber stays more or less intact. If you need adjustability, you can imbed fibers in epoxy blocks and lap

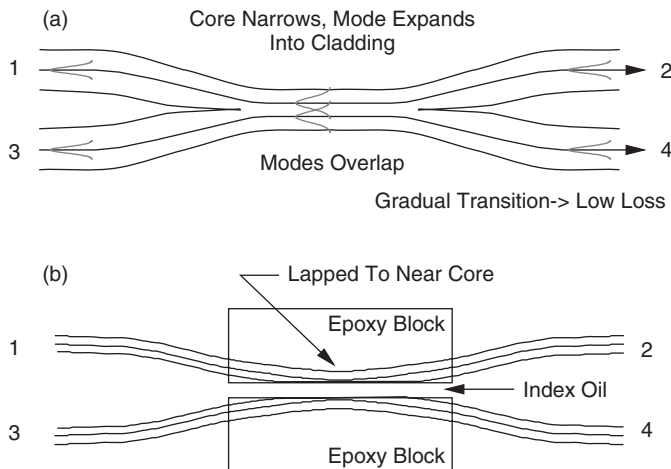


Figure 8.7. Fiber couplers: (a) tapered fused coupler and (b) lapped coupler.

them until the mode fields reach the surface; putting two such fibers together makes an adjustable coupler.

8.7.2 Fiber Gratings

A single-mode fiber is an ideal application for Bragg gratings. For a given wavelength, we know exactly what k_z is, and it's sufficient to put in a grating with $\mathbf{k}_G = 2k_z\hat{\mathbf{z}}$ to make a very effective mirror; in wavelength terms, the Bragg wavelength λ_B is

$$\lambda_B = 2n\lambda_G. \quad (8.29)$$

As with plane gratings, the resolving power of a weak Bragg grating ($R < 10\%$) is $R = mN$, the number of lines in the grating times the diffraction order (providing, of course, that the grating is weak enough that the far end gets illuminated). A weak Bragg grating exhibits a strong, narrow spectral dip, which is insensitive to polarization (see Figure 8.8).

You can use fiber Bragg gratings like mirrors, for example, to make fiber lasers or interferometers. You can use them as filters, say, to confine the pump light in a fiber amplifier or laser. And you can use them as fiber sensors, because stretching the fiber or changing its index will cause the reflection peak to move.

In the beginning, people made these fiber gratings by etching the cladding, but nowadays they're made by the photorefractive effect: strong UV (248 nm) irradiation causes

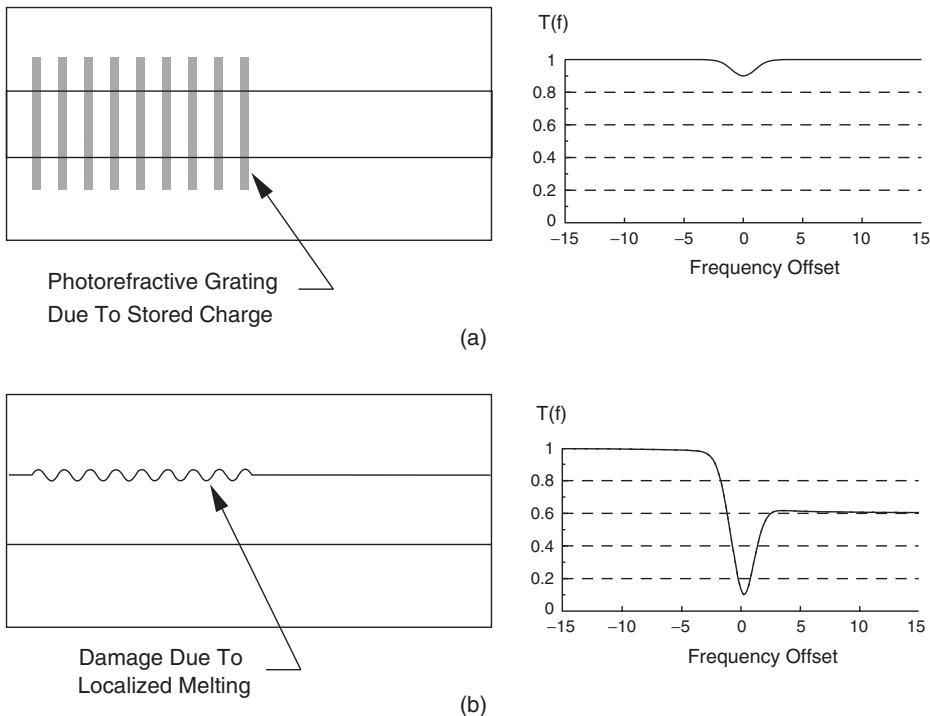


Figure 8.8. Fiber Bragg grating: structure and spectral behavior in reflection and transmission.

permanent changes in the index of the glass. Ordinary *Type I* fiber gratings (Figure 8.8a) have a small index change that extends more or less uniformly through the core, which makes the coupling to cladding modes very weak, normally a very desirable property.

A short-pulse, high rep rate excimer laser shining through a mask writes gratings in a fiber fed past from reel to reel. This is potentially very inexpensive in large quantities, although if you buy one unit, it isn't much different from a 25 mm square plane grating at this writing (\$500).

The peak reflectance of a Type I fiber Bragg grating is typically below 10%, though the newer commercial ones are long enough to reach 90% or more. The tuning of the grating depends on strain and temperature; the interaction between the two is slight, so the shift of the peak can be predicted from (8.24) and (8.27).

8.7.3 Type II Gratings

It is also possible to make extremely strong Bragg gratings, in which the grating lines are conceptually more like a stack of dielectric films; these Type II gratings are made by actually melting the core-cladding boundary with much higher laser fluence. They are asymmetrical and hence couple light into cladding modes very strongly for wavelengths shorter than λ_B ; the loss can be as much as 90% (Figure 8.8b).

Like dielectric stacks, these gratings can achieve very high, flat reflectivities over a moderately broad band—99.8% has been reported.

8.7.4 Fiber Amplifiers

Fibers doped with rare-earth ions are widely used as amplifiers; erbium at 850, 990, and 1550 nm, neodymium at 1.06 and 1.32 μm , and holmium at 1.38 μm . They boost the signal level without needing to be converted to electrical signals and back again, as in a repeater. The spontaneous emission increases the noise, so you can't use too many in a row. The spontaneous emission has shot noise, and it also leads to severe *beat noise*, caused by its interference with the desired signal. This is similar to the coherence fluctuations and double Rayleigh scattering problems, and shows up because there's no spatial averaging to get rid of it as there is in bulk optics. You can't do much about signal/spontaneous emission noise, but a narrow filter will mostly eliminate spontaneous/spontaneous beats, which can be much worse.

Erbium-doped fiber amplifiers (EDFAs) are pumped with special diode lasers at 980 or 1490 nm; the pump light is fed in via a coupler and stopped with a Bragg grating, or a coated facet, or sometimes just allowed to exit along the output fiber, where it will ultimately be attenuated by the fiber loss.

Don't be tempted to see fiber amplifiers as equivalent to packaged RF amplifiers in coax systems, because the resulting systems are not nearly as stable; a garden-variety RF amplifier gives you 20 or 30 dB of isolation from output to input, whereas a fiber amplifier gives none whatever. In fact, it gives less than 0 dB isolation, since the gain is not intrinsically directional—it amplifies the light going the other way just the same amount. RF folks would say that a 20 dB EDFA has *minus* 20 dB of isolation, which would make them unhappy, and for good reason. Part of the skill of designing fancy fiber systems is knowing how to use the minimum number of expensive isolators to get good stability.

8.7.5 Fiber Lasers

By adding Bragg mirrors tuned to the signal wavelength, putting mirror coatings on the facets, or by turning it into a ring resonator with a coupler, a fiber amplifier can be made into a solid state fiber laser. The simplicity of the idea is its main advantage, but in principle it can be of great benefit in long path difference interferometers such as Fabry–Perot types, because the linewidth can be made narrower than the typical 10–100 MHz of a single frequency diode, and the coherence length correspondingly longer.

8.7.6 Fiber Polarizers

Polarizers are an example of something that’s hard to do in fiber but easy in bulk. Persuading a fiber to actually absorb the unwanted polarization requires giving it anisotropic loss, for example, by lapping the cladding down to get to the mode field, and then putting on a metal coating to absorb the TM polarization. The TM mode excites surface plasmon waves in the metal, which are enormously lossy. Metal fiber polarizers can have open/shut ratios of 50 dB and losses of 0.5 dB.

You can get polarizing (PZ) fiber, in two types. One works by the single-mode equivalent of frustrated TIR; if the evanescent field falls off at slightly different rates for s and p polarizations in the cladding, then a localized cladding region with high absorption or low refractive index can cause one polarization to be attenuated differently from the other. The other kind works by the difference in the bend edge in the two polarizations; a few meters of carefully aligned coiled fiber sends most of the more weakly guided mode off into the cladding. Usually we just put a bulk polarizer before or after the fiber. Walkoff plates are especially popular for this because one beam is undeviated, which makes them easy to make and to align.

8.7.7 Modulators

We’ve already encountered integrated optic Pockels cells, which are the most common modulators used with fibers. Fiber phase modulators are often made by wrapping many turns around a large piezoelectric tube, but these *fiber stretchers* can’t go much faster than 50 kHz. They also have a lot of bend birefringence unless you anneal the fiber in the coiled state (which is hard since the piezo won’t stand the 800°C annealing temperature).

8.7.8 Switches

The most common type of intensity modulator or switch is the integrated optic Mach–Zehnder type, in which a phase modulator in one arm of a waveguide Mach–Zehnder interferometer causes a bright or dark fringe to occur at the output waveguide. Interferometers are always four-port devices, even though these ones don’t look like it; the third and fourth ports are radiation into and out of the substrate when the two waves are out of phase. You can also do it by putting fibers on piezos, and moving them in and out of alignment. People have built N -way switches in that fashion, but it’s a hard way to make a living.

8.7.9 Isolators

Fiber isolators are also extrinsic devices; the two most common are Faraday and acousto-optic, which we’re familiar with from bulk optics. Faraday isolators rely on polarizers

to get their directional selectivity, so use PM fiber or mount the isolator right next to the laser, where the polarization is very stable. Faraday rotator mirrors provide some isolation as well, providing a polarizing beamsplitter is used to separate out the returned light.

8.8 DIODE LASERS AND FIBER OPTICS

Both of these technologies are wonderfully useful, but they require much more engineering than one at first expects. Neither is a quick fix for an implementation problem. Especially seductive is the ease of hacking up a connectorized fiber system: just twist the connectors together, and you're done.

The parallel between this and doing signal processing with a table covered in boxes and RG-58/U patch cords is real but not perfect; the main differences are that the length of the optical fiber is at least millions of wavelengths, whereas the coax is usually less than 1, and that electronic amplifiers provide isolation, whereas almost all optical components are inherently bidirectional. An optical fiber setup is more closely analogous to a whole cable TV or telecommunications cable system, with miles and miles of coax. Stray cable reflections can cause severe ghost images in TV reception, much the way etalon fringes foul up measurements.

Even there, the analogy breaks down somewhat; the reflections from the cable network don't cause instability or frequency pulling of the sources, because they're well isolated; diode lasers are much more susceptible. And besides, coax doesn't have polarization or mode problems, or lose signal when you bend it. There's no coax network long enough to get coherence fluctuations from a typical oscillator with a 1 Hz linewidth, either.

Diode lasers and fiber optics work better apart than together, unless you're prepared to spend some real dough (\$1000 or more) for decent isolators, or to destroy the coherence by large UHF modulation, and even then, you have the coherence fluctuation problem we've already discussed.

8.9 FIBER OPTIC SENSORS

It is important to distinguish a fiber sensor from a sensor with fibers in it. There are right and wrong reasons for putting fiber optics in your setup. The right reason is that there are unique physical measurements you can do with the fiber, serious safety concerns, or compelling cost advantages, and eventually lots of similar sensors are going to be made, so that it's worth putting in the engineering time to get it really right. Borderline reasons are that you need a really good spatial filter, and pinholes aren't enough, or that the system needs 100 kV of voltage isolation. The wrong reason is that your bulk optical system is inconvenient to align, and it would be nice to use fibers to steer the light around to where it's needed. If you're building sensors for others to use, you'll quite likely spend a lot of late nights in the lab, wishing you hadn't done that. Where conventional sensors (optical or not) can do the job, you should do it that way almost always.

Fiber sensors can do some really great and unique things, but they're just one of many tools for the instrument designer. If you really like fibers, and want to use them everywhere, you face the same type of problem as the man in the proverb who only has a hammer—everything looks like a nail.

8.9.1 Sensitivity

The basic trade-off in fiber sensors is sensitivity versus convenience and low cost. It's often very fruitful to go the low cost route, since the enormous dynamic range of optical measurements is not always needed, and the market potential may be enormous. Do realize, though, that you are making that trade-off, so that if it isn't low cost, there's often no point.

8.9.2 Stabilization Strategy

If we are going to make a fancy fiber sensor, we have to be able to predict its performance. Elsewhere, we've been asking "What's the SNR going to be?" and calculating it from first principles and measured device parameters. With fiber sensors, we're not that fortunate, because there are two prior questions. The first is: "How is it going to be stabilized?" Every sensitive fiber sensor is seriously unstable unless you figure out a way to stabilize it, and this instability is the primary bogey in fiber sensor work. Scale factors drift, operating points go all over the place, and all on a time scale of seconds or minutes—unless you figure out how to avoid it. Some sort of common-mode rejection scheme is usually needed, e.g. a fiber strain gauge with an unstrained reference fiber in the same package, two wavelengths, or two modes. Some of these schemes work a lot better than others.

8.9.3 Handling Excess Noise

The next question is: "How do we get rid of the excess noise?" After stability, the worst problem in laser-based fiber sensors is their huge technical noise, caused by demodulation of source phase noise by etalon fringes and scattered light. Only after dealing with these two issues do we get to the point of worrying about our SNR in the way we do elsewhere. Fiber measurements are intrinsically squishier than bulk optical ones (see Section 11.4.2).

8.9.4 Source Drift

LEDs drift to longer wavelengths with increasing T , by as much as several parts in $10^4/^\circ\text{C}$. Diode lasers drift less (at least between mode jumps) because they are constrained by their resonators. Any fiber sensor that is sensitive to changes in the source spectrum will have to neutralize these effects somehow. The easiest way is to temperature-control the source, but temperature compensation or normalization by an independent source spectrum measurement will often work too.

8.10 INTENSITY SENSORS

Most mainstream fiber sensors at present are based on changes in received intensity, at or near DC. There are many kinds, but only a few physical principles: microbending loss from pressing against corrugated plates, evanescent coupling between nearby fiber cores, and misalignment.

Intensity sensors more or less ignore phase and polarization, which makes them fairly insensitive but pretty trouble-free (for fiber devices). The limitations on their sensitivity




Figure 8.9. Microbending sensors.

come from the smallness of the signal and their vulnerability to interference from other losses, e.g. microbending outside the sensor and source drift and noise.

Intensity measurements can help us normalize the output of fiber interferometers. Sampling the beams before recombining them gives us the intensity information separately, which can be subtracted or divided out.

8.10.1 Microbend Sensors

Microbending increases loss, with the attenuation being very roughly linear in the curvature. Microbending sensors exploit this effect in a variety of ways (see Figure 8.9).

They aren't very sensitive. Most cause only a few percent change in transmission, full scale, and they tend to be fairly badly nonlinear; oddly shaped deviations of tens of percent from the best straight line are common, and that isn't too confidence-inspiring since we often don't know its origin. Microbend vibration sensors are a good technology—the AC measurement avoids most of these troubles if you're careful.

8.10.2 Fiber Pyrometers

Black body emission follows the Stefan–Boltzmann equation, so an IR fiber with a decent black body on one end and a thermal-IR detector makes a nice thermometer for high temperature applications (up to 800 °C).

8.10.3 Fluorescence Sensors

The upper-state lifetime of many fluorescence transitions goes down as T goes up. Because the decay is exponential in time, measuring the decay time normalizes out the intensity shifts. A cell full of fluorophore solution or a fluorescent coating on the fiber facet are the common ways of doing it.

An unselective detector followed by a log amp and a bandpass differentiator is all you need; this is the same basic idea as cavity ring-down spectroscopy. The most popular fiber for this is Nd:glass, whose lifetime varies from 210 μs at 0 °C to 170 μs at 250 °C, which is quite a convenient range. The repeatability of an individual, undisturbed sensor is about 0.5°, but the exact dependence of the lifetime on temperature seems to vary a lot from sensor to sensor (as much as $\pm 3^\circ$ nonlinear shift from sensor to sensor) so that for high accuracy, individual calibration over the full range appears necessary.

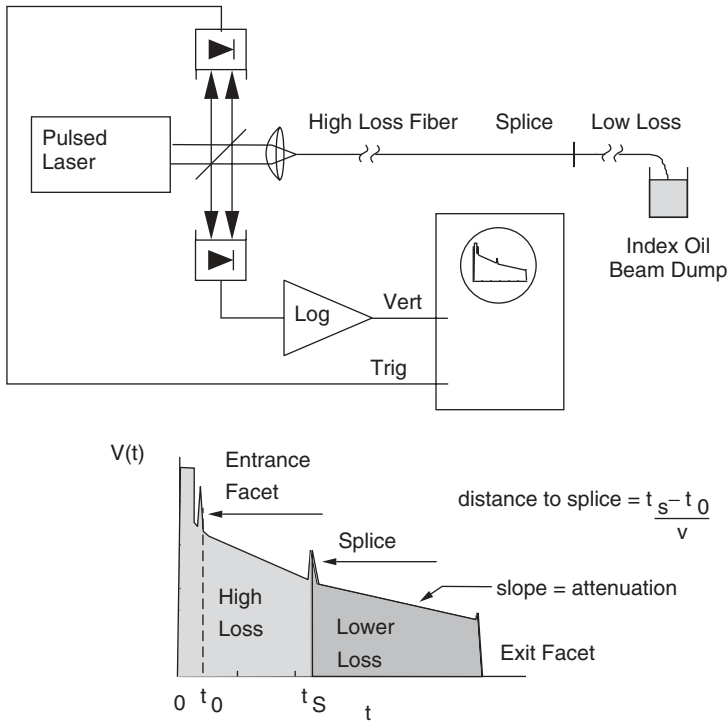


Figure 8.10. Optical time-domain reflectometry.

Fluorescence sensors based on fluorophores immobilized in fiber coatings are also used for oxygen detection; O_2 diffusing into the fiber coating quenches the fluorescence. (Of course, normal fluorescence measurements are often done using fibers as light pipes, but that isn't a fiber sensor in the sense used here.)

8.10.4 Optical Time-Domain Reflectometry

Optical time-domain reflectometry (OTDR) works a lot like radar: you send a tall, narrow pulse down the fiber, and look at what comes back, as a function of time delay. The result is a plot of $\log(r)$ versus t , that is, distance. As shown in Figure 8.10, we get a constant Rayleigh scatter signal, which decays at twice the fiber loss rate, and reflection spikes from splices and other discontinuities. A steeply dropping slope means excess loss. Chemical sensors, microbend sensors (e.g., fiber with a helically wrapped member inside the armor), and even porous cladding sensors for water can all be interrogated by OTDR. Polarization and phase are more or less ignored.

8.11 SPECTRALLY ENCODED SENSORS

8.11.1 Fiber Bragg Grating Sensors

Bragg grating sensors work by illuminating the fiber with a broadband source such as an LED, and using a spectrometer to watch the absorption peak move when the fiber

is stretched[†] or heated, in accordance with (8.24) and (8.27). Stretching the fiber moves the absorption peak by a combination of the change in length, the piezo-optic effect, and the change in the waveguide diameter through Poisson's ratio. The resolution of the measurement depends on the quality of the spectrometer and the length of the grating. This *spectral encoding* makes for a nice stable sensor, at least once the temperature dependence is removed. Fiber Bragg gratings escape nearly all the problems of other fiber sensors; they are insensitive to polarization, phase instability, excess loss, and etalon fringes (if you use fiber-coupled LEDs, you even escape alignment). They do require an external spectrometer, and the higher its resolution, the higher the sensitivity of the sensor. Because of the low étendue, you don't get that much light, but on the other hand you don't need much, because these sensors are typically measuring mechanical and thermal things, which are slow.

By putting several gratings (as many as 10 or 20 of different periods) into the same fiber at different positions, a single fiber can do simultaneous strain measurements at many positions. This idea is used especially in smart structures, which take advantage of Bragg grating sensors' ability to make measurements in otherwise totally inaccessible locations, such as the interior of reinforced concrete bridge abutments.

The strain sensitivity of this technique isn't bad; from Section 8.5.11, it's about

$$\frac{1}{\lambda_B} \frac{\partial \lambda_B}{\partial \epsilon} = 0.78, \quad (8.30)$$

so that a grating stretched by 0.1% of its length (1000 μ strain) changes its wavelength by +780 ppm. The total allowed strain isn't very big, 1000–2000 μ strain, so a sharply peaked wavelength response is obviously important. People have tried using interferometers to read out Bragg gratings, but that throws away the main advantage of the technique: its insensitivity to everything but strain and temperature. One interesting approach is to use a Mach–Zehnder delay discriminator to measure the average wavelength of the light reflected from the grating; although this is obviously highly sensitive to the exact line-shape of the reflection, it does get around the resolution limit imposed by spectrometer pixels. Fitting a grating instrument with a split detector or a lateral effect cell would be another way to do this.

The accuracy of fiber strain gauges isn't bad, a percent or so. A more serious problem is temperature drift; the tempo of λ_B is

$$\frac{1}{\lambda_B} \frac{\partial \lambda_B}{\partial T} = +6.7 \text{ ppm}/^\circ\text{C}, \quad (8.31)$$

so that a 1 $^\circ\text{C}$ change gives the same signal as a 9 μ strain stretch—which is on the order of 1% of full scale. To leading order, the two effects are both linear in λ , so separating them is difficult; the best method seems to be the one used in resistive strain gauges, that is, comparison with an unstrained grating at the same temperature. Various two-wavelength schemes have been tried, which use the fiber dispersion to make the strain and temperature coefficients different. These all fail in practice, either because the wavelengths are so far apart that the fiber is no longer single mode, or because the coefficients don't shift enough to give a well-conditioned measurement.

[†]Of course, you can do it with a tunable laser and a single detector, but then all the usual fiber problems reappear, and sufficiently widely tunable lasers don't grow on trees either.

Current laboratory accuracies are about $1 \mu\text{strain}$ or $0.1 \text{ }^\circ\text{C}$, but field accuracies are less because it's hard to make a good, cheap readout system. The main challenge in low resolution field systems is in maintaining a good enough bond between the strained element and the fiber that strain is transferred accurately. Quartz is very rigid stuff; glue isn't, especially not glue that's gotten hot or wet.

8.11.2 Extrinsic Fabry–Perot Sensors

Spectral encoding is also the idea behind many types of extrinsic Fabry–Perot sensors, for example, diaphragm-type pressure sensors, which look at the interference of light reflected from the fiber end and from a nearby diaphragm. Their low finesse and narrow spacing make their fringes broad enough to interrogate with white light or LEDs.

Other examples are thermometers based on refractive index drift in silicon.

8.11.3 Other Strain Sensors

Extrinsic piezo-optic strain sensors are sometimes used with fiber coupling, but this technology seems not to be becoming mainstream. Other strain-based sensors include magnetostriction-type magnetic field sensors, in which a strain sensor is bonded to a magnetostrictive element. It's pretty easy to put an unstrained companion on one of these, e.g. by bonding one with indium solder or brazing, and the other with elastomer.

8.11.4 Fiber Bundle Spectrometers

Grating spectrometers and FTIRs are big clunky things, so it's natural to want to couple them with fibers into inaccessible places. It's especially nice to use linear bundles, where the fibers are arranged in a row to fill the entrance slit. This can work OK if you pay extreme attention to the pupil function you're imposing on the measurement. An integrating sphere and a multimode bundle are stable enough for good spectra, but if you simply bounce light from one fiber off a surface and into another fiber, you're just asking for ugly measurement artifacts. You also give up a lot of light unless you use a lot of fibers; a grating spectrometer's étendue is nothing to write home about, but a fiber's is worse. It also tends to be unstable.

Two of the author's colleagues abandoned fiber-bundle spectroscopy after giving it a good run in an *in situ* application, in favor of mounting miniature monochromators right on the instrument head—they couldn't get enough photons otherwise.

One application where fibers actually *increase* the amount of light you get is in astronomical fiber plate spectroscopy, where a computer-controlled milling machine drills holes in a plate just where the stars will be imaged. Sticking multimode fibers into the plate brings all those widely separated sources into one spectrometer—a nice example of putting your étendue where you need it.

8.11.5 Raman Thermometers

The Raman effect is a bilinear mixing between vibrational levels in the fiber material and the optical field, and is similar in character to the acousto-optic interaction we talked about in Section 7.10.4. A photon can emit a phonon, and so be downshifted (Stokes

shift), or it can absorb a phonon and be upshifted (anti-Stokes). The Stokes/anti-Stokes intensity ratio depends on the occupation number of the vibrational states, and hence on the temperature. The ratio changes by a few tenths of percent/°C at a few hundred wave numbers offset, but the signal is weak; an OTDR setup with a good filter and a decent PMT or APD gives a few meters' resolution at 1 °C; by using photon counting and very narrow pulses, you can get 5 °C accuracy with 10 cm spatial resolution.

8.11.6 Band Edge Shift

Absorption and fluorescence spectra are often highly temperature dependent. The band edge of a semiconductor moves toward long wavelength with temperature, while remaining sharp, so it can be used as an extrinsic fiber thermometer. These don't get better than 1 °C, and they're almost as nonlinear as thermistors, but they can work from room temperature to 800 °C.

8.11.7 Colorimetric Sensors

Fiber chemical sensors can be made by cladding a fiber in a microporous sol-gel glass with a chromophore (color-changing chemical) mixed in. The chromophore is immobilized in tiny pores in the glass. The target chemical species diffuses in and combines with the chromophore, changing its absorption spectrum. This trick is used in commercial pH sensors, for instance.

8.12 POLARIMETRIC SENSORS

8.12.1 Faraday Effect Ammeters

Currents produce magnetic fields, which in turn produce Faraday effect polarization shifts in fibers. This underpins all-fiber Faraday rotation ammeters, in which the small Verdet constant of silica ($V = 4.6 \times 10^{-6}$ rad/A) is overcome by using many turns; with N turns wrapped around the conductor, the Faraday shift is

$$\theta = V \cdot N \cdot I. \quad (8.32)$$

This of course works only if the two circular polarizations remain distinct. If they are strongly coupled, as by bend birefringence, the Faraday effect is *quenched*; light in the fast and slow circular polarizations change places several times inside the coil, so that the phase difference tends to cancel out (a bit like twisted pair wiring). Annealing the coiled fiber can help a lot with this, or you can quench the quenching by twisting the fiber as it is coiled. The resulting circular birefringence prevents the linear birefringence from building up enough to couple the fast and slow circular polarizations much; the circular birefringence just gives rise to a fixed phase shift between polarizations, which isn't much of a problem. The fibers are relatively short, so double Rayleigh scatter isn't usually serious. The path difference between the two polarizations is small, so source PM doesn't turn into polarization jitter, and as long as the etalon fringes are well controlled, the FM-AM conversion is thus minor; the predominant noise is source intensity noise, $1/f$ noise in the front end amplifier, and the usual huge amount of low frequency junk due to temperature and bending.

In the scheme of things, nobody's about to use a fiber sensor when a better and cheaper technique such as a current transformer or a series resistor and isolation amp is applicable; thus these sensors are used only inside high voltage transmission systems. This is helpful from the noise point of view, since the current is AC, and therefore so is the measurement.

8.12.2 Birefringent Fiber

The change in birefringence with strain, pressure, and temperature is roughly proportional to the static birefringence, so people have tried using polarimetry with high birefringence fiber to sense pressure and vibration. This works at some level, but the temperature dependence is so strong, and the operating point so unstable, that such techniques appear to have little to recommend them.

8.12.3 Photonic Crystal Fiber

Recently, fibers have been developed whose cross sections are not solid but include arrays of air holes, making them 2D photonic crystals. Their structure gives these holey fibers unusual properties. Some are highly nonlinear, for example, those used in femtosecond supercontinuum comb generation; others have very wide mode fields that are very constant with wavelength, or actually guide the light inside one of the holes, rather than in the glass. They haven't been used much in fiber sensors yet, but that's likely to change soon.

8.13 FIBER INTERFEROMETERS

Fiber interferometers produce an order-unity change in output for a 1 radian phase shift, just as bulk ones do. Most of the usual kinds can be constructed in fiber: Michelson, Mach-Zehnder, Fabry-Perot, and so on.

The big challenge is that fiber interferometers are sensitive to *everything*, especially temperature. This lack of selectivity makes it hard to have much confidence in the data, unless great efforts are made. Temperature gradients are especially serious.

At the root of the problem is that the phase shift due to interfering effects tends to go as the length of the fiber; since it goes into the exponent rather than the scale factor, at some point the signal-to-noise ratio cannot be improved any more by using more fiber.

Fiber interferometers place immense demands on the stability and coherence of the laser. Ordinary diode lasers are not even in the running for this use, because their linewidths are in the tens of megahertz, and their drift, mode hopping, and mode partition noise make them very difficult to use barefoot. External-cavity stabilized diode lasers are a good choice, as are solid state devices such as diode-pumped YAGs.

The instability restricts the dynamic range of fiber interferometers. You can use them over a wide range by fringe counting, or over a narrow one by AC modulation, but you can't easily splice them together into a single wide-range measurement as you can with bulk. Even though the path length may be a factor of 1000 larger in the fiber case, fringe-counting resolution of 10^{-6} m/1 km is far poorer than the 1 Hz shot noise limit of 10^{-14} m/1 m (5 mW HeNe).

8.13.1 Single Mode

If single-mode fiber were really single mode, instead of degenerate two-mode, and that mode were really lossless, building single-mode fiber interferometers would be a piece of cake. Although neither is true, still single-mode fiber is so much superior to multimode for interferometry that it's the only game in town. We therefore make the best of it.[†]

8.13.2 Two Mode

The availability of communications fiber that supports two modes at some readily accessible wavelengths has led people to try using the mode coupling or differential delay or polarization shift between modes as a sensor. This is an admirable example of actually trying to realize the hitherto mythical cost savings due to economies of scale, instead of just talking about them. This works OK if you babysit it, but getting exactly two modes forces us to work in a regime where the second mode is somewhat leaky, and hence sensitive to manufacturing variations in the fiber, as well as to handling, bending, and so on. Higher order modes are not well preserved in splices, and especially not in connectors. It's a tough business. At present, avoid two-mode and multimode interferometric sensors if you want your gizmo to be usable by anyone but you, but stay tuned—if someone figures out how to do it well, this trick could really take off.

8.14 TWO-BEAM FIBER INTERFEROMETERS

By using fiber couplers instead of beamsplitters, it is quite straightforward to make Michelson, Sagnac, and Mach–Zehnder fiber interferometers, as shown in Figure 8.11. There are two basic approaches to using them; directly as sensors, or as delay line discriminators for optical frequency measurements. A lot of blood has been spilled making intrinsic interferometric fiber sensors, and it's just a very hard business to be in. Fiber gyros have been in development for many years and have only recently overcome their many, many second-order problems to become a mainstream technology. That's good if you're cranking out PhD theses, but not so good if you're building sensors for a living. If you're in the latter class, beware.

8.14.1 Mach–Zehnder

Mach–Zehnders are the workhorses, but they're horribly unstable if made of fiber, for all the usual reasons. Integrated optics ones are much better. All-fiber Mach–Zehnders are commonly used as delay line discriminators in high velocity Doppler shift measurements, but require a lot of babysitting due to polarization funnies.

8.14.2 Michelson

Fiber Michelsons are attractive because you can use FRMs, which eliminates the polarization drift and allows passive interrogation—the fringe visibility in a Michelson with decent FRMs normally stays above 95%, which is pretty good. The main residual problems are the few degrees of rapid polarization and phase wobble with (average)

[†]Every mode is its own interferometer, and you can't keep them all in phase.

temperature, and temperature gradients, which cause large amounts of phase drift in the relative phases of the two arms. A fiber Michelson built with FRMs and interrogated with the modulation-generated carrier approach works adequately for AC measurements, where the phase drift and fringe visibility can be taken out in software without screwing up the operating point.

8.14.3 Sagnac

Sagnac interferometers use waves propagating in the forward and reverse directions, more or less fixing the phase instability problem.

The static phase delays in the two directions are the same except for the Faraday effect, Berry's phase, and the *Sagnac effect*, which causes a phase shift in a rotating reference frame. An N -turn coil of cross-sectional area A , turning at angular velocity Ω , produces a phase shift between clockwise and counterclockwise propagating light of

$$\Delta\phi \approx \frac{8N\pi\Omega A}{\lambda c}. \quad (8.33)$$

The instability of fiber is so very well suppressed in fact that you don't get any output below a certain rotation speed—the backscattered light dominates until the Sagnac phase shift gets sufficiently large. Transient effects of course break this symmetry because they arrive at different times unless they occur exactly midway on the path.

Sagnac interferometers based on big coils of fiber have a lot of bend birefringence, just like Faraday effect ammeters of Section 8.12.1.

8.15 MULTIPLE-BEAM FIBER INTERFEROMETERS

8.15.1 Fabry–Perot

There are two kinds of fiber Fabry–Perots: *intrinsic* ones, where the fiber is entirely within the cavity, and *extrinsic* ones, where at least part of the cavity is outside the fiber (Figure 8.11).

An intrinsic F-P is a stretch of fiber with a Bragg grating or a mirrored facet at each end. It is sensitive to all the usual F-P things, minus misalignment, plus the usual fiber things, which is a really bad trade. Polarization instability prevents the use of intrinsic F-Ps in unattended situations and severely limits the attainable finesse; if you don't mind babysitting, however, their potentially very long path length makes them extremely sensitive detectors of tuning, strain, and temperature, *provided* that your laser linewidth is narrow enough. Because the cavity is normally long, the free spectral range is short; a 10 m piece of fiber has an FSR of 10 MHz, and if its finesse is 50, the FWHM of the passband is 200 kHz. Diode lasers typically have linewidths of 10–100 MHz, so the fringes show up in the spectrum rather than in the intensity, making all-fiber F-Ps hard to use except with very fancy lasers.

Low finesse intrinsic F-Ps can be made by cleaving the desired length, depositing a TiO_2 film on each end, and fusion-splicing the result back into the main fiber. Reflectivities are 1–10%. Because of their low finesse, they are optically much more like unbalanced Michelsons or Mach–Zehnders, but need no couplers. Another trick is the fiber loop mirror of Figure 8.11c, in which a cross-connected fiber coupler sends part

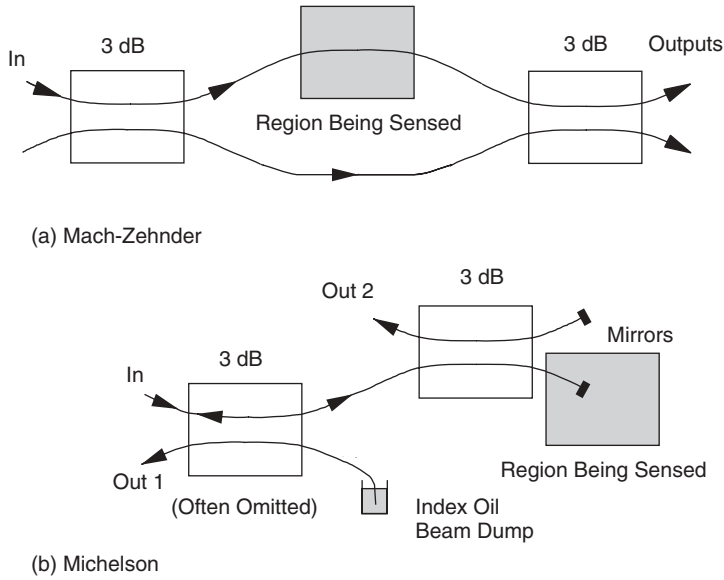


Figure 8.11. Two-beam fiber interferometers.

of the light back on itself. If the power coupling ratio is $x : (1 - x)$ and we neglect depolarization, the reflectance is

$$R = x(1 - x) \sin^2(\phi). \tag{8.34}$$

This simple fiber device can be used as a two-beam Sagnac interferometer itself, or as a mirror in a larger system such as a fiber laser. Since the reflectance is a very strong function of wavelength, temperature, and polarization, some sort of tuning will probably be required.

Extrinsic F-Ps come in two flavors: fiber delay lines with one or more bulk optic mirrors, and fiber coupled F-Ps, where the fiber just brings light to the real Fabry–Perot. An extrinsic F-P with one Faraday rotator mirror corrects the polarization problems of the original, and in the process makes the effective length of the cavity polarization independent, thereby removing one source of ambiguity. (See Figure 8.12.)

8.15.2 Ring Resonator

A ring resonator is very similar to an F-P, except that its transmission looks like the reflection of an F-P. It has all the same problems, too. If you launch light going both ways, you can use a ring as a resonant fiber-optic gyro. A good technique to eliminate polarization problems in a ring resonator is to use PM fiber, but splice it rotated 90° , so that the two polarizations exchange identities each turn. This actually makes the periodicity twice as long, but produces two bumps per period; adjacent passbands are slightly different.




Figure 8.12. Fiber Fabry–Perot sensors: (a) extrinsic, with external diaphragm; (b) intrinsic, with multilayer coating spliced into fiber; and (c) fiber loop (Sagnac interferometer used as a mirror).

8.15.3 Motion Sensors

Doppler-type motion sensors are a good match for fiber interferometry; you put a collimator on the output of one fiber of a Michelson interferometer and point it at the thing you want to measure. Absolute phase stability is not usually needed, and providing the polarization instability is controlled, all the low frequency headaches can be filtered out, and fringe-counting accuracy is usually good enough. Near-field applications such as fiber catheter flow cytometry don't need the collimator, but far-field ones do, and a bit of retroreflecting tape is usually a big help in getting light back into the single-mode fiber.

If the radial velocity is too high to count fringes ($\mathbf{v} \cdot \mathbf{k}/2\pi > 100$ MHz or so), you can use the fiber interferometer as a delay discriminator on the returned light instead.

8.15.4 Coherence-Domain Techniques

There is a family of techniques for measuring displacement and distributed parameters such as temperature without needing a fast photon-limited OTDR setup (Figure 8.13). They rely on the idea of white-light fringes. The measured output of an interferometer is a short-time-averaged autocorrelation between the signal and LO, which means that for a very wideband source fringes appear only when the path lengths of the two arms are very close. The key feature of coherence-domain measurements is that, apart from polarization funnies, the autocorrelation of the signal is preserved pretty well through the fiber. Some dispersion exists, of course, but that is compensated by the equal-length reference arm. Wideband sources such as LEDs have only a single autocorrelation peak,

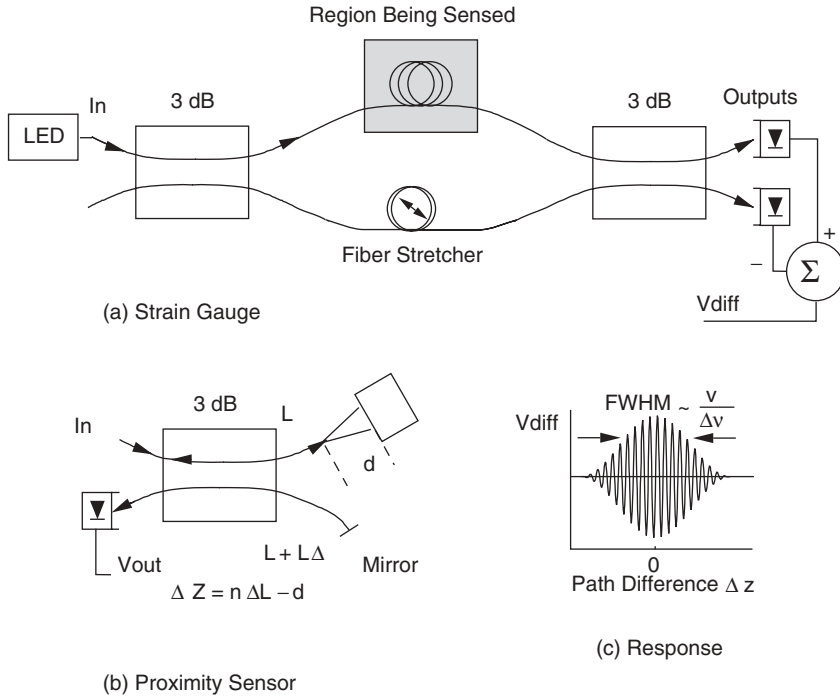


Figure 8.13. Coherence-domain techniques: (a) strain gauge, (b) OCDR proximity sensor, and (c) typical response curve.

unlike Fabry–Perots and highly coherent interferometers. Thus by stretching the reference arm until the white-light fringes appear, we can measure the distance to a scatterer unambiguously.

The accuracy is correspondingly less, because we have to locate the center of the autocorrelation peak, which of course has a fast carrier frequency imposed on it; thus unless we can use an optical *I/Q* mixer to get the true envelope, there is always a one-fringe error band.

Because the autocorrelations of multimode lasers tend to be about 10^3 fringes wide, the SNR needed to get one-fringe accuracy is far higher than with a fringe-counting interferometer: 50–60 dB instead of 15 dB, which is a serious limitation. An LED with a 5% bandwidth is rather better, but on the other hand, you don’t get as much light. One way to get around it is to synthesize an autocorrelation function that is much less smooth, and *in a known way*. For example, we can use a slightly unbalanced Mach–Zehnder to put fringes on the autocorrelation, as shown. This can also be done by using two wavelengths, or a filter with two passbands on a single wideband source. It is analogous to a dual-wavelength interferometer, where the fast fringes give high resolution and the long beat length disambiguates the 2π phase wraps. This is an example of a multiple-scale measurement, which we discuss in Section 10.4.2. Coherence-domain techniques are good for low resolution length measurements, such as noncontact surface sensing, for interrogation of features in fiber that are far too close for OTDR (e.g., 1 mm apart), and for disambiguation of interferometric data. They’re also good for measurements in turbid or multiple-scattering media such as tissue.

8.16 PHASE AND POLARIZATION STABILIZATION

Getting a stable operating point in a fiber sensor is a bit of a challenge, as we've seen. There are lots of methods in use, of which the following are the best.

8.16.1 Passive Interrogation

If the paths are sufficiently unequal (>1 cm), the phase and amplitude can be interrogated separately using a tunable laser. To avoid the linewidth problems we had with the Fabry–Perot, the laser should be tunable over many times its linewidth, and the nominal path difference adjusted accordingly—for example, $\Delta L = 1.5 - 2.5$ cm for a CD-type diode laser whose tuning range is 1 cm^{-1} . This allows stabilizing the operating point electronically, by tuning the laser.

One good way to do it is with the modulation-generated carrier approach, as in Section 13.9.6. This technique requires modulating the phase of the light by ± 2.6 radians or so, where the J_1 and J_2 terms in the Bessel expansion of the phase modulation are equal. Because you can combine in-phase and quadrature (I and Q) signals so as to ignore the slow phase drift completely (perhaps with a feedback loop to get the balance exactly right), this is a pretty effective solution to the phase problem. You do have to keep the modulation stable at that value, and the SNR suffers surprisingly little, since 85% of the interference term energy is in the first and second harmonics.

8.16.2 Frequency Modulation

Sufficiently rapid FM modulation can help a bit in reducing the coherence fluctuation noise, by transforming the widely different time delays of the twice-scattered components into large frequency shifts so that most of that contribution lies outside the measurement bandwidth. The modulation frequency needs to be several times the laser linewidth, and the modulation index should be large. Still, as we saw in Section 4.7.4, this technique isn't as good as we'd like. Very rapid chirps turn distance into frequency shifts, which is often better than simple FM.

8.16.3 Fringe Surfing

You can use a fiber stretcher or other phase modulator to surf on a fringe, essentially phase-locking to the fringe drift. This scheme, discussed in Section 10.7.7, is the first method everybody thinks of. Unfortunately, the phase will in general walk a long way over long times, far further than phase modulators can accommodate; for instance, a YAG laser resonator that changes by 0.001% in length will change the delay of a 50 m piece of fiber by 0.001% of 70 million wavelengths, or 1400π radians, and a 1.5°C temperature change in the fiber will do the same. Thus fringe surfing works only in the lab unless your measurement can be interrupted for a second or two at irregular intervals to reacquire lock. Because of the limited bandwidth of most such setups (fiber stretchers resonate in the hundreds of hertz to low kilohertz), they are seriously microphonic, too—bump the table and the loop will lose lock, to come back to rest somewhere else tens of milliseconds later.

As with a PLL phase demodulator (Section 13.9.5), you can take as the system output either the detected output from the optical phase detector, before the loop integrator, or the control voltage to the fiber stretcher, after the integrator. Unfortunately, the control

voltage here is useless for measurements. There are multiple operating points on the fiber stretcher, so every time the loop resets, there will be a huge voltage step on that output; furthermore, the actuator is nonlinear, so there is no easy way of reconnecting a run of data that has been interrupted by a reacquisition.

8.16.4 Broadband Light

Using broadband light and spectrally encoded data makes a lot of things easier, provided the source drift is taken care of. The main drawback is the pixel resolution of the spectrometers used to decode the spectrum.

8.16.5 Ratiometric Operation

One way to get rid of the scale factor drift is by comparing the signal to a known one; for example, a Faraday effect gaussmeter could use an electromagnet to excite the coil stably at a known frequency; the ratio of the two signals would correspond to the ratio of the applied field to the standard one. This requires only high linearity, and not high stability. Avoid trying to use two wavelengths, because etalon fringes among other things are highly tuning sensitive.

8.16.6 Polarization-Insensitive Sensors

Polarization drift can be greatly reduced by using Faraday rotator mirrors, as we've seen, and although this isn't magic, it keeps the polarization drift well enough controlled that the operating point doesn't move around much. The residual drift will still need correcting for in many cases, but that can be done in postprocessing.

One source of residual polarization drift is polarization-dependent losses, for example, Fresnel reflection at angled fiber cleaves. These sorts of effects produce errors by coupling the supposedly orthogonal modes together, thus ruining the orthogonality on which the FRM trick depends. Another is light that doesn't get transformed in the FRM, for example, facet reflections and front surface reflections, and multiple reflections in the Faraday crystal. A third is errors of the FRM itself, for example, misadjustment, dispersion, drift, and nonuniformity.

8.16.7 Polarization Diversity

The modulation-generated carrier approach is an example of phase diversity; it is also possible to use polarization diversity, although this is much more complicated. One method is to combine the ideas of Sections 6.10.10 and 15.5.4; after recollimating the output light, use bulk optics to split it into two or three copies. Use polarizing prisms rotated through different angles to select different projections of the two fields, and pick whichever one has the best fringe visibility. This takes care of the problem of misaligned major axes, leaving only the usual relative phase problems.

8.16.8 Temperature Compensation

Separating temperature from other effects usually involves measuring temperature independently and using a calibration curve. Another property of the fiber (e.g., Raman

backscatter) can be used, or an IC temperature sensor. The particular problems with this are generating the calibration curve, and figuring out what temperature to use, in the face of significant gradients across the dimension of the fiber system.

8.16.9 Annealing

You can anneal fiber that has been wound round a mandrel by heating it up to 800 °C for a while and cooling it slowly, being sure not to stretch it by letting the mandrel expand too much. That gets rid of the bend birefringence, but it tends to be hard on the jacket, which makes it unattractive unless you really, really need to, e.g. in Faraday sensors for power transmission ammeters. It is wise to leave the fiber undisturbed on the mandrel afterwards.

8.17 MULTIPLEXING AND SMART STRUCTURES

Some types of intrinsic fiber sensor lend themselves to multiplexing, which permits the use of many sensors with only one readout system, and perhaps with only one interrogating fiber. Multiplexing can be done in all the standard electronic ways; time-division (TDM), frequency-division (FDM), and code-division (CDM), plus a uniquely optical one, *coherence multiplexing*. The details of this are really beyond our scope, being network and interconnection issues, but they are discussed in Udd and elsewhere.

8.18 FIBER SENSOR HYPE

As the old saw runs, “It’s not what you don’t know that hurts you, it’s what you do know that ain’t so.” Fiber sensors are currently still fashionable, though not as modish as they were. In part, this is because of the advantages already enumerated, although as we’ve already seen, there’s less than meets the eye about some of them. The author is somewhat sceptical of fiber sensor claims in general, because the sheer volume of hype obscures any clear-eyed assessment of their strengths and weaknesses, and because the discussion in fiber sensor papers—even review papers and book chapters—seldom seems to include comparisons to bulk optic alternatives. We’re now in a position to enumerate and critique a few.

1. *Fiber Sensors Offer Unique Capability.* Some types certainly do. Fiber gyroscopes, Bragg grating sensors, distributed sensors like Raman thermometers, and arguably Faraday effect current sensors all have capabilities that are difficult or impossible to replicate with bulk optics. Fiber-optic smart structures, oil-well thermometers, and fiber catheter devices all have an important place.

2. *Fiber Sensors Are Cheap.* They could be, in principle, but sensitive ones almost never are—people who are used to the prices of electronic sensors are liable to keel over from pure sticker shock. Some intensity-based systems, such as fluorescent fiber thermometers, really are intrinsically cheap. Those tend to be the ones where the fiber is just used to pipe the light around, which is fair enough. Distributed systems, such as fiber Bragg grating strain gauge systems, where the cost of the interrogation system is

amortized across hundreds of sensors, might one day be cheap on a per-sensor basis, which is why they are interesting for fiber-optic smart structures.

Cheapness is often said to be a consequence of economies of scale in the telecommunications business, but it ain't so. Telecommunications fibers are not single mode at the wavelengths of cheap sources and detectors, and the cost of fiber transmission systems is not driven by the cost of lasers and detectors, as fiber sensor costs are. Besides, the *real* economies of scale are in diode laser-based consumer products—which are 100% bulk optics. (Compare the cost of a 1.5 μm diode laser to a 670 nm one of the same power level, if you doubt this.)

Simplicity and low cost can rapidly go away when we start dealing with phase and polarization instability, etalon fringes, FM–AM conversion, and extremely low étendue; we often find ourselves using \$10,000 single-frequency diode-pumped YAG lasers to get decent performance from \$50 worth of fiber.

It's also common to find apples-to-oranges comparisons; if a sensor has to work in situ, compare a fiber system to a miniaturized and cost-reduced bulk optic system, not to some behemoth with an argon laser and a table full of Newport mounts. Think in terms of DVD drives and digital cameras.

3. *Fibre Sensors Are Highly Sensitive.* This is sometimes true compared with nonoptical sensors, but is more or less entirely false in comparison to bulk optics. Leaving aside the expensive kinds, like gyros and Faraday rotator mirror interferometers, the really unique fiber sensors are all relatively low sensitivity devices. What's more, the performance of even fancy fiber sensors falls far short of the fundamental physical limits.[†] Fibers are highly sensitive to temperature, bending, vibration, you name it, and that leads to enormous amounts of low frequency drift and noise. High frequency performance is limited by the multiplicative effects of the low frequency stuff, including drift of the scale factors and operating points—things we'd never tolerate in a DVM—and by demodulation of source phase noise by etalon fringes and double Rayleigh scattering.

There are an awful lot of interferometric fiber sensor papers where somebody is bragging about his 120 dB dynamic range in 1 Hz. That isn't a bad number, but in reading these, remember that a bulk optic interferometer with a \$5 diode laser and a \$10 laser noise canceler just sits there at the shot noise level forever, unattended, with a dynamic range of 160 dB in 1 Hz; and that performance isn't hype, it's available in commercial hardware.[‡] Don't forget to ask about their relative stability and $1/f$ noise, either.

[†]This isn't entirely a fiber problem, of course.

[‡]To be fair, getting the noise canceler down to \$10 requires building your own—the commercial ones are over \$1000. The good part of this is that it's easy to do.

Optical Systems

A pile of rocks ceases to be a rock pile when somebody contemplates it with the idea of a cathedral in mind.

—Antoine de Saint-Exupéry, *Flight to Arras*

9.1 INTRODUCTION

An optical system is a collection of sources, lenses, mirrors, detectors, and other stuff that (we hope) does some identifiable useful thing. We've talked about the pieces individually, but until now we haven't spent much time on how they work together. This chapter should help you to think through the behavior of the system from end to end, to see how it ought to behave. That means we need to talk about the behavior of light in systems: practical aberration and diffraction theory, illumination and detection, and how to calculate the actual system output from the behavior of the individual elements.

9.2 WHAT EXACTLY DOES A LENS DO?

In Section 4.11.2, we looked at the Gaussian (i.e., paraxial) imaging properties of lenses. We were able to locate the focus of an optical system, calculate magnification, and generally follow the progress of a general paraxial ray through an optical system by means of multiplication of *ABCD* matrices.

Here, we concentrate on the finer points, such as the aberrations of an optical system, which are its deviations from perfect imaging performance. We will use three pictures: the pure ray optics approach, where the aberrations show up as ray spot diagrams where not all the rays pass through the image point; the pure wave approach, where the aberrations are identified with the coefficients of a polynomial expansion of the crinkled wavefront, derived from exact calculation of the wave propagation through the system; and a hybrid ray/wave picture. (See Figure 9.1.)

The hybrid picture is messy but useful and, in fact, is the basis of most "wave optics" models. It takes advantage of the fact that ray optics does a good job except near focus or in other situations where diffraction is expected to be important. Accordingly, we trace rays to the vicinity of the exit pupil from a single object point, construct a wavefront

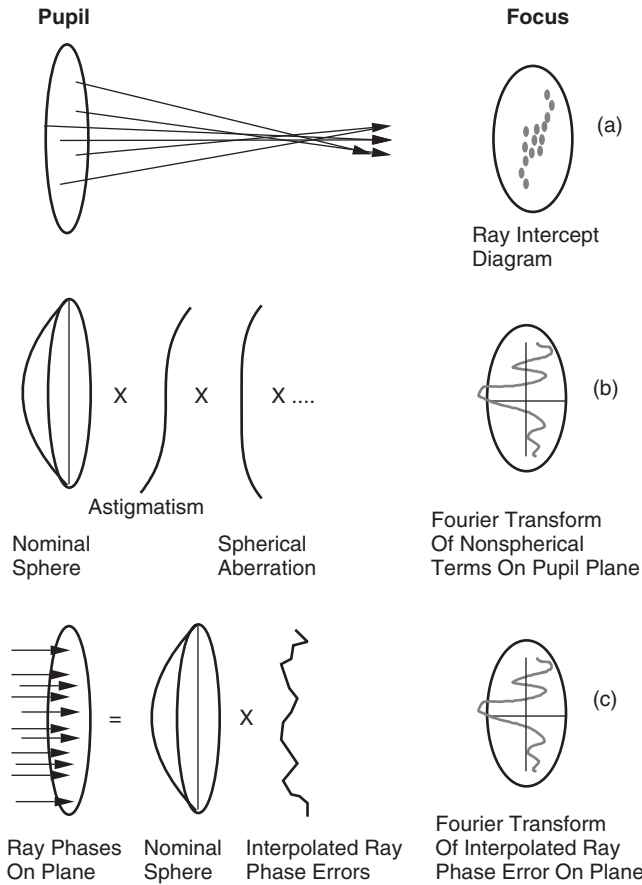


Figure 9.1. Three ways of looking at an imperfect optical system: (a) ray spot diagram, (b) wave-front polynomial expansion, and (c) wave aberration of rays.

whose phase is determined by the calculated propagation phase along the ray paths, and then use wave optics from there to the focus. This is unambiguous unless rays traversing widely different paths get too close to one another. By now you'll recognize this as another case of putting it in by hand, which is such a fruitful approach in all of engineering.

The different pictures contain different information. A ray tracing run is very poor at predicting the shape of the focused spot, but contains lots of information about the performance of the system across the field of view. For example, field curvature and geometric distortion show up clearly in a ray trace, since different field angles are presented at once, but tend to disappear in a pure wave propagation analysis, where only a single field position can readily be considered at a time.

9.2.1 Ray Optics

Ray optics assumes that all surfaces are locally planar, and that all fields behave locally like plane waves. To compute the path of a ray encountering a curved surface, we

notionally expand the ray into a plane wave, and the surface into its tangent plane. We then apply the law of reflection and Snell's law to derive the direction of the reflected and refracted \mathbf{k} vectors, and these become the directions and amplitudes of the reflected and refracted rays. This is a first-order asymptotic approach, valid in the limit $ka \rightarrow \infty$, where a is the typical dimension of the surface (e.g., its radius of curvature, diameter, or whatever is most appropriate). There is a slight additional subtlety. A single ray, being infinitesimally thin, transports no energy; to find out the field amplitudes, we must consider small bundles of rays, or *pencil beams*, occupying an element of cross-sectional area $d\mathbf{A}$, measured in a plane normal to their propagation direction. Conservation of energy requires that the product of the ray intensity $I d\mathbf{A}$ be constant along the axis. Curved surfaces and any refraction or diffraction will in general cause $|d\mathbf{A}|$ to change, either by anamorphic magnification or by focusing. Thus in computing the contribution Di of a given ray bundle to the intensity at a given point \mathbf{x}' from that at \mathbf{x} , we must multiply by the Jacobian,

$$dI(\mathbf{x}') = dI(\mathbf{x}) \left| \frac{d\mathbf{A}}{d\mathbf{A}'} \right|. \quad (9.1)$$

If the incoming illumination is spatially coherent, we must instead sum the (vector) fields, which transform as the square root of the Jacobian,

$$d\mathbf{E}(\mathbf{x}') = d\mathbf{E}(\mathbf{x}) \sqrt{\left| \frac{d\mathbf{A}}{d\mathbf{A}'} \right|}. \quad (9.2)$$

Going the other way, for example, computing a specular reflection by starting from the obliquely illuminated patch to the propagating beam, we have to put in the reciprocal of the obliquity factor—otherwise energy wouldn't be conserved on reflection from a perfect mirror. (See Section 9.3.5.) The mathematical way of putting this is that the Jacobian of the oblique projection equals the ratio $\cos\theta_2/\cos\theta_1$. We saw this effect in radiation from planar surfaces in Section 1.3.12, and it shows up in wave optics as the obliquity factor (see Sections 9.2.1 and 9.3.4).

9.2.2 Connecting Rays and Waves: Wavefronts

In order to move from one picture to another, we have to have a good idea of their connections. The basic idea is that of a *wavefront* (Figure 9.2). Most people who have had an upper-level undergraduate physical optics class will picture a wavefront as a plane wave that has encountered some object (e.g., a perforated plane screen or a transparency) and has had amplitude and phase variations impressed upon it. While this picture isn't wrong, it also isn't what an optical designer means by a wavefront, and the differences are a frequent source of confusion, especially since the same diffraction integrals are employed and the conceptual differences are seldom made explicit.

A physicist will tell you that a wavefront is a surface of constant phase, which can have crinkles and ripples in it, whereas a lens designer will say it's the deviation from constant phase on a spherical surface centered on the Gaussian image point. In actual fact, whenever people actually calculate imaging with wavefronts (as opposed to waving

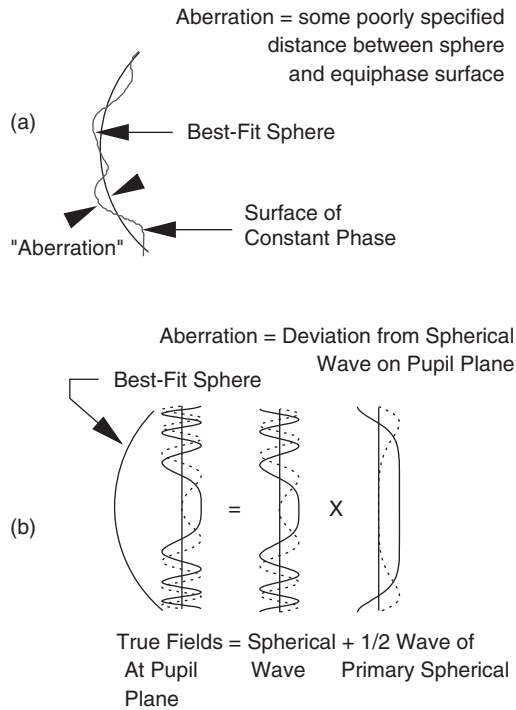


Figure 9.2. Wavefront definitions: (a) surface of constant phase and (b) deviation from a prespecified spherical wave on a plane.

their arms) the real definition is the phase deviation, on a plane, from a spherical wave:

$$W(\mathbf{x}) = K \arg(e^{ik|\mathbf{x}-\mathbf{x}_0|}\psi(\mathbf{x})), \tag{9.3}$$

where K is a constant that expresses the units in use: radians, waves (i.e., cycles), or OPD in meters. As long as the deviations from sphericity are sufficiently small and slow, the two descriptions are equivalent.

Deviations from the perfect spherical wave case are called *aberrations*; aberration theory is in its essence a theory of the way phase propagates in an optical system. Amplitude and polarization information is not given equal consideration, which leads to the total disregard of obliquity factors, among other things. A vernacular translation is, “it gives the wrong answer except with uniform illumination and small apertures, but is close enough to be useful.”

Aside: Phase Fronts on the Brain. From what we know already of diffraction theory, concentrating solely on phase like this is obviously fishy; an amplitude-only object such as a zone plate can destroy the nice focusing properties, even while leaving the phase intact initially. Neglect of amplitude effects is our first clue that aberration theory fundamentally ignores diffraction. The software packages sold with some measuring interferometers have serious errors traceable to this insistence on the primacy of phase information. The author’s unscientific sample is not encouraging; he has had two such units, from

different manufacturers, manufactured ten years apart. Both were seriously wrong, and in different ways.

9.2.3 Rays and the Eikonal Equation

We have the scalar wave equation, which allows us to predict the fields everywhere in a source-free half-space from an exact knowledge of the full, time-dependent, free-space scalar field on the boundary of the half-space. In the limit of smooth wavefronts and wavelengths short compared to D^2/d (where D is the beam diameter and d the propagation distance), we can neglect the effects of diffraction. In this limit, the gradient of the field is dominated by the $i\mathbf{k} \cdot \mathbf{x}$ term, and each segment of the wavefront propagates locally as though it were its own plane wave, $\psi_{\text{local}}(\mathbf{x}) \approx A \exp(i\mathbf{k}_{\text{local}} \cdot \mathbf{x})$.

Phase is invariant to everything, since it's based on counting; we can look on the phase as being a label attached to a given parcel of fields, so that the propagation of the field is given by the relation between \mathbf{x} and t that keeps ϕ constant. That means that the direction of propagation is parallel to $\mathbf{k}_{\text{local}}$,

$$\mathbf{k}_{\text{local}} \approx \frac{-i \nabla \psi}{\psi}, \quad (9.4)$$

which gives us a natural connection between rays and wavefronts.

If we take a trial solution for the scalar Helmholtz equation,

$$\psi(\mathbf{x}) = A(\mathbf{x})e^{ik_0 S(\mathbf{x})}, \quad (9.5)$$

applying the scalar Helmholtz equation for a medium of index n and taking only leading order terms as $k_0 \rightarrow \infty$ suppresses all the differentials of A , which is assumed to vary much more slowly than $\exp(ik_0 S)$, leaving the *eikonal equation*

$$|\nabla S(\mathbf{x})|^2 = n^2(\mathbf{x}), \quad (9.6)$$

where the eikonal $S(\mathbf{x})$ is the optical path length (it has to have length units because $k_0 S$ must be dimensionless). Once we have S , we can get the propagation direction from (9.4). What's more, Born and Wolf show that in a vector version of this, the Poynting vector lies along ∇S , so in both pictures, (9.6) is the natural connection between rays and waves.

This connection is not a 1:1 mapping between rays and waves, however. The eikonal can be used to attach rays to a wavefront, then trace the resulting rays as usual, but that procedure doesn't lead to the same results as propagating the field and then computing the eikonal, because the whole eikonal idea breaks down near foci and caustics, as well as having serious problems near shadow boundaries. The approximation becomes worthless at foci because the phase gradient vectors can never really cross; that would require a field singularity, which is impossible for the wave equation in a source-free region. Another way to say this is that the optical phase can be made a single-valued function of position in any given neighborhood, and therefore its gradient is also a single-valued function of position. So identifying rays purely as gradients of the phase cannot lead to rays that cross. For example, the eikonal equation predicts that a Fresnel amplitude zone plate (see Section 4.13.2) has no effect on wave propagation other than blocking some of the light.

The eikonal approximation shares the wave optics difficulty that there's no simple way to include multiple source points in a single run; the total field has only one gradient at each point.

9.2.4 Geometrical Optics and Electromagnetism

Geometrical optics (GO) is an asymptotic electromagnetic theory, correct in the limit $k \rightarrow \infty$. Like most asymptotic theories (e.g., the method of steepest descents), it requires some funny shifts of view. We go back and forth between considering our ray to be infinitesimally narrow (so that all surfaces are planes and all gradients uniform) and infinitely broad (so that the beam steers like a plane wave). The way a GO calculation goes is as follows:

1. Pick some starting rays, for example, at the centers of cells forming a rectangular grid on the source plane. Assign each one an amplitude and phase. (You can do GO calculations assuming incoherent illumination, but you're better off computing the amplitudes and phases of each ray and applying random phasor sums to the results—it's computationally cheap at that stage and avoids blunders.)
2. Assuming that the true optical field behaves locally like a plane wave, apply the law of reflection, Snell's law, and so on, to follow the path of the ray to the observation plane.
3. Add up the field contributions at the observation plane by computing the optical phase (the integral of $\mathbf{k} \cdot d\mathbf{s}$ over the ray path) and amplitude from each ray and summing them up. Remember to apply the Jacobian—if you imagine each ray occupying some small patch of source region, the area of the patch will in general change by the time it gets to the observation plane. This isn't mysterious, it's just like shadows lengthening in the evening. Field amplitudes scale as the reciprocal square root of the area. If the ray directions at the observation plane are similar, you can add up the fields as scalars. Otherwise, you'll need to be more careful about the polarization. Nonplanar rotations of \mathbf{k} will make your polarization go all over the place.

You can trace rays in either direction, so if only a limited observation region is of interest, you can start there and trace backwards to the source. Either way, you have to make sure you have enough rays to represent the field adequately. You have to think of a GO calculation as involving nondiffracting rectangular pencil beams, not just rays; in general, the patches will overlap in some places, and you have to add all the contributions in complex amplitude.

In an inhomogeneous but isotropic medium, the geometric optics laws need to be generalized slightly. The local direction of wave propagation is the gradient of the phase. This leads to the eikonal equation (9.6) or the curvature equation (9.12), which are differential equations giving the change of the ray direction as a function of distance along the ray. Note that the Jacobian has to be carried along as well if you want to get the correct answer for the field amplitudes. If the medium is anisotropic as well as inhomogeneous, life gets a good deal harder, as you have to carry along the polarization state and any beam walkoff as well as the \mathbf{k} vector and Jacobian as you go. If you have sharp edges, caustics, or shadows, geometric optics will give you the wrong answers there—it ignores diffraction, will exhibit square-root divergences at caustics and edges, and will predict zero field in the shadow regions.

9.2.5 Variational Principles in Ray Optics

Some propagation problems yield to a less brute-force approach: calculus of variations. If the medium is nonuniform, so that n is a function of \mathbf{x} , we need to put an integral in the exponent instead,

$$\psi(\mathbf{x}) = A(\mathbf{x}) \exp \left[-i\omega t + ik_0 \int_{\text{path}} n(\mathbf{x}) ds \right], \quad (9.7)$$

as in the eikonal approximation (9.5). Since the ray path depends on n , we don't know exactly what path to do the integral over, so it doesn't look too useful. In fact, we're rescued by *Fermat's principle*, a variational principle that states that

$$\delta S = \delta \int_{\text{path}} n(\mathbf{x}) ds = 0; \quad (9.8)$$

that is, the integral has an extremum on the true ray path P . The path yielding the extremum is said to be an *extremal*. Fermat called it the *principle of least time*, which assumes that the extremal is a global minimum. There's obviously no global maximum—a given path can loop around as much as it wants—so this is a good bet.

We solve variational problems of this sort by imagining we already know the parametrized $\mathbf{P} = \mathbf{x}(u)$, where $\mathbf{x}(0)$ is the starting point \mathbf{x}_0 , $\mathbf{x}(u_1)$ is the end point \mathbf{x}_1 , and u is a dummy variable. We demand that a slight variation, $\epsilon \mathbf{Q}(u)$ (with $\mathbf{Q} \equiv 0$ at the ends of the interval), shall make a change in S that goes to 0 faster than ϵ [usually $O(\epsilon^2)$] as $\epsilon \rightarrow 0$. Since the arc length is all we care about, we can parameterize the curve any way we like so we'll assume that \mathbf{x} is a continuous function of u and that $d\mathbf{x}/du \neq 0$ in $[0, u_1]$. Thus

$$\int_0^{u_1} ((n(\mathbf{x}) + \epsilon \mathbf{Q} \cdot \nabla n) |\dot{\mathbf{x}} + \epsilon \dot{\mathbf{Q}}| - n(\mathbf{x}) |\dot{\mathbf{x}}|) du = O(\epsilon^2), \quad (9.9)$$

where dotted quantities are derivatives with respect to u . Since it's only the ϵ term we're worried about, we series-expand the squared moduli, cancel the zero-order term, and keep terms of up to order ϵ , which yields

$$\int_0^{u_1} \left(\frac{n\dot{\mathbf{x}} \cdot \dot{\mathbf{Q}}}{|\dot{\mathbf{x}}|} + |\dot{\mathbf{x}}| \mathbf{Q} \cdot \nabla n \right) du = 0, \quad (9.10)$$

which isn't too enlightening until we notice that it's nearly a total derivative, with one term in \mathbf{Q} and one in $\dot{\mathbf{Q}}$. Integrating by parts, and using the fact that $\mathbf{Q} = 0$ at the ends and is continuous but otherwise arbitrary, we get the result

$$\frac{\nabla n - \dot{\mathbf{x}}(\dot{\mathbf{x}} \cdot \nabla n)}{n} = \frac{(\ddot{\mathbf{x}} - \dot{\mathbf{x}}(\dot{\mathbf{x}} \cdot \ddot{\mathbf{x}}))}{|\dot{\mathbf{x}}|^2}, \quad (9.11)$$

which can be written more neatly by changing back to arc length and using the convention that ∇_{\perp} is the gradient perpendicular to $\dot{\mathbf{x}}$, yielding the *curvature equation*

$$\frac{\nabla_{\perp} n}{n} = \frac{d^2 \mathbf{x}}{ds^2} \Big|_{\perp}. \quad (9.12)$$

This says that the curvature of the path is equal to the perpendicular gradient of $\log n$, which makes a lot of physical sense, since we don't expect the path to change when we (say) double the refractive index everywhere.[†]

9.2.6 Schlieren Effect

One interesting consequence of the curvature equation (9.12) is that light waves steer like tanks: they turn toward whichever side goes more slowly. Accordingly, a refractive index gradient causes the wave to bend, the *schlieren effect*. Since $dn/dT < 0$ for gases, a temperature gradient in air produces schlieren, which is why there are mirages. On a hot, sunny day, the ground is warmer than the air, so $dn/dz < 0$ and light bends upwards; an image of a patch of sky appears near the ground in the distance, looking like a pool of water. At sea, with the opposite sign of dT/dz , a ship becomes visible before it crosses the geometrical horizon. More complicated gradients can cause multiple images, as in the beautifully named *fata Morgana* (after Morgan Le Fay, King Arthur's nemesis); ghostly shorelines with fantastic mountains can appear in the middle of the ocean. (Who said there was no poetry in optics?)

There are a couple of examples that show up more often in instruments: *thermal lensing*, which is pretty much like a mirage, and gradient-index (GRIN) optics, as we saw in Sections 4.13.1 and 8.3.5. Thermal lensing in nonaqueous liquids can be a big effect—big enough to be a sensitive laser spectroscopy method. A visible probe laser traverses a long path in some solvent, coaxially with an infrared pump beam. An axially symmetric temperature gradient results in progressive defocusing of the probe beam, which can be detected very sensitively with a masked detector. Water is a disappointing solvent for thermal lensing, with a low dn/dT and a high thermal conductivity.

9.2.7 The Geometrical Theory of Diffraction

For objects whose typical dimension a is large compared to a wavelength, the ordinary ray optics laws (the law of reflection and Snell's law) apply with high absolute accuracy except near shadow boundaries and places where rays cross—foci and caustics. (The relative accuracy is of course also very bad inside shadows, where geometric optics predicts zero fields.) For such large objects, it is reasonable to apply a local correction in these situations, the *geometrical theory of diffraction* (GTD), formulated by Keller,[‡] Ufimtsev,[§] and others. Like ray optics, GTD is an asymptotic theory valid in the limit $ka \gg 1$, but it lets us include higher order terms to get better accuracy. The basic idea is that illuminated bodies follow geometrical optics (GTD) or physical optics (PTD) except within a wavelength or two of shadow boundaries and sharp edges. Large objects ($ka \gg 1$) can be looked on as arrangements of flats, curved regions, and edges,

[†]Extremals that minimize some smooth functional such as (9.9) are called *weak extremals*, because the true minimum may not be continuous. If discontinuous and unsmooth functions are considered, the result is called a *strong extremal*. The strong extremal is usually a global minimum, never a maximum, but sometimes just a local minimum. If you don't know any calculus of variations, consider learning it—it isn't difficult, and it's a great help in optical problems. The book by Gelfand and Fomin (see the Appendix) is a good readable introduction.

[‡]J. B. Keller, Geometric theory of diffraction. *J. Opt. Soc. Am.* **52**, 116–130 (1962).

[§]Pyotr Y. Ufimtsev, *Method of Edge Waves in the Physical Theory of Diffraction*. Available at <http://handle.dtic.mil/100.2/AD733203>.




Figure 9.3. GTD and PTD calculations combine geometrical or physical optics calculation with a correction factor due to the vector diffraction from edges and shadow boundaries: (a) edge waves from discontinuities and (b) creeping waves from curves.

and the scattered fields can be decomposed into sums over those individual contributions. Locally these can be described as flat, ellipsoidal, cylindrical, wedge-shaped, or conical—all shapes for which rigorous analytic solutions exist, at least in the far field. The beautiful trick of PTD is to take each of these canonical cases, solve it twice—once rigorously and once by physical optics—and then subtract the two solutions, yielding the edge diffraction contribution alone. For sharp edges, this is expressed as a line integral around the edges, turning a 3D problem into a 1D problem. In most calculations the diffracted contributions are further approximated as *diffracted rays*. (Curved surfaces give rise to *creeping rays*, which are harder to deal with—the F-117A stealth fighter is all flats and angles because it was designed using 1970s computers.)

The two kinds of diffracted rays are shown in Figure 9.3: *edge rays*, which emanate from each point of discontinuity, such as a corner or an edge, and *creeping rays*, generally much weaker, which emanate from shadow edges on smooth portions of the surface.

So the way it works is that you do the calculation via geometrical or physical optics, which ignores the edge contributions, and then add in the vector diffraction correction in a comparatively simple and computationally cheap way. Complicated geometries will have important contributions from multiple scattering, leading to a perturbation-like series in which n th order terms correspond to n -times scattered light.

These approximations are usually very complicated, but on the other hand, they contain information about *all* incident and scattered angles, *all* positions, and *all* wavelengths, in one formula. The information density in that formula dwarfs that of any numerical solution, and frequently allows algebraic optimization of shapes and materials, which is very difficult with numerical solutions. This makes GTD and PTD well suited for design problems, especially with computer algebra systems available for checking.

GTD approximations tend to diverge as $x^{-1/2}$ at shadow boundaries, caustics, and foci, which of course are points of great interest. The same idea, local approximation by analytically known results, can be used to get a uniform asymptotic approximation, valid everywhere. For details, see the IEEE collected papers volume[†] and the excellent monographs of Ufimtsev and of Borovikov and Kinber referenced in the Appendix.

[†]Robert C. Hansen, ed., *Geometric Theory of Diffraction*, IEEE Press, New York, 1981.

9.2.8 Pupils

The entrance pupil is an image of the aperture stop, formed by all the elements ahead of the stop, and the exit pupil is the image of the same stop formed by all succeeding ones. Thus they are images of one another, and each point in the entrance pupil has a conjugate point in the exit pupil. Since nobody really knows what a lens does, we rely on this property heavily in wave optics calculations of imaging behavior. The most consistent high-NA approach to the problem is to use the ray optics of thin pencil beams to construct the fields on the pupil plane, and then propagate from there to the image using the Rayleigh–Sommerfeld integral. Remember the ray/wave disconnect: in the wave picture, *pupil* refers to the Fourier transform plane, not to the image of the aperture stop. (Like most ray/wave terms, the two are related but generally not identical.)

Not all imaging optical systems possess proper pupils; for example, a scanning system with the x and y deflections performed by separate mirrors lacks one unless intervening optics are included to image one mirror onto the other. An optical system without a pupil is not a shift-invariant system, so that Fourier imaging theory must be applied with caution.

9.2.9 Invariants

There are a number of parameters of an optical beam which are invariant under magnification. One is the state of focus: if an object point is 1 Rayleigh range from the beam waist, its image will be at 1 Rayleigh range from the waist of the transformed beam (neglecting diffraction). This is because the longitudinal magnification of an image is not M but M^2 .

The best known is the *Lagrange invariant*, which we've encountered already as the conservation of étendue. You can get this by putting two rays as the columns of a 2×2 matrix R . No matter what $ABCD$ matrix you hit R with, the determinant of the result is equal to $\text{Det}(R)$: $x_1\theta_2 - x_2\theta_1$ is invariant in the air spaces in any paraxial optical system. If we generalize to the case $n \neq 1$, the $ABCD$ matrix that goes from n_1 into n_2 is

$$\begin{bmatrix} 1 & 0 \\ 0 & n_1/n_2 \end{bmatrix}, \quad (9.13)$$

whose determinant is n_1/n_2 , so the generalized Lagrange invariant L is

$$L = n(x_1\theta_2 - x_2\theta_1). \quad (9.14)$$

The more usual form of this is the theorem of Lagrange, where for a single surface between media of indices n_1 and n_2 ,

$$n_1x_1\theta_1 = n_2x_2\theta_2. \quad (9.15)$$

Another invariant is the number of resolvable spots, which is the field of view diameter or scan distance measured in spot diameters; if we take the two ends of the scan to be the two rays in the Lagrange invariant, the range goes up as the cone angle goes down, and hence the spot size and scan angle grow together.

9.2.10 The Abbe Sine Condition

The Lagrange invariant holds for paraxial systems, but not for finite apertures. Its most natural generalization is the *Abbe sine condition*,

$$n_1 x_1 \sin \theta_1 = n_2 x_2 \sin \theta_2, \quad (9.16)$$

which we don't get for free. (Optical design programs include *offense against the sine condition* (OSC) in their lists of aberrations.) A system obeying the sine condition is said to be *isoplanatic*, and has little or no coma.[†] Like other aberration nomenclature, this term has a related but not identical meaning in the wave picture: an optical system is said to be isoplanatic if its transfer function does not vary with position in the image. You can see that this is a different usage by considering vignetting; a few missing rays won't make the sine condition false, but they will certainly change the transfer function.

Aside: NA and f-Number. It's possible to get a bit confused on the whole subject of numerical aperture and *f*-number, because there are two competing definitions of *f*# in common use. One, historically coming from photography, is

$$f\# = D/EFL = 1/(2 \tan \theta),$$

where *D* is the pupil diameter, θ is the half-angle of the illuminated cone, and EFL is the effective focal length (just focal length *f* to us mortals). There's no clear upper limit to this number—light coming in from a hemisphere effectively has an infinite radius at any nonzero focal length, so $f\# = \infty$.

The other definition, coming from microscopy, is

$$f\# = 1/(2 \sin \theta) = 0.5/NA,$$

assuming $n = 1$. Since $NA \leq 1$ in air, in this definition a hemispherical wave would be coming in at $f/0.5$. The two are equivalent for small NA and distant conjugates, so they're often confused. Photographers care most about image brightness, since that determines exposure, so the quoted *f*# on the lens barrel actually applies on the *image* side of the lens, and is nearly constant as long as the object distance $d_o \gg f$. Microscopists care most about resolution, so microscope NA is quoted on the object side, where it's also nearly constant because of the small depth of focus. The two definitions express the same information, but confusion is common when we don't keep them straight. (The author recommends the 0.5/NA definition as being closer to the imaging physics as well as giving a simpler exact formula for image brightness, since $n^2 \Omega' = \pi (NA)^2$.)

9.3 DIFFRACTION

There is not enough space in this book to treat diffraction in complete detail. For purposes of measurement systems, diffraction is important in four ways: in imaging; in gratings

[†]Optical system design is full of forbidding terms like that, but don't worry—half an hour's work and you'll be obfuscating with the best of them.

and holograms; in spatial filtering; and in vignetting, the incidental cutting off of parts of the beam by the edges of optical elements, apertures, and baffles. We've already covered the ordinary Huyghens–Fresnel theory in Section 1.3, so this section concentrates on the finite-aperture case.

9.3.1 Plane Wave Representation

Monochromatic solutions to the scalar wave equation in free space can be expressed exactly as sums of plane waves of different \mathbf{k} . The \mathbf{k} -space solution is exactly equivalent to the real-space solution; no approximation is involved. Thus if we have a focused beam, and we know its plane wave spectrum exactly, we can calculate its amplitude and phase at any point (\mathbf{x}, t) we like. It's important to hold on to this fact in discussing diffraction; once we have specialized to the scalar case, there are no further approximations in the actual wave propagation calculation. The additional approximations of diffraction theory involve how spatial Fourier coefficients on surfaces couple into plane waves, and how an obstruction in a free-space beam modifies its plane wave spectrum.

The easiest case is a plane boundary, because different plane waves are orthogonal on that boundary; thus a Fourier transform of the fields on the surface, appropriately weighted, gives the plane wave spectrum directly. Life gets significantly harder when the boundary is nonplanar. There are a handful of other coordinate systems in which the Laplacian separates, but the only three useful ones are Cartesian, cylindrical, and spherical. (Ellipsoidal coordinates are a special case of spherical for electrostatics, but not for electrodynamics.) Generally, though, unless you're a glutton for punishment, you have to choose among plane interfaces, asymptotically large spheres, and numerical solution.

9.3.2 Green's Functions and Diffraction

The study of diffraction is based on the idea of the *Green's function*, which is the response of a system consisting of a linear partial differential equation plus boundary conditions to a source term of $\delta(\mathbf{x} - \mathbf{x}')$. There is so much confusion around as to what the origins and limitations of diffraction theory are that it seems worth going through the math here. The following discussion follows Jackson fairly closely, so look there for more detail if this is unfamiliar. We solve the equation for the Green's function, and then we can solve the equation by a superposition integral of the source term $f(\mathbf{x}')$ (neglecting boundary terms),

$$\psi(\mathbf{x}) = \iiint_{\text{all space}} f(\mathbf{x}') G(\mathbf{x}, \mathbf{x}') d^3x'. \quad (9.17)$$

The usual case in diffraction theory is a bit more complicated, in that we actually have the field (rather than a source) specified on some surface, which may or may not be one of the boundaries of the space. Boundary conditions couple to the normal derivative of the Green's function, $\hat{\mathbf{n}} \cdot \nabla G$. (Don't confuse n the refractive index with \mathbf{n} the unit vector normal to the surface.)

We'll specialize to the Helmholtz wave equation, so the defining equation for G is

$$(\nabla^2 + k^2)G(\mathbf{x}, \mathbf{x}') = -\delta^3(\mathbf{x} - \mathbf{x}'). \quad (9.18)$$

The two Green’s functions of interest are G_0 , the one for free space,

$$G_0(\mathbf{x}, \mathbf{x}') = \frac{\exp(ik|\mathbf{x} - \mathbf{x}'|)}{4\pi|\mathbf{x} - \mathbf{x}'|} \tag{9.19}$$

and G_+ , the one for Dirichlet boundary conditions on the plane $z = 0$,

$$G_+(\mathbf{x}, \mathbf{x}') = \frac{\exp(ik|\mathbf{x} - \mathbf{x}'|)}{4\pi|\mathbf{x} - \mathbf{x}'|} - \frac{\exp(ik|\mathbf{x} - \mathbf{x}''|)}{4\pi|\mathbf{x} - \mathbf{x}''|}, \tag{9.20}$$

where \mathbf{x}'' is the mirror image of \mathbf{x}' .

Green’s theorem is a straightforward corollary of the divergence theorem,

$$\iiint_V (\phi \nabla^2 \psi - \psi \nabla^2 \phi) d^3x = \iint_S \left(\phi \frac{\partial \psi}{\partial n} - \psi \frac{\partial \phi}{\partial n} \right) dA, \tag{9.21}$$

where surface S encloses volume V . If we choose $\phi = G_+$, and make S the plane $z = 0$ plus a hemisphere off at infinity, then by applying the wave equation and the definition of G , we get the *Rayleigh–Sommerfeld integral*,

$$\psi(x) = \frac{1}{i\lambda} \iint_{z=0} \frac{\exp(ik|\mathbf{x} - \mathbf{x}'|)}{|\mathbf{x} - \mathbf{x}'|} \left(1 + \frac{i}{k|\mathbf{x} - \mathbf{x}'|} \right) \frac{\mathbf{n} \cdot (\mathbf{x} - \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} \psi(\mathbf{x}') d^2x'. \tag{9.22}$$

A limiting argument shows that the contribution from the hemisphere goes to 0.

If we choose G_0 instead, we get the *Kirchhoff integral*,

$$\psi(\mathbf{x}) = -\frac{1}{4\pi} \int_{z=0} \frac{\exp(ikR)}{R} \left[\nabla' \phi + ik \left(1 + \frac{i}{kR} \right) \frac{\mathbf{R}}{R} \psi \right] \cdot \mathbf{n}' dA'. \tag{9.23}$$

These scary-looking things actually turn out to be useful—we’ll revisit them in Section 9.3.6.

Ideally what we want is to find the exact plane wave spectrum of the light leaving S for a given plane wave coming in, because that makes it easy to do the propagation calculation. Getting the correct plane wave spectrum is easy for a planar screen, since different plane waves are orthogonal on a plane, and because we can use the correct Green’s function in the planar case (the Rayleigh–Sommerfeld theory). For more complicated boundaries, life gets very much harder since analytically known Green’s functions are rare, and plane waves are not then orthogonal on S , so we can’t just Fourier transform our way out of trouble. The surfaces of interest are usually spheres centered on some image point, so we’d need to expand in partial waves, and then find the plane wave spectrum from that. Fortunately, there’s an easier way.

Aside: Theory That’s Weak in the Knees. One problem for the outsider coming to learn optical systems design is that it’s a pretty closed world, and the connections between the scalar optics of lens design and the rest of optics are not clearly brought out in books on the subject, or at least those with which the present author is familiar—it isn’t at all obvious how a given ray intercept error influences the signal-to-noise ratio, for example. This is not helped by the uniformly inadequate presentation of the theoretical

underpinnings, which almost always base Fourier optics on the Fresnel approximation and aberration theory on a sloppy use of the Huyghens propagator.

A charitable interpretation of this is that it is an attempt to make the subject accessible to undergraduates who don't know anything about Green's functions. Yet it is unclear how they are aided by such sloppiness as defining the wavefront on a spherical reference surface near the exit pupil, then doing the integrals as though it were a plane.

Some claim that this is the Kirchhoff approximation (it isn't), and others unapologetically toss around the (paraxial) Huyghens integral on the spherical surface, even for large-aperture lenses. The funny thing about this is that, apart from neglect of obliquity, they get the right result, but for the wrong reasons. It matters, too, because the confusion at the root of the way the subject is taught damages our confidence in our results, which makes it harder to calculate system performance with assurance. If you're an optics student, ask lots of rude questions.

9.3.3 The Kirchhoff Approximation

Usually we have no independent way of measuring the actual fields on the boundary and are reduced to making a guess, based on the characteristics of the incoming wave. The *Kirchhoff approximation* says that, on surface S , the fields and their derivatives are the same as the incoming field in the unobstructed areas and 0 in the obstructed ones. This turns out to work reasonably well, except for underestimating the edge diffraction contribution (see Section 9.2.7). The degree of underestimate depends on the choice of propagator (see below); empirically the Kirchhoff propagator does a bit better on the edge diffraction contribution than the Rayleigh–Sommerfeld propagator.

You can find lots more on this in Stamnes, but the net is that these physical optics approximations work pretty well for imaging and for calculating diffraction patterns, but it won't get fine details right, for example, the exact ripple amplitude, and will underestimate the field in the geometric shadow regions. You need GTD or PTD to do that properly.

9.3.4 Plane Wave Spectrum of Diffracted Light

In Section 1.3, we used the Huyghens propagator, which in real space is

$$\Theta(x, y, z) = \frac{-i}{\lambda} \iint_P \Theta(x', y', z') \frac{\exp\left(ik \frac{(x-x')^2 + (y-y')^2}{2(z-z')}\right)}{(z-z')} dx' dy', \quad (9.24)$$

where P is the xy plane, and in \mathbf{k} -space is

$$\Theta(x, y, z) = \iint_{P'} \Theta(u, v)|_{z=0} e^{i(2\pi/\lambda)(ux+vy)} e^{-i(2\pi z/\lambda)(u^2+v^2)/2} du dv, \quad (9.25)$$

where P' is the uv plane.

If a beam gets cut off sharply, it scatters strong fringes out to high angles. Being a paraxial approximation, the Huyghens integral requires very large values of $z - z'$ to be used in that situation. The Rayleigh–Sommerfeld result (9.22) is the rigorously correct scalar solution for a correctly given $\psi(\mathbf{x})$ on the plane $z = 0$, because it is based on the correct Green's function for a half-space above that plane. To get the \mathbf{k} -space

representation (angular spectrum), we choose \mathbf{x} to be on the surface of a very large sphere of radius R , and neglect the constant term $-ie^{ikr}/R$, which yields

$$\psi(u, v) = \frac{w \operatorname{circ}(1-w)}{\lambda} \iint_P \exp\left(\frac{i2\pi}{\lambda}(ux' + vy')\right) \psi(x', y') dx' dy', \quad (9.26)$$

where u and v are the direction cosines in the x and y directions as before, $w = (1 - u^2 - v^2)^{1/2} = k_z/k$, and $\operatorname{circ}(x)$ is 1 for $0 \leq x < 1$ and 0 otherwise. It is clear from this equation that the \mathbf{k} -space solution is the Fourier transform of the fields on the boundary, multiplied by a factor of $-ik_z = 2\pi iw/\lambda = ik \cos \theta$, where θ is the angle of incidence of the outgoing light. A heuristic way of looking at this uses a pencil beam rather than a plane wave. A circular beam coming from a surface at an incidence angle of θ occupies an elliptical patch on the surface, whose area is $\pi a^2 \sec \theta$. On this patch, the field strength is not diminished by spreading out (different places on the long axis of the patch are seeing the same beam at different times), so the *obliquity factor* $w = \cos \theta$ is required to counteract the tendency of the integral to become large as the angle approaches grazing. (We saw this as the Jacobian in Section 9.2.1 and in the drinking-straw test of Section 1.3.12.)

The \mathbf{k} -space Kirchhoff integral is similar,

$$\psi(u, v) = \frac{(w_{\text{inc}} + w) \operatorname{circ}(1-w)}{2\lambda} \iint_{P'} \exp\left(\frac{i2\pi}{\lambda}(ux' + vy')\right) \psi(x', y') dx' dy', \quad (9.27)$$

which is just the same as the far-field Rayleigh–Sommerfeld integral except for the obliquity factor. The Neumann boundary condition case, where $\hat{\mathbf{n}} \cdot \nabla \psi$ is specified on the boundary, yields the same Fourier transform expression with an obliquity factor of w_{inc} . The three propagators are all exact since they predict the same fields if the source distribution is correct—they differ only when we make an inaccurate guess at $\phi(x, y)$.

9.3.5 Diffraction at High NA

Diffraction from apertures in plane screens can be calculated for all z by assuming that the field is the same as the incident field in the aperture, and zero elsewhere. In imaging problems, the screen has reflection or transmission amplitude and phase that depend on position. If we just take the incident field as our guess, we wind up suppressing the high-angle components by the obliquity factor (see Section 9.2.1), so in fact we have to put the reciprocal of the obliquity factor into the illumination beam in order for energy to be conserved (i.e., multiply by the Jacobian of the inverse transformation). This is physically very reasonable, since the screen could have a transmission coefficient of 1 (i.e., not be there at all), in which case the plane wave components had better propagate unaltered.

If the illumination beam has high NA, then the obliquity factors of the plane wave components of the illumination beam will be different, and that has to be taken into account. If the object has only low spatial frequencies, and doesn't have large phase variations due to topography, then each plane wave will be scattered through only a small angle, so that $\cos \theta$ doesn't change much, and the obliquity factors cancel out. This effect is partly responsible for the seemingly unaccountable success of Fourier optics at high NA.

As we discussed in Section 1.3.9, the simple thin object model used in diffraction imaging theory is a complex reflection coefficient, which depends on \mathbf{x} and not on \mathbf{k} . Height differences merely change the phase uniformly across the pupil. This works fine as long as the maximum phase difference across the pupil is smaller than a wave, i.e., we're within the depth of focus, and providing we take a weighted average of the phase shift over the entire pupil, i.e., the phase shift with defocus isn't $k_z z$ anymore (see Example 9.4).

9.3.6 Propagating from a Pupil to an Image

We're now in a position to say what the exact scalar field propagator is between a pupil and an image. Consider the exit pupil plane of an optical system, with a mildly wrinkled wavefront that is basically a spherical wave centered on the nominal image point \mathbf{x}_0 ,

$$\psi(\mathbf{x}') = \tilde{A}(\mathbf{x}') \frac{e^{-ik|\mathbf{x}' - \mathbf{x}_0|}}{|\mathbf{x}' - \mathbf{x}_0|}, \quad (9.28)$$

where the *pupil function* \tilde{A} is a complex envelope that carries the amplitude and phase information we care about. (In a little while it will be apparent that the natural variables for expressing \tilde{A} are the direction cosines u and v , just as in the paraxial theory.) We're interested in the structure of the image, so we use the Rayleigh–Sommerfeld integral to propagate to $\mathbf{x}_1 = \mathbf{x}_0 + \boldsymbol{\zeta}$, where $|\boldsymbol{\zeta}|$ is assumed to be small compared to $|\mathbf{x}' - \mathbf{x}_0|$. We further assume that $1/(k|\mathbf{x}' - \mathbf{x}|) \ll 1$, that is, we're in the limit of large Fresnel number, which allows us to discard that term (which turns out to represent the evanescent fields and pupil edge diffraction), so we write (where P' is the uv plane as before)

$$\psi(\mathbf{x}) = \frac{1}{i\lambda} \iint_{P'} \frac{\exp(ik|\mathbf{x} - \mathbf{x}'|)}{|\mathbf{x} - \mathbf{x}'|} \frac{\exp(-ik|\mathbf{x}_0 - \mathbf{x}'|)}{|\mathbf{x}_0 - \mathbf{x}'|} \tilde{A}(\mathbf{x}') \frac{\mathbf{n} \cdot (\mathbf{x} - \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d^2x'. \quad (9.29)$$

Note that we haven't made any assumptions about small angles or slowly varying envelopes—apart from the scalar field and discarding the evanescent contributions, this is an exact result. Providing that $\boldsymbol{\zeta}$ is small compared to $|\mathbf{x} - \mathbf{x}'|$, we can ignore it in the denominator, but since it isn't necessarily small compared to $1/k$, we have to keep it in the exponent. Doing a third-order binomial expansion of the exponent, we get

$$|\mathbf{x}_0 + \boldsymbol{\zeta} - \mathbf{x}'| - |\mathbf{x}_0 - \mathbf{x}'| = \boldsymbol{\zeta} \cdot \frac{\mathbf{x}_0 - \mathbf{x}'}{|\mathbf{x}_0 - \mathbf{x}'|} + \boldsymbol{\zeta} \cdot \frac{\left(\boldsymbol{\zeta} - \frac{(\mathbf{x}_0 - \mathbf{x}')(\boldsymbol{\zeta} \cdot (\mathbf{x}_0 - \mathbf{x}'))}{|\mathbf{x}_0 - \mathbf{x}'|^2} \right)}{2|\mathbf{x}_0 - \mathbf{x}'|} + O(\zeta^3). \quad (9.30)$$

The first term is the phase along the radial vector, which as usual is going to turn into the kernel of a Fourier transform; the second is the phase due to the length of the vector changing. (Note that there is no radial component of the phase in order ζ^2 .) If we were in the paraxial case, we'd just forget about terms like that, or at most say that the focal surface was a sphere centered on \mathbf{x}' , but the whole point of this discussion is that $\mathbf{x}' - \mathbf{x}_0$ be allowed fractional variations of order 1, so we can't do that.

What we do need to do is restrict ζ . In order for the ζ^2 term to be small compared to $1/k$, it is sufficient that

$$|\zeta| \ll \sqrt{\frac{\lambda |\mathbf{x}' - \mathbf{x}_0|}{\pi}}. \quad (9.31)$$

Since we imagine that the pupil function has been constructed by some imaging system, the rays have been bent so as to construct the spherical wavefront. For consistency, we must thus put in the inverse of the obliquity factor, and the $\hat{\mathbf{n}} \cdot (\mathbf{x}_0 - \mathbf{x}')$ term then goes away to the same order of approximation as neglecting ζ in the denominator.[†] We also transform into direction cosines, so that $(dx, dy) = |\mathbf{x}_0 - \mathbf{x}'|(du, dv)$, which leaves a pure inverse Fourier transform,

$$\psi(\mathbf{x}) = \frac{1}{\lambda} \iint_P e^{i\mathbf{k}\mathbf{u}\cdot\zeta} \tilde{A}(\mathbf{u}) du dv. \quad (9.32)$$

For a pupil-image distance of 20 mm and a wavelength of 0.5 μm , this Fraunhofer-type approximation is valid in a sphere of at least 100 μm in diameter, *even at* $NA = 1$. In order to cause the image to deviate seriously from the Fourier transform of the pupil function, there would have to be hundreds of waves of aberration across the pupil, so that for all interesting cases in imaging, where the scalar approximation applies, Fourier optics remains valid. This applies locally, in what is called the *isoplanatic patch*, and does not imply that the whole focal plane is the Fourier transform of the whole pupil, as it is in the paraxial case, because that depends on things like the field curvature and distortion of the lens, and the different obliquities at different field positions.

This analysis applies backwards as well, in going from a small-diameter object to the entrance pupil of the lens, although if the object is a material surface and not itself an aerial image, the usual thin-object cautions apply. The combination of the two shows that an object point is imaged to an image point, and that the point spread function of the system is the Fourier transform of the pupil function \tilde{A} , at least in the large Fresnel number limit.

This is really the main point: the imaging of an object point into an image point via a pupil is controlled by Fourier optics in all cases, and for an imaging system faithful enough to deserve the name, the amplitude PSF of the imaging operation really is the Fourier transform of the pupil function \tilde{A} , regardless of NA.

Example 9.1: High-NA Fourier Optics—Metal Lines on Silicon at $NA = 0.95$.

Figures 9.4 and 9.5 show the Fourier optics result versus experiment for a 90 nm tall line of gold on silicon. Even though the scalar Fourier optics approximation to high-NA imaging is a fairly sleazy one, it nevertheless works extremely well in practice.

Example 9.2: When Is the Paraxial Approximation Valid? The special case of a perturbed spherical wave is very important in applications, but the usual Fourier optics result is more general; the far-field pattern is the Fourier transform of the pupil function. What is the range of validity of that approximation?

[†]For a system of unit magnification, this cancellation is exact when both the object-to-pupil and pupil-to-image transforms are computed; when the magnification is not 1, the pupil function \tilde{A} will need some patching up, but that's not a fundamental objection at this point.




Figure 9.4. Heterodyne microscope image of Au lines on Si, 515 nm, 0.90 NA: amplitude.




Figure 9.5. Au lines on Si: phase.

Comparison of the Huyghens integral with the Kirchhoff and Rayleigh–Sommerfeld ones shows two differences: the Huyghens integral omits the obliquity factor, and for a plane wave component $\exp(i2\pi(ux + vy)/\lambda)$, the Huyghens integral replaces the true phase shift $kz(w - 1)$ by the first term in its binomial expansion, apparently limiting its use to applications where this causes a phase error of much less than 1 (it is not enough that it be much less than kz because it appears in an exponent). If we require that the next term in the binomial expansion be much less than 1, we find that this requires that

$$|z| \gg \sqrt[3]{\frac{x^4}{\lambda}}. \tag{9.33}$$

This restriction is necessary for general fields, but *not* for paraxial ones. The slowly varying envelope equation is

$$\frac{d^2\Theta}{dx^2} + \frac{d^2\Theta}{dy^2} + 2ik\frac{d\Theta}{dz} = 0. \quad (9.34)$$

Its validity depends solely on the initial conditions; a sufficiently slowly varying envelope will be accurately described by this equation for all z . For slowly varying Θ and small $z - z'$, the error in the phase term does indeed become large, but a stationary phase analysis shows that the large- x contribution to the integral goes strongly to zero as $z \rightarrow z'$, due to the rapidly varying phase factor, so that the integral remains valid for all z , and the Huyghens integral is not limited to far-field applications. This is perhaps easier to see in the spatial frequency representation.

If we take $\Theta_\alpha(\mathbf{x}) = e^{i\alpha x} e^{i\gamma z}$ in (9.34), requiring the phase error to be small compared to 1 leads to (9.33) for fixed α of order k , and an absolute phase error that grows secularly with z , as one would expect. This is not a deadly error, as it amounts only to a deviation of the field curvature from spherical to parabolic; if we take as our reference surface a parabola instead of a sphere, it goes away; it may make the calculated optical path incorrect, and in applications where that matters, it should be checked by comparison with the Rayleigh–Sommerfeld result.

For fixed z , the restriction can be applied to α instead:

$$|\alpha| \ll \sqrt[4]{\frac{k^3}{z}}. \quad (9.35)$$

This is easily satisfied for small z as well as large.

9.3.7 Telecentricity

As Figure 9.6 illustrates, a telecentric optical system is one in which the principal ray is parallel to the optical axis. This means that, roughly speaking, the axis of the cone of light arriving at the image or leaving the object is not tilted, and is equivalent to saying that the pupil is at infinity. An optical system can be telecentric in the object space, the image space, or both.




Figure 9.6. A telecentric optical system.

This property is of more practical interest than it may sound. In a telecentric system, tilting the sample or moving it in and out for focusing does not change the magnification—the image is an orthographic projection, like an engineering drawing. Light reflected from a plane sample, such as a microscope slide or a flat mirror, retraces its path. Both of these properties are very useful for scanning or imaging interferometers such as shearing interference microscopes.

In a telecentric imaging system with telecentric illumination, the illumination diagram is independent of position; all points in the field of view are imaged with light covering the same range of angles. Because both the illumination and collection NAs are constant, the range of spatial frequencies received is the same everywhere too. These two properties together give telecentric systems nearly space-invariant point spread functions. This is of great benefit when interpreting or postprocessing images, for example, in automatic inspection systems. Obviously a telecentric system can have a field of view no larger than its objective (the last optical element on the outside), and usually it's significantly smaller.

9.3.8 Stereoscopy

Stereoscopic vision requires the ability to look at a scene from two different directions and synthesize the resulting images. This is different from merely binocular vision. A binocular microscope presents the same image to each eye, whereas a properly stereoscopic microscope splits the pupil into two halves, presenting one half to each eye. Since pupil position corresponds to viewing angle, this reproduces the stereo effect. Splitting the pupil reduces the resolution, but the gain in intuitive understanding is well worth it.

9.3.9 The Importance of the Pupil Function

Pupil functions don't get the respect they deserve. The point spread function $h(\mathbf{x})$ of an optical system is the Fourier transform of the pupil function $\hat{A}(\mathbf{u})$, and the point spread function ranks with the étendue as one of the two most important attributes of an imaging system. The pupil function is the filter that is applied to the spatial frequency spectrum of the sample to generate the image.

In signal processing, we choose our filter functions very carefully, so as to get the best measurement, but this is less often done in optics, which is odd since optics are much more expensive. One reason for it is confusion of two quite different objects, both called *transfer functions*, and both giving rise to *point spread functions* (PSFs). The confusion has arisen because, for historical reasons, the one less clearly connected to the electromagnetic field quantities \mathbf{E} and \mathbf{B} has staked out the high ground.

9.3.10 Coherent Transfer Functions

When we describe the actions of an optical system in terms of the plane wave decomposition of the scalar optical field E , and apply Fourier transform theory to describe how a sinusoidal component of the field distribution at a sample plane propagates to the image plane, we are using the *coherent transfer function* (CTF) of the system. The CTF is the convolution of the illumination and detection pupil functions, because the amplitude PSF of the measurement is the product of the illumination and detection PSFs. Most of the

time, one of the two is far more selective than the other, so the CTF and the broader of the two pupil functions are often interchangeable.

The CTF is the right description of a translationally invariant phase-sensitive optical system; this class includes holography setups, scanning heterodyne microscopes, and phase shifting imaging interferometers, as well as any system producing an aerial image, such as binoculars. To determine the output of such a system, multiply the Fourier transform of the sample's complex reflection coefficient, as a function of position, by the instrument's CTF, and take the inverse transform. This is of course mathematically equivalent to convolving the sample function with the instrument's 2D amplitude point spread function. Since the optical phase information is preserved, digital postprocessing can be used to transform the complex fields in a great variety of ways.

The net effect is that provided you measure both phase and amplitude, on a sufficiently fine grid and at sufficiently high SNR, you can do with postprocessing anything you could do on an aerial image with optical elements; this is a remarkable and powerful result.

Example 9.3: Heterodyne Microscope. A heterodyne microscope is basically a heterodyne Michelson interferometer, using an AO deflector as its beamsplitter, and with a microscope objective in one or both arms. Some versions use separate lenses, and some send both beams down to the sample through the same lens, as shown in Figure 9.7. A uniform pupil has an illumination pupil function $L = \text{circ}((u^2 + v^2)/\text{NA}^2)$, which transforms to an illumination PSF of

$$l(\chi) = \frac{J_1(\pi\chi)}{\chi}, \tag{9.36}$$

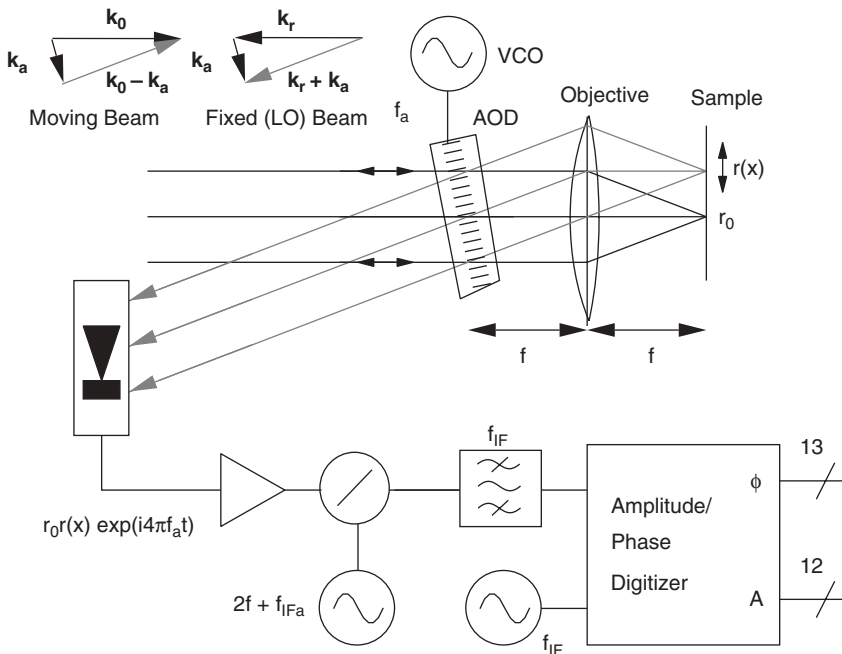


Figure 9.7. Heterodyne confocal microscope.

where $\chi = r\text{NA}/\lambda$. The coherent detector uses interference with a nominally identical beam $s(\chi)$ to produce the AC photocurrent

$$\begin{aligned} i_{\text{AC}} &= 2\mathcal{R} \operatorname{Re} \left\{ \iint_{\text{det}} d^2x \psi_{\text{LO}} \psi_s^* \right\} \\ &= 2\mathcal{R} \operatorname{Re} \left\{ \exp(-i \Delta\omega t) \iint_{\text{det}} |\psi_{\text{LO}}(\mathbf{x})| |\psi_s(\mathbf{x})| \exp(i \Delta\phi(\mathbf{x}, t)) dA \right\}, \end{aligned} \quad (9.37)$$

as in Section 1.5. By the power theorem, this dot product can be computed in the pupil or the image, or anywhere in between. For our purposes, it is easiest to see what happens if we compute it at the sample surface. There, the two jinc functions are superimposed and multiplied by the local complex reflection coefficient \tilde{r} of the sample S . Thus the total complex AC photocurrent is

$$\tilde{i}_{\text{AC}} = \mathcal{R} \iint_{\text{sample}} l(\mathbf{x}) \tilde{r}(\mathbf{x}) s^*(\mathbf{x}) d^2x, \quad (9.38)$$

which if both beams are unaberrated and focused on \mathbf{x} is

$$\tilde{i}_{\text{AC}}(\mathbf{x}) = \mathcal{R} \lambda \iint_S \tilde{r}(\mathbf{x}') \left[\frac{J_1((\pi \text{NA}/\lambda)|\mathbf{x} - \mathbf{x}'|)}{|\mathbf{x} - \mathbf{x}'| \text{NA}} \right]^2 d^2x', \quad (9.39)$$

so by construction, the amplitude PSF of such a microscope is

$$g(\mathbf{x}) = \left[\frac{J_1(\pi \chi)}{\chi} \right]^2. \quad (9.40)$$

The CTF of this microscope is the *Chinese hat function*,

$$H(\omega) = \frac{2}{\pi} \left(\cos^{-1}(\omega) - \omega \sqrt{1 - \omega^2} \right), \quad (9.41)$$

whose name makes perfect sense if you look at Figure 9.8 and remember that it's cylindrically symmetric. This function has a cusp at 0 and extends out to $\omega = 2\text{NA}$. In talking about the CTF, we're subtly sliding into the thin-object Fourier optics approximation, where a spatial frequency component at $v = 2\text{NA}/\lambda$ scatters light coming in at u all the way to $-u$, which can still just make it back through the pupil.

The line spread function is

$$\begin{aligned} l_2(x) &= \frac{8\pi \text{NA}}{\xi^2} \mathbf{H}_1(\xi) \\ &= \frac{16\text{NA}}{\pi} \sum_{m=0}^{\infty} \frac{(-1)^m \xi^{2m}}{(2m+1)!!(2m+3)!!} \end{aligned} \quad (9.42)$$




Figure 9.8. CTFs of a heterodyne interference microscope before and after Fourier postprocessing.

where $\xi = 2kxNA$, and $\mathbf{H}_1(x)$ is the Struve function of order 1 (see Abramowitz and Stegun). This function has an asymptotic series for large x ,

$$l_2(x) \sim \frac{16NA}{\pi} \xi^{-2} - \frac{8NA(\cos \xi + \sin \xi)}{\sqrt{\pi} \xi^5} + O(\xi^{-4}) \text{ monotonic} + O(\xi^{-4.5}) \text{ oscillatory}, \tag{9.43}$$

which is a distressingly slow falloff. The slowness is due to the cusp at the origin and the higher order nondifferentiability at the outer edges. Because the heterodyne system preserves phase information, this cusp can be removed by digital filtering in a postprocessing step (see Section 17.7.1). Even a very gentle filter can make a big difference to the settling behavior; for example, $F(u) = \cos^2(u/NA)/G(u)$, which turns the Chinese hat function into a von Hann raised cosine (see Section 17.4.9). This filter removes the cusp and makes the edges go to 0 quadratically, and as Figures 9.8–9.10 show, the step response settles at its final value when the uncorrected LSF is still 5% away. It does this with essentially 0 dB noise gain, so there’s no penalty whatever.

A different filter, which boosts the wings of the transfer function further, can yield results that approach those expected from a microscope operating at half the wavelength, provided the noise is sufficiently low[†]; the 10–90% edge rise going over a $\lambda/6$ phase step can be reduced from 0.45λ to 0.19λ , that is, 90 nm for $\lambda = 514$ nm. (Phase edges are a bit sharper than pure amplitude ones, since the amplitude dip in the middle sharpens the edge; the pictures in Figures 9.4 and 9.5 were preternaturally sharp because the step happened to be close to $\lambda/4$ tall.) Since an optical image is one of the few real examples of a band-limited function (because waves with spatial frequency higher than $1/\lambda$ cannot propagate to the detector), this is as much as can be achieved in a model-independent fashion.

[†]P. C. D. Hobbs and G. S. Kino, Generalizing the confocal microscope via heterodyne interferometry and digital filtering. *J. Microsc.* **160**(3) 245–264 (December 1990).

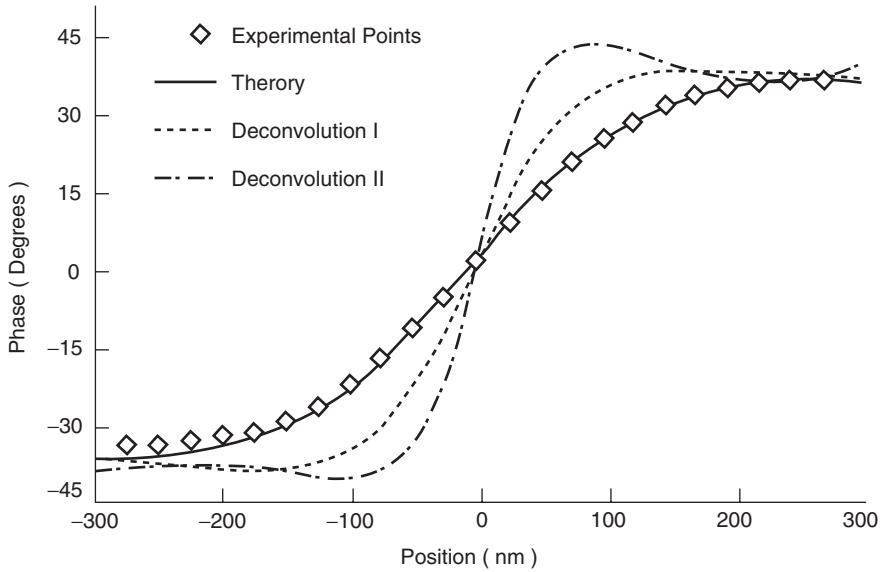


Figure 9.9. Experimental and theoretical phase plots for a heterodyne confocal microscope looking at an aluminum-on-aluminum step, 80 nm tall, before and after deconvolution.




Figure 9.10. Theoretical step response of a heterodyne confocal microscope to an amplitude step, and several deconvolutions.

Example 9.4: Modeling Defocus. Another way to illustrate the difference between coherent and incoherent optical systems, and the value of coherent postprocessing, is the compensation of defocus. In an incoherent system, it is impossible to distinguish between positive and negative defocus, because (apart from aberrations) the difference is only the sign of the phase shift, which gives rise to no intensity change. Although there are minor differences in the behavior of lenses depending on the direction of the defocus, this does not change the basic point.

In a coherent system, on the other hand (provided that we measure the amplitude and phase independently, with adequate spatial resolution and signal-to-noise ratio), we can numerically refocus an image, or focus on any given depth. This can be done as follows: decompose the beam into plane waves, multiply by $\exp(-ikz\sqrt{1-u^2-v^2})$, where u and v are the x and y direction cosines as usual, then put it back together. This is a kind of convolution filter.

The author's former colleagues, Paul Reinholdtsen and Pierre Khuri-Yakub, used this idea with a confocal acoustic microscope to remove blurring caused by out-of-focus structures, by numerically defocusing an in-focus image of the interfering top surface. Looking at a quarter, they were able to read QUARTER DOLLAR on the back, right through the coin, by defocusing George Washington and subtracting him out.

Performing the convolution and taking the real part, we can get the (complex) vertical response of a confocal reflection microscope (where the phase shift is doubled):

$$\tilde{i}(z) = 2\pi \int_0^{\text{NA}} d\omega \exp\left(-i2kz\sqrt{1-\omega^2}\right). \quad (9.44)$$

Here we've assumed that the pupil function is uniform, so that the obliquity factors in transmit and receive cancel out exactly, and that the medium is air. With a change of variable from $\omega = \sin\theta$ to $r = \cos\theta$, this becomes

$$\tilde{i}(z) = 2\pi \int_r^1 \exp(-i2kzr')r' dr'. \quad (9.45)$$

This is easily done by partial integration, but the result is a mess. We can get a good memorable result accurate to about 0.2% up to $\text{NA} = 0.5$ by setting the factor of r' outside the exponent to 1 and computing the envelope and carrier:

$$\tilde{i}(z) = 2\pi(1-r)\text{sinc}(rz/\lambda) \exp(-i2krz), \quad (9.46)$$

that is, the amplitude response is a sinc function and the phase shift is not $2kz$ but is reduced by a factor $[1 - (\text{NA})^2]^{1/2}$. The exact result shows that the phase slope reduction reaches a factor of 2 at $\text{NA} = 1$.

9.3.11 Optical Transfer Functions

The CTF is not the most commonly encountered transfer function in the literature. The more usual *optical transfer function (OTF)* is another beast altogether, and it's important to keep the distinction crystal clear. The OTF predicts the *intensity* distribution of the image based on that of the sample, with certain assumptions about the spatial coherence of the illuminator, i.e., the statistical phase relationships between the various Fourier components. There is no 1:1 correspondence to the propagation of plane wave components through the optical system. As we'll see, the OTF isn't a proper transfer function.

The intensity[†] of the beam is $\psi\psi^* \cos \theta$. Since the propagation of \tilde{A} to the image plane is governed by the CTF H , the autocorrelation theorem gives us the OTF O :

$$O(u, v) = H(u, v) \star H(u, v). \quad (9.47)$$

The OTF for an ideal system whose pupil function is $\text{circ}[(u^2 + v^2)/(\text{NA})^2]$ is our old friend the Chinese hat; the circ function is real and symmetric, so its transform is real and symmetric, and therefore its self-convolution equals its autocorrelation. (This is only true in focus, of course.)

The OTF and CTF each presuppose a concept of spatial frequency, but it must be understood that these two concepts do not map into each other in a simple way. Intensity is related to the squared modulus of the field variables; this nonlinearity results in the field amplitude spatial frequencies of the CTF undergoing large-scale intermodulation and distortion in the process of becoming the optical intensity spatial frequencies of the OTF. In particular, the width of the OTF is twice that of the CTF, but that does not imply the ability to resolve objects half the size. In discussing the OTF, we still use the variable names u and v , but do be aware that they no longer correspond directly to pupil plane coordinates, nor to the direction cosines of the plane wave components of ψ . (This is another example of a problem that's endemic in optics: reusing nomenclature in a confusing way.)

Being autocorrelations, optical transfer functions always droop at high spatial frequencies, and since intensity is nonnegative, OTFs must always have a maximum at zero. Interestingly, the OTF can go negative at intermediate values of spatial frequency, leading to contrast inversion for objects with periodicities falling in that region, an effect called *spurious resolution*. The OTF is purely real for symmetric optical systems but can exhibit phase shifts in systems lacking an axis of symmetry.

The justification for centering on the OTF is that, with thermal light, the phases of image points separated by more than a couple of spot diameters are uncorrelated, so there is no utility in keeping the phase information. This is of course fine if an in-focus image is being projected on film or an intensity detector, which is intrinsically insensitive to optical phase, but is inadequate for an aerial image or a phase-preserving system like a phase shifting interferometer or a laser heterodyne system, where the phase information still exists and can be very important, not least in fixing the imperfections of the image.

Perhaps the most intuitive way of capturing the distinction is that the OTF is not changed by putting a ground-glass screen at the image, whereas the CTF has its phase scrambled. Classical lens and optical systems designers use the OTF as one of their primary tools, which explains some of the difficulty encountered by workers in the different fields when they talk to each other.

Aside: Nonuniqueness of the Intensity Pattern. Since the relative phases of the plane waves in the CTF are lost when going to the OTF, any two patterns whose fields differ only in phase will produce the same intensity pattern, for example, positive and negative defocus.

[†]Well, irradiance, to be exact—the author resists the Humpty-Dumpty approach to radiometric nomenclature that goes around redefining preexisting terms to mean something else, in this case *intensity* being used for total power per steradian through some surface.

9.3.12 Shortcomings of the OTF Concept

The classical formulation of the optical transfer function is not a good analogue to transfer functions as used in circuit theory, ordinary differential equations, and so forth, although it might superficially look like it.

The behavior of fields is much more intuitive than that of image irradiance, because the fields exist throughout the optical system, whereas image irradiance doesn't. There are other ways in which the OTF isn't really a transfer function, the most important one being that you can't compute the OTF of two systems in cascade by simply multiplying the individual OTFs.

For example, consider a 1:1 relay system consisting of two lenses of focal length f , spaced $4f$ apart, as in Figure 12.1a. With an object at $-2f$ from the first lens, there will be a good image at the center of the system and another one at $2f$ past the second lens. If we choose the reference plane for the individual OTFs to be the center, everything works reasonably well. On the other hand, if we choose it to be off center, the image at that plane will be out of focus, leading to an ugly OTF, falling off very rapidly from zero spatial frequency. The second lens will also be defocused, leading to another ugly OTF, so their product will be ugly squared. This is exactly the right answer, *provided we put a diffuser at the reference plane*.

In real life, of course, an odd choice of reference plane doesn't affect the system operation at all—the defocus of the first half is undone by the defocus of the second, leading to a good image. The OTF gets this wrong, but the CTF gets it right—the phase curvatures of the two CTFs compensate correctly, and you get the right answer.

Lest anyone say that this is just silly, that nobody would set up a calculation that way, let's go a bit deeper into the problem. A symmetric optical system such as this 1:1 relay has no odd-order wave aberrations, because the second half's aberrations cancel out the first half's. (The even orders add.) Computing the overall OTF by multiplying the two half-OTFs will get this wrong, because the phase information is lost, so all the aberrations add in RMS instead of directly. Odd-order contributions will be overestimated, and even-order ones underestimated. Yet this weird OTF thing is called “the transfer function” and tossed about as though it had physical meaning. Beware.

9.3.13 Modulation Transfer Function

The modulation transfer function (MTF) is the magnitude of the OTF, normalized to unity at zero spatial frequency, and is most commonly used to describe the resolution performance of lenses, while not considering their photon efficiency.

9.3.14 Cascading Optical Systems

Under appropriate assumptions, when two optical systems are cascaded, their transfer functions are multiplied to get the transfer function of the cascade. If there is a diffuser, image intensifier, television system, or other phase-randomizing device between the two, use the OTF or MTF. Otherwise, use the CTF.

9.3.15 Which Transfer Function Should I Use?

This depends on the properties of the illuminator, and to a lesser degree on those of the detector. The assumptions leading to the derivation of the OTF are: an illuminator

with very low spatial coherence, and a detector that is sensitive only to intensity, such as a television camera or photodiode, with no phase reference (as in an interferometer). The resulting near-total loss of phase information severely limits the opportunities to gain from postprocessing, although the work of Fienup and others has demonstrated that some phase information can often be retrieved.

Example 9.5: OTF of an Ideal CCD Camera. As an example of the use of the OTF, consider a CCD camera with square pixels of pitch δ , a 100% fill factor, $QE = 1$ everywhere, and negligible bleed of one pixel into its neighbor. This is a spatial analogue of the sampled-data systems we'll encounter in Section 17.4.3, so although the detector is not shift invariant, we lose no information about the true OTF as long as the pixel pitch obeys the Nyquist criterion, and it is still sensible to talk about the OTF and MTF of such a system. The detector sensitivity pattern is $\text{rect}(x/\delta) \text{rect}(y/\delta)$, which is unaltered by squaring. Since u and x/λ are the conjugate variables, the detector CTF is the product of x and y sinc functions scaled by δ/λ , and its OTF is the same, so the OTF of the lens/CCD system is

$$\text{OTF}_{\text{tot}}(u, v) = \text{OTF}_{\text{lens}}(u, v) \left(\frac{\delta}{\lambda}\right)^2 \text{sinc}\left(\frac{u\delta}{\lambda}\right) \text{sinc}\left(\frac{v\delta}{\lambda}\right). \quad (9.48)$$

(We couldn't use this detector coherently without an LO beam, of course, so we can think of the spatial filtering action corresponding to this CTF as occurring on a Gaussian surface just above the detector.)

9.4 ABERRATIONS

A lens is often thought of as imaging an object plane on an image plane, but really it images a volume into another volume. A perfect imaging system would image every point in its object space to a corresponding point in its image space. The fidelity of the image would be limited only by diffraction, and in the transverse direction, it would be perfectly faithful geometrically as well. Squares would come out square (no distortion or anamorphic errors), and a flat object would give rise to a flat image (no field curvature), but unless the magnification was unity, the longitudinal magnification would nonetheless differ from the transverse magnification. Paraxial theory, whether the ray model (as in *ABCD* matrices) or the field model, always predicts perfect imaging, apart from defocus. We therefore expect the aberrations to turn up in the higher order terms.

Unfortunately, the algebra gets ugly in a hurry when we're dealing with exact ray tracing or scalar wave propagation; there are lots of square roots. Good quality optical systems ought to have phase aberrations that are small compared to the total phase delay through the system, so we anticipate that a power series expansion will yield useful simplifications. This power series is really not that well defined, because higher orders yield higher spatial frequency information that will eventually be corrupted by edge diffraction and vignetting, so that the aberration series is really a high-order polynomial plus some hard-to-treat residual, which we will assume is small.[†]

[†]This is rather like the distortion polynomial of Section 13.5.

Nobody uses high-order analytical aberration theory anymore. Lens designers use the low-order terms as conveniences, but rely on computer ray tracing and (manually guided) numerical optimization of an intelligently chosen starting configuration. For the system designer, high-order aberrations are of peripheral concern as well.

Aberrations are most obtrusive in wide field, high-NA optics, such as lithographic lenses and fast telescopes. Lots of instruments use lenses like that, but they are seldom fully custom designs, because the engineering and tooling costs would be astronomical. Thus the heroic lens design is someone else's problem—the rest of us mostly live in the low-NA, narrow field region behind those fancy lenses. Instrument designers need to know how aberrations propagate, what produces them, and how to avoid doing it. For this use, the lowest order aberrations are generally enough. For the same reason, we'll ignore the pure ray picture entirely and center on phase shifts of the plane wave components of a focused spot at an arbitrary field position.

9.4.1 Aberration Nomenclature

Aberration theory is somewhat separate from the rest of optics, because it is primarily used by lens designers, who have been doing much the same sort of work for 100 years, in sharp distinction to workers in most of the rest of optics. This is not to disparage the great strides lens design has made in that time, but it remains true that the formalism of classical aberration theory is not clearly related to the rest of the optical world, and a number of the terms used have different meanings than elsewhere. For example, to most practical optics folk, *defocus* means that the focal plane is offset from where it should be. In the paraxial approximation, a defocus d translates to a quadratic phase (time delay) across the pupil,

$$\Delta t_{\text{defocus}} \approx \frac{nd}{c} \left(1 - \frac{u^2}{2} \right), \quad (9.49)$$

whereas the real phase delay is $k_z z$, which in time delay terms is

$$\Delta t_{\text{defocus}} = \frac{nd}{c} \cos \theta = \frac{nd}{c} \sqrt{1 - u^2}, \quad (9.50)$$

which of course contains all even orders in u .

In wave aberration theory, *defocus* means the quadratic expression (9.49), *even at large NA*. A pure focus shift in a large-NA system thus comes out as aberrations of all even orders, including spherical aberration and so on, even though a twist of the focus knob will restore the image completely. Those of us who use physical intuition heavily must guard against being led astray by this sort of thing.

Aberrations with the same name in the two different pictures do not correspond uniquely to one another; we've already seen the problem with defocus, but it also exists in other places. The names of aberrations have only mnemonic value—once again, if you expect everything to make sense together, you'll wind up chasing your tail.

As a way of connecting aberration theory with ordinary experience, let's calculate the effects of introducing a plane-parallel slab of dielectric into a perfect, converging spherical wave of limited NA.




Figure 9.11. A plane-parallel slab of dielectric introduced into a plane wave.

9.4.2 Aberrations of Windows

Figure 9.11 shows the \mathbf{k} vector of a plane wave incident on a plane-parallel slab of dielectric constant n_2 . The refracted wave travels farther, and through a different index material. In the wave picture, this one is easy; the phase shift is just $(k_{z2} - k_{z1})d$. Let's use the hybrid picture, where we calculate the phase difference along the ray paths. Inspection of the figure shows that the change in the propagation time due to the presence of the slab is

$$\Delta t = \frac{d}{c} \sec \theta_2 (n_2 - n_1 \cos(\theta_1 - \theta_2)), \quad (9.51)$$

since a translation perpendicular to \mathbf{k} has no effect on a plane wave. (If this isn't obvious, it's worth spending a bit of time on. This is a move that the hybrid picture relies on a good deal.) Without loss of generality, if the slab's faces are parallel to the (x, y) plane, and the incident plane wave has direction cosines $(u_1, 0)$, then Snell's law requires that $u_2 = n_1 u_1 / n_2$ (we've used $u_1 = \sin \theta_1$). Writing (9.51) in terms of u_1 , we get

$$\Delta t = \frac{d}{c} \left[n_2 \sqrt{1 - \left(\frac{n_1 u_1}{n_2} \right)^2} - n_1 \sqrt{1 - u_1^2} \right], \quad (9.52)$$

which (comfortingly enough) is the same as the wave result. This obviously has terms of all even orders in u_1 . Let's look at the first three orders, Δt_0 to Δt_4 :

$$\Delta t_0 = \frac{d}{c} (n_2 - n_1), \quad \Delta t_2 = \frac{d}{c} \frac{u_1^2}{2} (n_1 - n_1^2 / n_2), \quad \Delta t_4 = \frac{d}{c} \frac{u_1^4}{8} (n_1 - n_1^4 / n_2^3) \quad (9.53)$$

The higher order terms in the series expansion are equally simple, and all have the same form. Specializing to a 6 mm thick plate and BK7 glass ($n_d = 1.517$), we get the result shown in Figure 9.12. The bulk of this is obviously caused by the zero-order time delay and the focus shift, but considering that one cycle of green light takes only 1.8 fs even the residuals may be very large (the right-hand axis goes up to 12,000 waves). We

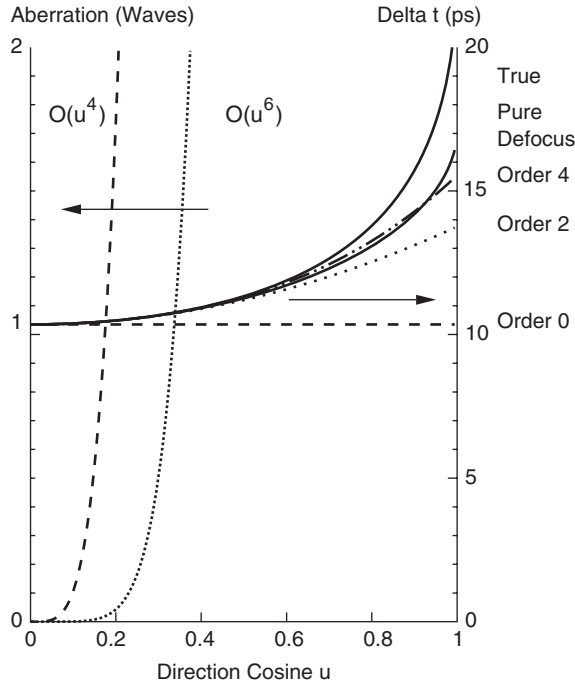


Figure 9.12. Differential time delay t suffered by a plane wave on passing through the dielectric plate of Figure 9.10, together with approximations of orders 0, 2, and 4. The curves at left are aberration residuals up to orders 2 and 4.

find out the effects of the aberrations on our nice focused beam by using the delay as a function of u and v to construct a phase factor to multiply our pupil function \tilde{A} , and then using (9.32) to get the new point spread function.

As we’ve already discussed, the u_1^2 term is usually called simply “defocus,” though Figure 9.12 shows up the clear distinction between this and real defocus; to keep things clear, we’ll call the quadratic term *paraxial defocus*. The term in u_1^4 is called *primary spherical aberration*. Spherical aberration is independent of field position, and so it occurs even for an on-axis beam.

The curves on the left show the aberrations of fourth order and greater, and of sixth order and greater, in units of waves (i.e., cycles) for 600 THz (500 nm) light, assuming that the pupil function of the beam is symmetric around $u = 0$. If that isn’t true, for example, a low-NA beam coming in at a big field angle, the true aberration is v times the spread of time delays across the pupil function, minus the best-fit defocus.

9.4.3 Broken Symmetry and Oblique Aberrations

Other aberrations, such as astigmatism and coma, show up as soon as we run a finite field angle, that is, move the center of the pupil function distribution away from $(0, 0)$. In an axisymmetric system like this plate or most lenses, these *oblique aberrations* are purely an effect of a shift of origin. (If the plate had some wedge angle, that would no longer be true.)

A residual effect of this broken symmetry is that if we move the origin to $(u_0, 0)$ (which loses us no generality), the pupil function is an even function of v . Thus the aberrations of a really symmetric system depend only on even powers of v , and by appropriate rearrangement of terms, that means they depend only on the cosine of the azimuthal angle θ ($u = \rho \cos \theta$, $v = \rho \sin \theta$). Manufacturing and assembly errors are in general asymmetrical and are frequently of the same order as the design residuals, so don't make too much of it.

If we move the center of the pupil function to $(u_0, 0)$, we're calculating the fields at a point $\mathbf{x} = (u_0 f / (1 - u_0^2)^{1/2}, 0, f)$, where L is the z distance from the pupil to the focus. For simplicity, we'll call this x coordinate the height h . The aberration polynomial coefficients get a tiny bit more complicated,

$$\Delta t_0 = \frac{d}{c}(a - b), \tag{9.54}$$

$$\Delta t_1 = -\frac{d}{c}\eta(\alpha a - \gamma b), \tag{9.55}$$

$$\Delta t_2 = -\frac{d}{2c}(\eta^2[(\alpha\beta + \alpha^2)a - (\gamma\delta + \gamma^2)b] + v^2[\alpha\beta a - \gamma\delta b]), \tag{9.56}$$

$$\Delta t_3 = -\frac{d}{16c}(\eta^3[(8\alpha\beta + \alpha^3)a - (8\gamma\delta + \gamma^3)b] + \eta v^2[8\alpha\beta a - 8\gamma\delta b]), \tag{9.57}$$

and so on, where $u = u_0 + \eta$, $\beta = 1/(n_2^2/n_1^2 - u_0^2)$, $\alpha = u_0\beta$, $\delta = 1/(1 - u_0^2)$, $\gamma = u_0\delta$, $a = n_1/\beta^{1/2}$, and $b = 1/\delta^{1/2}$. The coefficients of $\eta^i v^j$ are the aberration amplitudes.

9.4.4 Stop Position Dependence

One good way of reducing the effect of badly aberrated edge rays is to block them with a strategically placed stop. This may seem wasteful of light, but those rays weren't doing our measurement any good anyway, so they're no loss. This is one example of a case where the stops may be fairly far from the Fourier transform plane.

9.5 REPRESENTING ABERRATIONS

The standard method of representing the aberration coefficients of a wavefront is the *wave aberration polynomial*,[†]

$$W = \sum_{l,m,n=0}^{\infty} W_{2l+n,2m+n,n} h^{2l+n} \rho^{2m+n} \cos^n \phi, \tag{9.58}$$

where W is the optical path difference in meters (converted to $n = 1$ as usual). Think of kW as the (unwrapped) phase of \tilde{A} . Apart from the fact that the practical upper limit of this summation is very finite, it's moderately useful, although more mysterious looking in this form. The coefficients all have names up to order 6 or so (the order of a term

[†]Warren J. Smith, Optical design, Chap. 2 in J.S. Accetta and D.L. Shumaker, *The Infrared and Electro-Optical Systems Handbook*, Vol. 3.

TABLE 9.1. Seidel Aberrations

Piston	W_{000}
Tilt	$W_{111}h\rho \cos \theta$
(Paraxial) defocus	$W_{020}\rho^2$
Spherical	$W_{040}\rho^4$
Coma	$W_{131}\rho^3 \cos \theta$
Astigmatism	$W_{222}h^2\rho^2 \cos^2 \theta$
Field curvature	$W_{220}h^2\rho^2$
Distortion	$W_{311}h^3\rho \cos \theta$

is the sum of the exponents of ρ and h , not $\cos \theta$), which are listed up to order 4 in Table 9.1. Of the ones we haven't talked about, piston is just an overall phase shift, which we often don't care about, and tilt corresponds to a shift in the focal position.

The ray model came first. Ray aberrations are quoted as position errors in the focal plane; because the ray travels along ∇S , the same term shows up in one lower order in the ray model—astigmatism is a fourth-order wave aberration but a third-order ray aberration, which can cause confusion sometimes. We saw in Section 9.2.3 that the local direction of propagation is parallel to $\nabla \Phi$. The ray intercept error is

$$\Delta \mathbf{x} = -\frac{L}{n} \nabla(\text{OPL}). \quad (9.59)$$

The most common way to quote aberration contributions is in peak-to-peak waves over the full diameter of the pupil.

9.5.1 Seidel Aberrations

Wave aberrations up to order 4 are known as Seidel aberrations; their pupil functions are shown in Figure 9.13 and their functional forms in Table 9.1. Looking at the spherical aberration and astigmatism profiles, it is clear that the RMS wavefront error could be significantly reduced by *compensation*, that is, adding a bit of tilt to the coma and a bit of defocus to the spherical aberration so as to minimize $\langle \Delta \phi \rangle$. Compensated waveforms are shown in the bottom row of Figure 9.13.

9.5.2 Aberrations of Beams

Frequently in instruments we want to talk about the aberrations of a fixed laser beam, so it doesn't make much sense to talk about dependence on field angle or image height. In that case, the only relevant terms up to fourth order are paraxial defocus ρ^2 , astigmatism $\rho^2 \cos^2(\theta - \theta_0)$, spherical aberration ρ^4 , and coma $\rho^3 \cos(\theta - \theta_1)$. Since in general no symmetry constraint applies, the θ_i can be anything.

Aside: Zernike Circle Polynomials and Measurements. The Zernike polynomials are an orthogonal basis set for representing the optical phase *in a circular pupil*. This sounds like a great way of expressing measurement results—decomposing a wavefront into orthogonal polynomials is computationally cheap and well conditioned, and all. Unfortunately, their practical utility is zilch. Due to vignetting and beam nonuniformity, our pupils are almost never exactly circular or uniformly illuminated, and errors in the

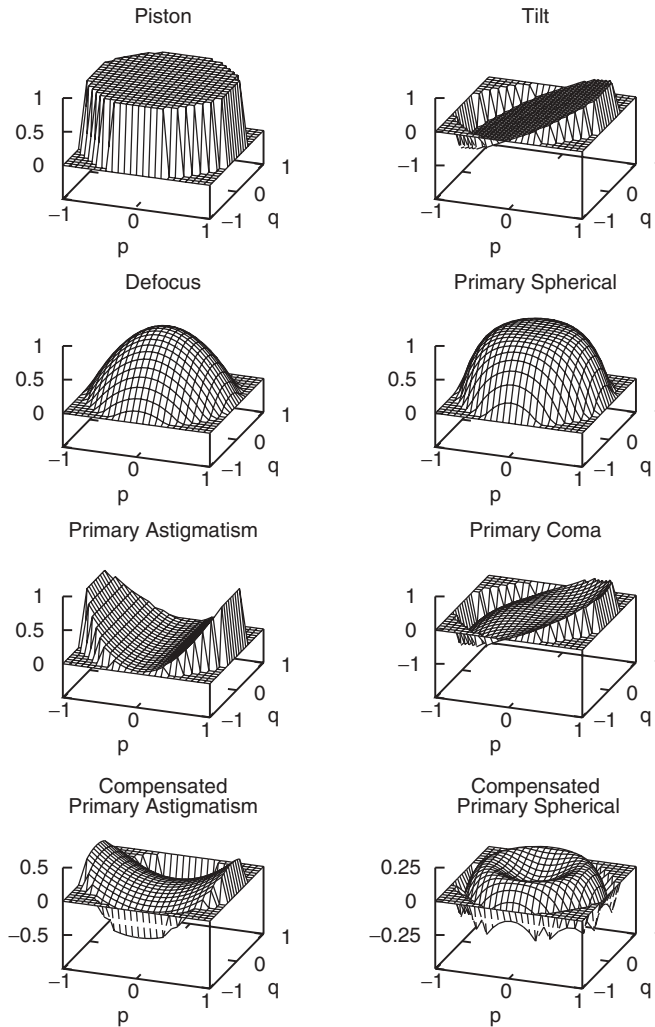


Figure 9.13. Seidel aberrations.

boundary conditions destroy the orthogonality. Defining the “true” Zernike coefficients is especially problematical when our measuring interferometer is intentionally using an elliptical clipping boundary, and perhaps choosing a different ellipse for each run. Even if the pupil stays circular, Zernikes are only obliquely connected to the beam quality measures we care about (e.g., defocus in diopters).

The power series coefficients stay reasonably still even as the bounding ellipse changes, and are pretty well connected to what we care about, so use that up to fourth order in the wavefront. If that isn’t accurate enough, do the calculation numerically.

9.5.3 Chromatic Aberrations

Different wavelengths see different values of n , and a time delay Δt produces different phase shifts. Thus all the coefficients of the aberration polynomial are wavelength

dependent. Changes in focal length with λ are *longitudinal chromatic aberration*, changes in magnification are *lateral chromatic aberration* or *lateral color*, and changes in the primary spherical term are *spherochromatism*. Few of these things are under our control as system designers, which isn't to say they aren't important.

9.5.4 Strehl Ratio

The Strehl ratio is the ratio of the central intensity of a focused spot to what it would be with the same amplitude distribution but zero phase error,

$$R = \left| \frac{\iint \tilde{A}(u, v) du dv}{\iint |\tilde{A}| du dv} \right|^2. \quad (9.60)$$

The Schwarz inequality guarantees that this ratio achieves 1 only when the phase error is indeed zero, and never exceeds that value. A Strehl ratio of 0.8 corresponds to Rayleigh's $\lambda/4$ criterion for a system that is diffraction limited.[†] When building focused-beam instruments, we frequently find that the electrical signal power goes as the square of the Strehl ratio, which is a convenient way of including aberration tolerances in our photon budgets. A useful approximation for the Strehl ratio is Marechal's formula,

$$R \approx \exp \frac{-\langle \Delta\phi^2 \rangle}{2\pi^2}, \quad (9.61)$$

where $\langle \Delta\phi^2 \rangle$ is the mean square phase error in rad^2 (remember to weight the mean square calculation by the intensity and area, and normalize correctly). If you prefer to use the phase p in waves (i.e., cycles), it's $\exp(-p^2/2)$, which is pretty easy to remember.

The Strehl ratio suffers from the same problem as equivalent width; if the peak is shifted from the origin, the ratio may come out very small, even for a good beam. Thus we often want to calculate the Strehl ratio after the tilt contribution (which moves the focus sideways) has been removed. Also, it can give some odd results with multiple transverse mode beams—because the different spatial modes have different frequencies, DC intensity measurements like Strehl ratio miss the rapidly moving fringe pattern in the beam, and so underestimate the far-field divergence. The Strehl ratio is an excellent quality measure for good quality beams, where the intensity profile has one main peak; for uglier ones, consider using Siegman's M^2 instead.

Example 9.6: ISICL Signal-to-Noise Calculation. The Strehl ratio shows up by other names in other fields. In antenna theory, it is called the *phase efficiency*, which is the ratio of the on-axis detected signal to that of a perfectly figured antenna, neglecting loss. Strehl ratio thus shows up as a multiplicative factor on the detected signal from focused-beam instruments. The ISICL sensor of Example 1.12 uses coherent detection with a Gaussian beam. Referring to Figure 1.15, a particle with differential scattering cross section $d\sigma/d\Omega$ crossing near the focus produces a scattered photon flux per steradian

$$\frac{dn_s}{d\Omega} = \frac{2\pi P_T (\text{NA}_T)^2}{\lambda hc} \frac{d\sigma}{d\Omega} R_T, \quad (9.62)$$

[†]Strictly speaking, the 0.8 corresponds to 0.25 wave of uncompensated spherical aberration.

where R_T and NA_T are the transmit beam's Strehl ratio and numerical aperture. When coherently detected by a similar LO beam, the detection NA is $\pi(NA_R)^2 R_R$ (the factor of R_R accounts for the dephased (aberrated) part of the LO power, which goes into the shot noise but not into the coherently detected signal). Doppler shift makes this an AC measurement, so from Section 1.5.2, the 1 Hz SNR is equal to the number of detected signal photoelectrons per second, which is

$$\text{SNR} = n = \frac{2\pi^2 \eta P_T (NA_T)^2 (NA_R)^2}{\lambda h c} \frac{d\sigma}{d\Omega} R_T R_R, \quad (9.63)$$

where η is the detector's quantum efficiency.

The detected SNR goes as the product of the Strehl ratios of the transmit and LO beams. This provides a natural connection between aberrations and signal level (and hence SNR), which is why the Strehl ratio is so useful for instrument designers.

9.6 OPTICAL DESIGN ADVICE

Many books have been written on lens design, and lots of software exists to help. That's not what we're talking about now, and is in fact beyond our scope. Optical design (in the restricted sense used here) is concerned with sticking lenses and other parts together to get the desired result. An analogy from electronics is IC design versus application circuit design, with the lens being like the IC: most of the time you use standard ones, but occasionally you have to do a custom one; it will be a better match, but will cost something to design and will usually be more expensive in production unless you have gigantic volumes.

Computerized exact ray tracing is not usually necessary in instrument design, although if you have easy-to-use software for it, you might as well—it doesn't make anything worse, after all. On the other hand, we often don't have the full optical prescription for the lenses, so exactness doesn't buy us a lot. Thick-lens paraxial ray optics is better than good enough for layout, and our hybrid wave optics model plus some simple aberration theory does the job of calculating image or beam quality, signal levels, and SNR.

If necessary, we can use ray tracing or numerical wavefront simulation to dial in the design once we've arrived at a close approximation, but it is needed surprisingly seldom since the fancy stuff is usually done at very low NA, where life is a lot easier. Apart from etalon fringes, using mostly collimated beams (or parallel light in imaging systems) makes it possible to add and remove components freely, with only minor effects on optical quality.

9.6.1 Keep Your Eye on the Final Output

In Section 1.7.1, we used an expression for the detected photocurrent as the optical system output, and that remains the right thing to do when aberrations and finite aperture systems are considered—you just use the real pupil function instead of the paraxial one, and pay attention to obliquity factors and the Strehl ratio. It's an excellent way to know just when our treatment of system aberrations is losing contact with instrument-building reality. As we saw in Example 9.6, it leads to a natural interest in the Strehl ratio, which appears in the photocurrent and SNR calculations. There are other things that matter

besides the photocurrent (e.g., geometrical distortion), but if you don't see how to relate the figure of merit under discussion to the actual instrument performance, find another way of describing it until you can. Lens designers produce lenses for a living, and we build electro-optical systems.

9.6.2 Combining Aberration Contributions

An optical system is made up of a series of imperfect elements, each contributing its own aberrations. In order to combine them, we note that all the pupils in the system are images of each other, and so the pupil functions multiply together. The exit pupil of each element is imaged at the output of the system by the (ideal) imaging action of all subsequent elements, and the resulting pupil functions multiplied together to get the total pupil function of the system. Watch out for magnification differences—those pupils won't all be the same size, and the smallest one wins.

9.7 PRACTICAL APPLICATIONS

9.7.1 Spatial Filtering – How and Why

Spatial filtering is the deliberate modification or removal of some plane wave components from a beam, and is completely analogous to the ordinary electronic filtering performed in a radio. It is normally done by using a lens to Fourier transform the beam, inserting at the focal plane a mask that is transparent in some places and opaque in others (such as a slit or a pinhole), and then transforming back with a second lens. It is widely used for cleaning up beams and for removing artifacts due to periodic structures in a sample (e.g., IC lines).

Spatial filtering using a pinhole can make a uniform beam from a Gaussian one but will waste around 75% of the light doing it. The smallness of the pinhole required for a good result with a given NA may be surprising—the first Airy null is *way* too big (see Example 9.9).

Spatial filters are not as widely used as one might expect, based on the analogy to electrical filters. They require complex mechanical parts, which is a partial explanation. The real trouble is tweakiness: they are difficult to align, easy to misalign, and sensitive to mechanical and thermal drifts; if ordinary translation stages are used, an expensive and very labor-intensive device results. These problems can be reduced by clever design, for example, by using a laser or an in situ photographic process (with a fixed optical system) to build the mask right inside the filter. If you need a pinhole spatial filter, use a big pinhole ($>20\ \mu\text{m}$ diameter in the visible) and correspondingly low NA to get the best stability. One other problem is that they nearly all have sharp edges, which isn't usually optimal.

9.7.2 How to Clean Up Beams

Laser beams are frequently rotten. Most gas lasers produce beams of reasonable quality, but these are often surrounded by multiple reflections, scattered light, and spontaneous emission. Diode lasers are worse; their beams suffer from astigmatism and are highly asymmetric. We may need to free these beams of their bad associations in order to make them useful, or to change their spatial distributions to something more convenient.

This is done through spatial filtering and the use of apertures (and occasionally special elements such as anamorphic prism pairs). Since both of these operations are performed by passing the beam through holes of various sizes, the distinction is somewhat artificial but is nonetheless useful: apertures are used on the beam before focusing, and spatial filters in the Fourier transform plane. A seat-of-the-pants test is that if it requires a fine screw to adjust, it's a spatial filter.

Putting an aperture some way out in the wings of the beam (say, four times the $1/e^2$ diameter) has little effect on its propagation characteristics, so use them freely to reduce artifacts. If the artifacts are close to the beam axis, it may be helpful to let the beam propagate for some distance before applying the aperture; a lens may be helpful in reducing the optical path length this might otherwise require (here is where it shades into spatial filtering). Appropriately placed apertures can turn a uniform beam into a good Gaussian beam, or chop off the heavily aberrated wings of an uncorrected diode laser beam.

Example 9.7: How Small Do I Have to Make My Aperture? Slits and pinholes are frequently used to render an instrument insensitive to the direction from which incident light originates. A monochromator (see Example 7.1) uses a slit to ensure that the light incident on its grating arrives from one direction only; this maximizes its spectral selectivity. There's obviously a trade-off between selectivity and optical throughput; because different positions across the slit translate to different incidence angles on the grating, a wider slit translates directly into lower spectral resolution.

More subtly, with a wide slit a change in the position or direction of the incoming light can cause apparent spectral shifts. This is because the slit doesn't obliterate the spatial pattern of the incoming light, it just clips it at the slit edges, and so broadens its angular spectrum. If the slit is made small enough, it can't shift far laterally, and we expect that the width of the diffraction pattern will swamp any likely angular change from the incoming light. On the other hand, using a slit that narrow can make things pretty dim. Exactly how narrow does it have to be?

As we saw in Section 5.7.9, scatterers tend to produce patterns that are aligned with the incident light, but smeared out in angle. The same is true of slits and pinholes; in the Fourier optics approximation, the (amplitude) angular spectrum of the incident light is convolved with the Fourier transform of the aperture's transmission coefficient.

If we illuminate the slit at an incidence angle θ , the main lobe of the sinc function is aligned with the \mathbf{k} vector of the incoming light, and the intensity along the normal to the slit will decrease as θ increases. Since x/λ and u are conjugate variables, if we require that a change of $\pm\delta$ radians cause a fractional decrease of less than ϵ in the central intensity of the diffraction pattern, the slit width w must obey

$$w \sin \delta < \sqrt{\frac{6\epsilon}{\pi}}, \quad (9.64)$$

so that for a relatively modest requirement, for example, a shift of less than 5% from a $\pm 10^\circ$ rotation, $w < 2.25\lambda$. It is difficult to use slits this small, but not impossible. Improving on this, such as requiring a 1% shift from a $\pm 30^\circ$ rotation, is impractical, as it requires a slit very small compared with a wavelength.

Similar considerations govern the sensitivity to lateral motion of a pinhole in a focused beam. The moral of this story is that although spatial filters can reduce the sensitivity of a measurement to angular shifts of illumination, a complete cure is not to be

TABLE 9.2. Effects on a Gaussian Beam of a Circular Aperture of Radius r Placed at the Beam Waist

r/w	$\Delta I(\text{nom})$	w_{best}	$\Delta I(\text{best})$
0.5	0.66	2.72	0.027
1.0	0.30	1.47	0.022
1.5	0.099	1.13	0.014
2.0	0.022	1.0245	0.0055
2.5	0.0028	1.003	0.0014
3.0	0.0002	1.0002	0.0001

found here. Single-mode fibers are dramatically better, because they control the direction from which the light hits the pinhole (see Section 8.2.2). Unfortunately, they'll often make your light source dramatically noisier as well, by turning FM noise into AM; see Section 8.5.13.

Example 9.8: How Small Can I Make My Aperture? On the other hand, if we have a beam that has artifacts out in its wings, such as most gas laser beams, we would like to make the aperture as small as possible without causing objectionable diffraction rings. If the beam is Gaussian, the ripple amplitude is very sensitive to slight vignetting. Table 9.2 gives the maximum absolute deviation ΔI from both the nominal and best-fit Gaussian beam, in the transform plane, due to apertures of different radius placed at the beam waist. Values have been normalized to a central intensity of 1.0.

Example 9.9: Using an Aperture to Make a Nearly Gaussian Beam. Let's consider the use of an aperture in the pupil to make a uniform beam into a reasonable approximation to a Gaussian beam. Although a uniform beam exhibits an extensive set of diffraction rings, an appropriately placed aperture can greatly reduce their amplitude. A uniform beam of amplitude E_0 and radius a at a wavelength λ gives rise to a far-field Airy pattern whose first null is at $\theta = 0.61\lambda/a$. If the beam is Fourier transformed by a lens of focal length f , so that the numerical aperture $\text{NA} = a/f$, the pattern is given by

$$E(r) = E_0 k \text{NA}^2 \text{jinc}(kr\text{NA}). \quad (9.65)$$

Figure 9.14 shows the result of truncating this pattern with a circular aperture at the first Airy null, and recollimating with a second lens. Here the radius of the original uniform beam was $100 \mu\text{m}$ and the wavelength was 530 nm . The Gaussian shown has a $1/e^2$ (intensity) radius of $95 \mu\text{m}$, and the sidelobes are below 1 part in 10^3 of the peak intensity. Only about 15% of the light is lost. Note that the peak intensity is twice that of the original uniform beam. Contrary to appearances, the total beam power has gone down—the high peak intensity covers a very small area since it's at the center. A graded neutral density filter, which is the competing technique, cannot increase the central value, so that it is limited to at most half this efficiency, and far less if we expect the Gaussian to drop to nearly zero before the edge of the beam is reached; on the other hand, it requires no lenses. (See Figure 9.15.)




Figure 9.14. Turning a uniform beam into a nearly Gaussian one with a pinhole of the same radius as the first Airy null.




Figure 9.15. The data of Figure 9.14 on an expanded scale.

9.7.3 Dust Doughnuts

In general, a bit of dust on a lens is no big deal. The exception is when the dusty surface is near a focus, in which case dust is extremely objectionable, leading to strong shadows called *dust doughnuts*. (The shadows are annular in shape in a Cassegrain telescope, hence the name.) How far out of focus does the dust have to be?

Assuming the dust is less than 100 μm in diameter, and that a 1% intensity error is acceptable, a focused spot has to be at least 1 mm in diameter at the dirty surface, so

the required defocus is

$$|\delta Z_{\text{defocus}}| \gtrsim \frac{1 \text{ mm}}{2 \cdot \text{NA}}. \quad (9.66)$$

9.8 ILLUMINATORS

This discussion is indebted to an SPIE short course, “Illumination for Optical Inspection,” by Douglas S. Goodman.

9.8.1 Flying-Spot Systems

A scanning laser microscope is an example of a *flying-spot* system, as opposed to a *full-field* or *staring* system, in which a larger area is imaged at once. The first flying-spot optical systems in the 1950s used an illuminated point on a cathode ray tube as the light source, and a PMT as the detector because the spot was so dim. That at least had the advantage of speed and no moving parts. Flying-spot systems are simple to analyze, because their PSFs are the product of the illumination and detection spots (whether amplitude or intensity is the appropriate description), and there is no problem with speckle or scattered light smearing out the image.

9.8.2 Direction Cosine Space

Full-field systems require a bit more work to specify accurately. We’re usually using thermal light from a bulb of some sort, so the illumination is coming from some reasonably wide range in (u, v) . Figure 9.16 shows a sample region illuminated by a cone of light, which is plotted on a large sphere (much bigger than the sample region of interest) and then projected down into the (x, y) plane. Since the exact radius R is of no significance, we normalize it out and plot the illumination in terms of u and v . Due to the curvature of the spherical surface, illuminated patches lying near the horizon are strongly foreshortened in area (by a factor of $\cos \theta$, where θ is the polar angle). This is quite useful actually, since the flux passing into the surface is reduced by the same factor due to the spreading out of the obliquely illuminated patch (equivalently, the apparent source area seen by any one point is reduced by the cosine). Figure 9.17 shows bright- and dark-field systems, with oblique and concentric illumination.

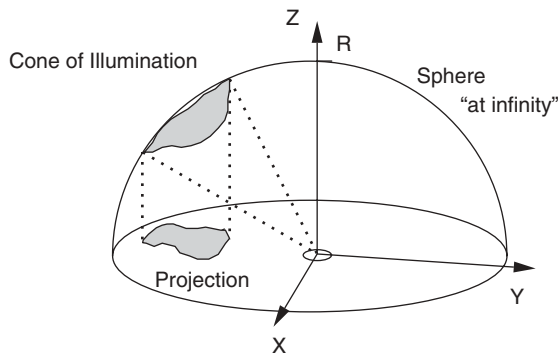


Figure 9.16. Direction cosine space.




Figure 9.17. The general type of illumination is determined by the overlap of the illumination and detection patterns in u,v space.

9.8.3 Bright and Dark Fields

Illuminators are divided into bright- and dark-field types, depending on whether the image of a featureless sample (e.g., a mirror in reflection or a glass plate in transmission) is bright or dark; this is equivalent to whether the illuminator is directly visible in the image. In direction cosine space, a bright-field system has illumination and collection patterns that substantially overlap, whereas they miss each other completely in a dark-field setup. It is possible to work in intermediate states, generically called dim field.

Dark-field images consist entirely of scattered light, so that highly scattering areas such as dust particles appear bright on a dark background. Bright-field ones see darker areas where light has been absorbed or scattered so far as to miss the collection lens. This means that increasing the NA tends to wash out the contrast in bright field—all that scattered light is collected and reassembled in the image, so only the absorption contrast is left. Dark-field illumination shows up phase objects and weak scatterers better than bright field, but the image intensity is quadratic in the phase shift ϕ , so it is not especially sensitive at small ϕ . In Sections 1.5 and 10.3.5 we discuss ways to get higher sensitivity for small signals.

9.8.4 Flashlight Illumination

The first crack at an illumination system is often to take a fiber bundle illuminator and point it at the sample, keeping it away from the rest of the optical system. This produces a kind of oblique dark-field illumination that varies in angular spectrum and intensity

across the sample. If the sample is diffuse (e.g., white paper), or quantitative results are unnecessary, this may work fine—and it's certainly quick.

9.8.5 Critical Illumination

Of more thoughtful illumination strategies, the simplest is critical illumination: just image the source down onto the sample. Any variations in the source appear directly in the image, so the source must be very uniform. A more subtle problem is that critical illumination is generally nonstationary; that is, the illumination diagram is different at different points in the field of view. This is because the image of the bulb radiates in all directions, and hence its angular spectrum will be vignetted in the pupil of the imaging lens, at least for points sufficiently far off-axis. This vignetting can be so severe that the edges of the field change from bright to dark field, as the specular reflection misses the collection lens completely.

9.8.6 Köhler Illumination

Köhler illumination overcomes the source nonuniformity problem by putting the sample at the Fourier transform plane of the condenser, or equivalently by imaging the source on the entrance pupil of the collecting lens. This strategy makes the illumination conditions stationary across the field (i.e., it doesn't change from bright to dark field the way critical illumination can), because all the unscattered light makes it through the pupil of the imaging lens without being vignetted.[†]

Köhler illumination tends to keep the spatial frequency bandwidth more nearly constant; the center of the light cone from each point goes through the center of the pupil. Vignetting will reduce the spatial frequency bandwidth at the edges but won't make the illumination change from bright to dark field across the sample, as can happen with critical illumination.

9.8.7 Testing Illuminators

Douglas Goodman suggests using the *ball-bearing test*. A sphere reflects a collimated beam into 4π steradians with equal power per steradian in all directions, and this of course works backwards too; you can see the whole illumination pattern in direction cosine space by looking at its image in a sphere. For looking at microscope illumination, mount the ball bearing on a flat black plate in reflection, or on a microscope slide with a tiny glue patch in transmission. You can see the microscope lens in the ball too, so you can align the two for concentricity that way. The direct readout in direction cosine space allows you to see what your illumination function looks like very easily, though you may have to stop down your collection NA to get a clear view of it.

In other situations, the interior of a ping-pong ball cut in half makes a very handy screen for projecting the illumination on. It's especially good as a scatterometer with diffracting structures and laser illumination; shine the laser through a hole in the top of the hemisphere and look at the diffraction spots. (Ping-pong balls can also be used as integrating spheres—see Section 5.7.8.)

[†]Recall that the pupil is usually at or near the transform plane.

9.8.8 Image Radiance Uniformity

We saw in Section 2.4.1 that a planar source with no preferred direction is Lambertian. Whenever such a diffuse planar object is imaged on a planar detector over a wide field, the edges of the image become darker. The darkening gets worse rapidly as the field angle θ increases; the image radiance goes as $\cos^4 \theta$, as we can verify by counting powers of the cosine. The object–pupil distance goes as $\sec \theta$, which gives us two cosine factors since the flux at the pupil goes as $1/r^2$. The projected pupil area goes as the cosine, which makes three, and Lambertian angular dependence of the object radiance makes four factors of the cosine altogether. Vignetting at places other than the pupil and the increased Fresnel losses at high angles make real systems somewhat worse in general, though there are ways of increasing the loss near the middle to flatten the curve a bit. The \cos^4 falloff is also why it's easier to get uniform illumination by spacing many sources at close intervals, rather than trying to use one source plus diffusers.

9.8.9 Contrast and Illumination

Imaging systems stand or fall by their contrast. A featureless image contains no information; one where the contrast comes from an unexpected mechanism (e.g., stray fringes) may contain even less—you're more confused and misled after the measurement than before.

The word *contrast* is used in two different senses. Contrast is defined as the ratio of p-p intensity change to mean intensity (flux density, irradiance) in the image, which is 0 for no contrast and 1 at perfect contrast:

$$C = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}. \quad (9.67)$$

The other sense denotes the source of the contrast: imaging can be done in phase contrast, interference contrast, and so on. Contrast is a normalized measure; dialing down the illumination power doesn't change the contrast, but of course the signal level and SNR will deteriorate.

The image contrast is a complicated function of the illumination type, collection, detection strategy, and sample. The dependence is not weak, either; different illumination will make bright areas look darker than dark areas. For example, smooth aluminum lines on rough, (nearly black) polycrystalline silicon carbide will look very bright when illuminated in bright field, but will look dark when illuminated obliquely, because their specular reflection causes the returned light to miss the objective lens, whereas the rough surrounding surface will scatter some light into the lens—a black surface can look lighter than a mirror. Other examples are everywhere. The scattering geometry is important too; linear features such as scratches or smears scatter very efficiently perpendicular to the line, but only weakly at other angles; if the illumination doesn't have any component perpendicular to the lines, the scattered light will be weak. (This phenomenon is familiar to us all—think of a smeared car windshield at night.)

The most important thing in choosing an illumination strategy for your measurement is to mess around with a whole lot of alternatives before making a final choice. Automatic inspection, machine vision, lithography, microscopy, and trace defect detection live and die by the quality of the illuminator and by how well it's matched to the characteristics of the problem. Poorly chosen illumination can make the software job many times harder

(or even impossible). Don't just buy a fiber illuminator, shove it in, and expect it to work well.

In the highest resolution optical microscopy, resolution is improved by using diffuse illumination and high-NA microscope objectives, because high spatial frequency components can take a plane wave component near grazing and send it back the way it came, so that the spatial frequency bandwidth is doubled. Unfortunately, most samples show almost no contrast when examined this way, so experimentation is needed even here.

9.8.10 Retroreflectors and Illumination

Any time you're trying to inspect a transparent or specularly reflecting object that's bigger than your lens, there's a problem with illumination uniformity and efficiency, as well as a serious case of changing illumination conditions with field position. One very useful way of solving this is to use a big chunk of retroreflecting material, with the illuminator and sensor optically superposed. Light from the source bounces off the specular surface, hits the retroreflecting material, and retraces its path to the collecting lens. This is a great way of keeping the illumination conditions the same even with a small light source—it's sort of a poor man's telecentric system. The angular spread of the returned beam is a degree or so, so there is a certain amount of ghosting and other artifacts, but for jobs like inspecting plastic film for defects, or looking at semiconductor wafers, it's a very useful trick sometimes (see Section 7.8).

Optical Measurements

My means are sane, my motive and my object mad.

—Captain Ahab, in *Moby Dick* by Herman Melville

10.1 INTRODUCTION

It's easy to be beguiled by reading people's papers into believing that optical instruments are complicated to fit them exquisitely for their purpose, but they aren't.[†] They are complicated in order to work around the limitations of available components, weird sample geometries, the idiosyncrasies and prejudices of experimenters,[‡] and the occasional inconvenient Law of Nature. If we could do stable, flexible, and convenient measurements at the quantum limit all the time with no problems, there would be many fewer types of instrument. Besides component limitations, the main problems come from noisy background signals and environmental effects. That isn't all bad—it provides instrument builders with job security and an unending string of fun problems to solve. There are some enduring principles underlying all this welter, and that's what we'll concentrate on here—background reduction, labeling signal photons to distinguish them from background, and exploiting all the changes that physical interactions can produce in light beams.

10.2 GRASS ON THE EMPIRE STATE BUILDING

As Miller and Friedman engagingly point out in *Photonics Rules of Thumb*, our job is like measuring the height of a tuft of grass on the roof of the Empire State Building, as shown in Figure 10.1; the desired signal is a small delicate thing, riding on a huge ugly background. (It's lucky for New York that the Empire State Building is actually quite nice looking, and doesn't shuck and jive the way our background signals often do.) This

[†]There's a wonderful paper on the tension between frankness and self-promotion, written by Edmond H. Weiss, "Professional Communication" and the "Odor of Mendacity": The Persistent Suspicion that Skillful Writing is Successful Lying'. *IEEE Trans. Professional Commun.* **38**, 3 (September 1995).

[‡]The author's have of course been rehearsed at some length in this book, but other people have them too.




Figure 10.1. Background signals.

whimsical image is a helpful one all through our discussion, so hang on to it. The idea is to make the grass as conspicuous as possible, by fertilizing it to make it grow taller, painting it fluorescent orange to make it stand out, moving it to ground level, climbing the building, or looking at it from a helicopter. In other words, by getting as much signal as possible, labeling signal photons somehow so they can be distinguished from background photons, and getting rid of the background itself, or at least its fluctuations.

10.2.1 Background, Noise, and Spurious Signals

Some background signals are inherent in the measurement, for example, speckle statistics in shearing speckle interferometry or shot noise in fluorescent measurements. Most of the rest come from the immediate environment (e.g., room lights) or the measurement system itself: source noise, stray reflections, and molecular (Rayleigh) scatter.

This junk comes in two flavors: additive and multiplicative. The distinction is based on what we consider to be our signal. In the grass image, the thermal expansion of the building is additive, since it just lifts all the grass the same amount as the building heats and cools, whereas wind is more or less multiplicative, because tall grass waves more than short grass. In spectroscopy, (additive) background shifts the baseline and (multiplicative) gain drift changes the peak heights. Source power variations produce both; the baseline is proportional to the source power, so its fluctuations appear as additive noise, and the peak heights are also proportional to source power, which is multiplicative noise. If we thought of the baseline as signal too, this would be a purely multiplicative noise source. Low frequency multiplicative junk puts noise sidebands on all your measurement signals and carriers, at all frequencies, which are going to stay there

and frustrate your attempts to make measurements at high SNR unless you get rid of them somehow.

10.2.2 Pedestal

A quiet additive background or *pedestal* is like the DC offset of a single-supply circuit. As long as it is really still, it doesn't cause any problem. Background fluctuations are far more likely to be the limiting factor. Unlike a quiet electronic bias level, optical background signals have full shot noise at least, and usually much more of other stuff too.

Occasionally the pedestal is so big that its fixed pattern noise dominates your dynamic range budget. An example is a HgCdTe area array in the thermal IR, where the integration time and signal linearity are limited by the nonuniform sensor seeing a huge 300 K background. The fixed pattern itself is a strong function of intensity, so careful calibration at a variety of scene brightnesses is required, but eventually the fixed pattern noise can be tamed. In most cases, the background shot noise can be made to dominate the fixed-pattern fluctuations and readout noise.

10.2.3 Background Fluctuations

Background fluctuations fall into two classes: photocurrent shot noise and everything else. Shot noise is spectrally white, so it lands on your signal irrespective of where the signal might be found, and has Gaussian amplitude statistics, so its average properties are very easily calculated; the only way to get rid of it is to reduce the background light, using baffles, filters, and so forth.

The other category includes source excess noise, circuit noise, and spurious modulation of the background signal (e.g., 100/120 Hz from room lights). We can usually get round this sort of thing by optical hacks and circuit improvements.

10.2.4 Noise Statistics

As we'll see in Section 13.6.14, shot noise really is Gaussian, to amazing accuracy, and so is Johnson noise, although it's harder to measure. Other types of noise are not so nice, especially $1/f$ noise and mode hopping, which tend to have a strongly popcorn character, that is, lots of fairly sharp step-like jumps, and wildly nonstationary statistics, with long bursts of noise coming at irregular intervals and going away again for awhile. Most measurements are more sensitive to extreme values than to the fuzz level.

10.2.5 Laser Noise

We talked about laser noise in Section 2.13; the two basic kinds are residual intensity noise (RIN) and phase noise. It appears as a purely multiplicative problem in dark-field measurements, and both additive and multiplicative in bright-field ones. It's often quite large, 10^{-3} – 10^{-5} , which will really degrade a measurement if we don't take steps.

Diode lasers have lower RIN power spectral density (at least if you can prevent them from mode hopping), but it extends from DC to a gigahertz and so is hard to avoid by rapid modulation; by contrast, a frequency shift of a few tens of megahertz will get you out of the excess noise region of gas and solid state lasers. If the laser has multiple

longitudinal modes, their beats will produce near-replicas of the baseband noise spaced near multiples of $c/(2n\ell)$, where ℓ is the cavity length and n is the real part of the refractive index at the laser wavelength.

Diode lasers have significantly worse phase noise than other types do, due to the low finesse and intrinsic instability of their resonators. DFB lasers are better than Fabry–Perot types, but still nowhere near the frequency stability of a HeNe. Demodulation of the phase noise and mode partition noise by etalon fringes and multiple scattering is a horrendous problem in bright-field diode laser measurements, as we saw in Section 2.5.3; Section 19.1.1 has a real tear-jerker that hinges on this point.

10.2.6 Lamp Noise

Lamps being thermal devices, one might think their noise would all be down in the very low baseband, but unfortunately that isn't always true; the thermal signal falls off as $1/f^2$ or slower, and plasma noise in arc lamps extends well into the kilohertz. Flashlamps exhibit pulse-to-pulse variations of 1% or so, which puts big noise sidebands on the repetition rate component (see Section 15.9.4).

10.2.7 Media Noise

There are many classes of measurements, for example, photon echo and other spectral hole-burning measurements, optical recording, photographic densitometry, and holographic reconstruction, in which the sample consists of some recording medium in which a pattern has been recorded. The noise budget has to include the noise added in the recording process.

Example 10.1: Media Noise in Spectral Hole Burning. The frequencies and line-shapes of the spectral lines of chromophores (i.e., atoms, molecules, and ions exhibiting spectral lines) depend on their environment. Line broadening is usefully divided into *homogeneous* and *inhomogeneous* broadening. Homogeneous broadening affects all the chromophores the same way, at least on average: collisions in a gas, orbital overlaps in a crystal (which lead to reduced upper state lifetime). Inhomogeneous broadening affects different chromophores differently: Doppler shifts in gases and the different Stark shifts seen by chromophores in different places in a crystal. In a nutshell, in a measurement of a single atom, homogeneous broadening shows up in the width of the transition, and inhomogeneous broadening shows up as a shift in the center frequency—it's the ensemble of many chromophores with different offsets that makes up the inhomogeneous lineshape.

Rare-earth ions are unique because their F-shell electrons are so well shielded by the valence electrons that they are hard to Stark-shift. This means that the homogeneous broadening of such ions when introduced into a suitable crystal can be very small indeed—below 100 Hz for some nominally forbidden transitions at low temperatures (4 K). In an inhomogeneously broadening medium, such as a YAG crystal, this narrow homogeneous linewidth means that there can be 10^3 – 10^6 resolvable frequencies. Since the homogeneous transition is narrow, the upper state lifetime must be long, leading to the possibility of *spectral hole burning*, in which a narrowband laser bleaches the absorption in a narrow bandwidth by pumping all the chromophores in that interval into the upper state. There are a lot of things one can do with this sort

of idea, but one problem is that with that many resolvable frequencies, there aren't that many chromophores per bin—perhaps only 10^6 . Unless the hole is fully saturated (all the chromophores are bleached), this medium is going to have frozen-in counting noise due to the stochastic nature of absorption. Such a medium cannot record any signal with more than 60 dB SNR, no matter what the readout system's photon budget may be—the desired signal isn't identical with the sample property being measured, so the media noise limits the attainable SNR, even though the readout can be made very quiet.

10.2.8 Electrical Interference

This is more of a problem for Chapters 16 and 19, but deserves mention here because sometimes it's easier to avoid localized background effects than to fight them; for example, trying to do sensitive measurements at harmonics of 50 or 60 Hz is a mug's game unless you're running battery powered in the Arctic—and even then, the Earth's magnetospheric activity shows peaks at those frequencies due to power line excitation at lower latitudes that puts 50 and 60 Hz sidebands on the auroras, so you still can't win. Similarly, if your measurement uses a Pockels cell modulated with 1.5 kV RMS of 10 MHz sine wave, it's going to be pretty difficult to make shot noise limited measurements right there. It's certainly possible, but in doing it you'll learn more about grounding and shielding than you probably care to. A slower subsidiary modulation (e.g., chopping the light) will give you signals at, say, $10 \text{ MHz} \pm 2 \text{ kHz}$, which is a lot easier to work with.

10.2.9 Electronic Noise

Electronic noise is dissected minutely in Chapter 18, because it is such a common source of measurement-destroying additive junk. The main effects are Johnson noise in a too-small load resistor, the technical noise of amplifiers, and $1/f$ noise due to not cleaning the circuit board, cracking an IC package by overheating, or just choosing the wrong components (e.g., thick film resistors). No matter how good your detector, a poor choice of front end will make it look bad. A few kinds of electronic noise are multiplicative, such as base resistance voltage noise in an unbiased cascode stage (Section 18.4.4) or denominator noise in an analog divider, and others go as the square root of the photocurrent, for example, the shot noise of a current mirror on a photodiode.

10.2.10 Quantization Noise

A final, and usually more or less additive, noise source is quantization. In Section 13.11.1 we'll see that it has an RMS value of $1/\sqrt{12}$ ADU, and that it will severely limit our dynamic range unless we reduce that range in analog before digitizing.

10.2.11 Baseband Isn't a Great Neighborhood

All the ruffraff of the frequency domain collects at baseband. The $1/f$ noise, power supply ripple, arc wander, table vibrations, wind noise, and so on, all happen down in the hertz to kilohertz region. There are two ways to deal with this; either clean up the neighborhood or move away yourself. Both are possible, but cleaning it up prevents it from following you, which it otherwise tends to do, since a lot of it is multiplicative.

Aside: Moving the Grass to Ground Level. The easiest method for measuring the grass is to send someone up to the roof with a trowel to bring it down to ground level, then measure it with a child's ruler. This is where smart measurement ideas are really important: they can make hard things easy by redefining the problem.

10.3 DETECTION ISSUES: WHEN EXACTLY IS BACKGROUND BAD?

Background signals are actually not an unmixed evil, and sometimes are an enormous help in making sensitive measurements. The key is their optical phase correlations with the desired signal. First, let's take a close look at noise in bright- and dark-field measurements.

10.3.1 Dark Field

A dark-field measurement is one where ideally the only light is signal. A dark-field microscope (see Section 9.8.2) is the classical example, but there are lots of others: fluorescence, scintillation, second harmonic generation (SHG). If there is enough light to override the circuit noise, dark-field measurements are terrific.

Unfortunately, that's a pretty stiff requirement, because the detected electrical power is quadratic in the optical power, so the sensitivity of the measurement drops off extremely badly at low light levels; dark-field measurements are almost always severely Johnson noise limited unless you're using electron multiplying detectors (PMTs, APDs, L^3 CCDs, or MCPs). Section 18.2 shows how careful you have to be in order to get to the shot noise limit with a $2\ \mu\text{A}$ photocurrent in a 1 MHz bandwidth; that $2\ \mu\text{A}$ corresponds to a 12 THz average photoelectron count rate. A photon-counting photomultiplier can see single photoelectrons (a good one can at least, providing it's cooled) and will get you up to around 10^6 – 10^7 photoelectrons/s, which is of the order of 1 pW, but that doesn't fill the gaping hole between there and 12 THz. In that empty space, we're stuck with analog mode PMTs and APDs, which are no fun at all when you try to make stable measurements—they're great for on/off, for example, optical fiber communications. Changing to bright field is often a good solution, although that changes the counting statistics, which sometimes matters (see Section 10.3.6).

10.3.2 Bright Field: Amplitude Versus Intensity Sensitivity

We've talked about the advantages of coherent detection in Sections 1.5 and 3.11.7, where we saw that this technique allows shot noise limited measurements at all received power levels. Bright-field measurements have the same benefits, because the background (if it's coherent with the signal light) functions as the local oscillator. One way of looking at this is as dynamic range compression. The interference term's electrical power is linear in the signal beam power, so a 1 pW optical signal winds up being 60 dB (electrical) below a $1\ \mu\text{W}$ signal, rather than 120 dB as in a dark-field measurement. This lifts small optical signals up above the electrical noise most of the time, and this is the key to the advantages of bright-field and dim-field measurements. Since one coherently added photon/s/Hz produces a current change equal to the RMS shot noise, and the shot noise can be made to dominate the circuit noise, the nasty hole goes completely away.

Remember that coherent detectors are highly selective in frequency, position, and angle, which helps reject incoherent background light (the good news) but restricts the optical signal the same way (the bad news).

10.3.3 Coherent Background

So far we've been thinking of additive backgrounds adding in intensity, or equivalently in photocurrent. When the signal and background beams are coherent, though, this is not the only possibility. A pure extinction measurement, where the signal is a small blip on the DC photocurrent from the unattenuated laser, has a shot noise floor of

$$\langle \Delta i^2 \rangle = 2eI_{DC}. \quad (10.1)$$

There are two ways of looking at this, namely, the *optical theorem* (also known as the *extinction sum rule*) and coherent detection of forward-scattered light. Let's describe the total field as $\psi_{\text{inc}} + \psi_{\text{scat}}$, where ψ_{inc} is the total field in the absence of the absorber.

10.3.4 Optical Theorem

By conservation of energy, the power lost to the beam has to equal the absorption plus the total integrated scatter. We can write the total power as

$$\iint |\psi_{\text{tot}}|^2 d\Omega = \iint |\psi_{\text{inc}} + \psi_{\text{scat}}|^2 d\Omega = \iint (|\psi_{\text{inc}}|^2 + |\psi_{\text{scat}}|^2 + 2\text{Re}\{\psi_{\text{scat}}\psi_{\text{inc}}^*\}) d\Omega. \quad (10.2)$$

For a lossless scatterer, $P_{\text{tot}} = P_{\text{inc}}$, so the last two terms sum to zero:

$$\iint \psi_{\text{scat}}\psi_{\text{scat}}^* d\Omega = -2\text{Re} \left\{ \iint \psi_{\text{scat}}\psi_{\text{inc}}^* d\Omega \right\}. \quad (10.3)$$

In other words, a small detector sensing extinction in the transmit beam receives the same amount of signal as a huge detector covering the whole $4\pi - \epsilon$ sr (excluding the main beam). The bad news is that the shot noise of the beam is also there, which reduces our sensitivity—but it's still a dramatic improvement over a dark-field detector of the same NA.

The other way of looking at it is by the coherent detection Rule of One from Section 1.5.2: *One* coherently added signal photon in *One* second gives an AC measurement with *One* sigma confidence in a *One* hertz bandwidth. Although the beam itself generally has a tiny solid angle Ω_{Tx} , so that the scattered signal power in that solid angle is impossibly small, the coherently added ψ_{inc} lifts it up to the same level as the huge dark-field detector would see. This operation is unable to improve the counting statistics of ψ_{scat} in the small Ω_{Tx} , but due to the Rule of One, it does as well as it can: the measurement sensitivity is limited by the ratio of the power in ψ_{scat} to its *own* shot noise in Ω_{Tx} (*not* the power in ψ_{scat} divided by the shot noise of ψ_{inc}). (For many years the author had to rederive this fact every time before using it, because it seemed so counterintuitive.)

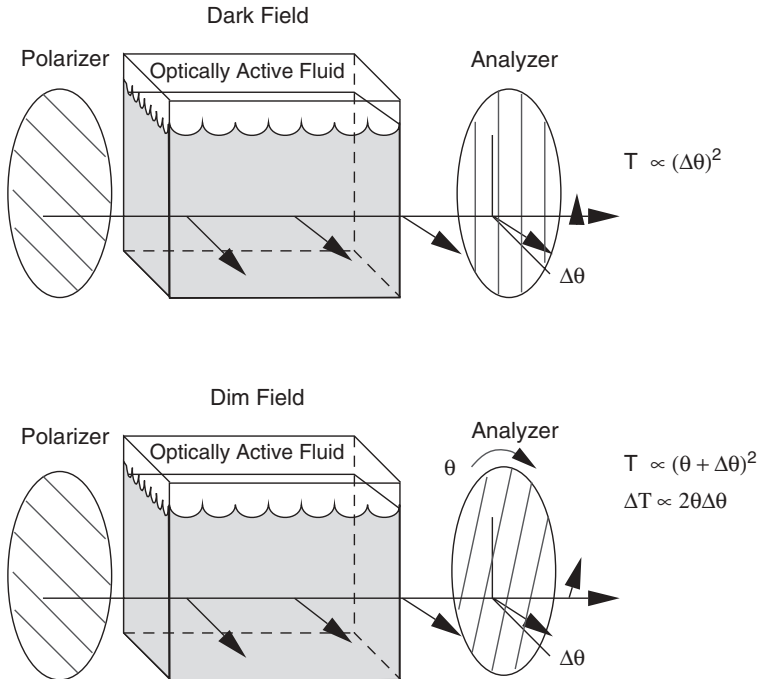


Figure 10.2. Dim-field measurements use a bit of coherent background as an LO signal. Here a crossed-polarizer measurement gets an SNR boost from uncrossing them a bit.

10.3.5 Dim-Field Measurements

Measurements that could be dark field, but where just a bit of coherent background is added are called *dim-field* measurements. An example is the slightly uncrossed crossed-polarizer measurement of Figure 10.2. The coherent background functions as a local oscillator, and so lifts the signal shot noise above the Johnson noise, but not much above; you stop when the multiplicative noise due to LO beam fluctuations becomes the limiting factor, or when you start rejecting your own signal by turning the analyzer too far.

10.3.6 Bright and Dark Fields Are Equivalent

The net of all this is that a bit of coherently added background doesn't hurt anything and may help a lot. As a matter of fact, there is no limit in principle to how much coherent background light you can add. At some point things start deteriorating, due to detector saturation and residual laser noise. We know how to get rid of those pretty well, except when the laser RIN and frequency noise get intermingled, for example, by vignetting or etalon fringes. Another source of trouble is rising electronic noise due, for example, to voltage noise in a cascode stage (Section 18.4.4), but that doesn't happen until we get up into at least hundreds of microamps of photocurrent, and it isn't too hard to get to the shot noise with that. Thus the RMS fractional uncertainty in the photocurrent is the same in bright and dark fields.

They do differ in one moderately important respect: the counting statistics of a bright-field measurement are Gaussian, and those of a photon-counting measurement are Poissonian. This leads to the Bright-Field Rule:

Bright and dark fields are exactly equivalent except for their counting statistics.

This equivalence means that we're free to choose bright or dark field based on other considerations. If you can get a big detection NA, have no background or scattered light problems, and can stand the large dynamic range, you're better off with dark field, whereas in the presence of strong optical or electronic background signals for which the selectivity and gain of bright field helps, or a large signal range, go with bright field or dim field. A noise canceler can get rid of the laser noise, or alternatively (for gas and solid state lasers) a heterodyne interferometer or FM system can avoid the worst of it.

Aside: Counting Statistics: Bright versus Dark Field. In Section 13.6.14 we'll discuss the differences, and how to set detection thresholds intelligently based on the confidence level required. The difference is less striking than it may look, since what's going on is really just the reapportionment of uncertainty; the zero signal state of a Poisson process has zero uncertainty, whereas the variance is fixed in the Gaussian case for all sufficiently weak signal beam levels. What's going on is a redistribution of statistical uncertainty from the finite signal case to zero signal.

10.3.7 Heterodyne Interferometry

The heterodyne interferometer is the queen of laser-based optical instruments (Figure 10.3). It is demanding and occasionally temperamental, but unique in its ability to function in extreme situations, for example, very strong backscatter that would cause interference and possibly mode hopping. The backscattered light is frequency shifted away from the cavity transmission peak, so it bounces harmlessly off the output coupler. Heterodyne measurements are normally 3 dB worse than shot noise limited homodyne or baseband measurements due to the image frequency noise, just as in radios (see Section 13.8.7). With gas or solid state lasers, an 80–200 MHz AOD can get us pretty well out of the laser noise and into the shot noise region, at least as far as additive effects go.




Figure 10.3. Heterodyne interferometers: (a) Michelson with separate beamsplitter, (b) Michelson with Bragg cell as beamsplitter, and (c) Mach-Zehnder.




Figure 10.4. Optical single-sideband (SSB) mixer.

10.3.8 SSB Interferometers

In a heterodyne interferometer, the frequency mixing occurs at optical frequency, with a fractional offset of maybe 1 part in 10^6 if we're lucky, which makes the usual radio trick of filtering out the image frequency next to impossible. This costs us 3 dB SNR, and if that doesn't bother you by now, you haven't been paying attention. Two ways round this are to use homodyne detection (IF at DC), which makes the signal and image frequencies coincide, so that only one bandwidth's worth of noise gets through, or to build an optical SSB detector[†].

Optical SSB mixers work just like the ones in Section 13.8.7. Combine two beams in orthogonal polarizations (Figure 10.4). Split the combination beam without polarizing it, for example, with a patterned metal film (you may need a bit of spatial filtering to take out the pupil function weirdness this causes). Put one-half through a quarter-wave plate aligned to one of the polarization axes, detect separately, using Wollastons with their outputs subtracted, and add or subtract the result. You'll get a zero-background output that is purely USB or LSB. Whether all this foofaraw is worth it depends on how valuable your photons are.

10.3.9 Shot Noise Limited Measurements at Baseband

You really can get to the shot noise limit in a bright-field measurement, even at baseband, and even if the laser is an ugly mass of RIN. The trick is to use a laser noise canceler, which will get you 50–70 dB rejection of additive laser noise at low frequency, and

[†]Bo F. Jørgensen, Benny Mikkelsen, and Cathal J. Mahon, Analysis of optical amplifier noise in coherent optical communication systems with optical image rejection receivers. *J. Lightwave Technol.*, **10**(5), 660–671 (1992).

40 dB out to 10 MHz or so. Using the ratiometric output of the canceler will also get rid of the multiplicative noise to the tune of 70 dB from DC to some kilohertz, where you need it most. It's a pretty ugly laser indeed that needs more than that. See Sections 10.8.6 and 18.6.3 for the details.

10.4 MEASURE THE RIGHT THING

The best way to get better data is to measure something closer to what you need to know. A photoacoustic measurement, where a chopped laser heats a surface and a laser interferometer detects the surface motion, could probably be improved by using a focused ultrasonic transducer to detect the thermal signal instead. There are lots of other examples, for example, using a Nomarski microscope instead of a dark-field microscope to look at weak phase objects such as unstained animal cells.

Example 10.2: Nomarski DIC Microscope. The Nomarski differential interference contrast microscope is a white-light shearing interference microscope with zero path difference between beams (Figure 10.5). The shear is very small, less than one spot diameter, and is produced by a modified Wollaston prism (see Section 6.7.4) in the pupil of the microscope objective. It produces an image that is roughly

$$\begin{aligned}
 I(x, y) &= |h(\mathbf{x}) * E_x(x, y, z(x, y)) + e^{i\phi} h(\mathbf{x}) * \delta(x - \xi) * E_y(x, y, z(x, y))|^2 \\
 &\approx |h_2(\mathbf{x}) * E(x, y, z(x, y))|^2 + 2 \cos\left(\phi + \xi \frac{\partial \theta}{\partial x}\right) \cdot |(h_2(\mathbf{x}) * E(x, y, z(x, y)))|^2,
 \end{aligned}
 \tag{10.4}$$

where θ is the phase of E , so that an amplitude and a phase gradient image are superimposed. The phase shift ϕ between the images can be varied by moving the prisms laterally to get the best contrast.

With thermal light, it is important that the illumination passes through a matching prism; two passes through the one prism works in reflection, but you need a second one in transmission. (Why?)

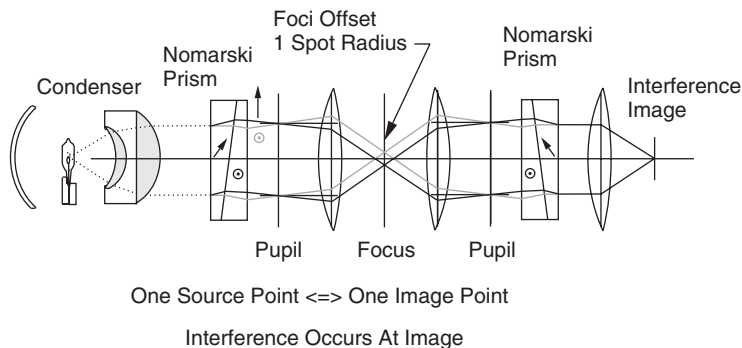


Figure 10.5. The Nomarski differential interference contrast microscope is a shearing microscope with a very small shear and zero path difference.

10.4.1 Phase Measurements

Always consider measuring phase, and think hard. Optical phase, modulation phase, any phase you can think of usually holds useful information. If your current measurement is phase insensitive, what phase would get you? And how hard would it be to get? Section 15.5 has lots on how to measure the phase of an RF signal accurately, so if you're doing a single-point AC measurement, it isn't a difficult matter.

Example 10.3: Attractive Mode Atomic Force Microscope (AFM). The attractive mode AFM of Figure 10.6 uses a vibrating cantilever to sense the force gradient between a very sharp tip and a sample, with a heterodyne interferometer to turn the tip vibration into phase modulation of the 80 MHz carrier. The force gradient acts like a change in the spring constant, and so changes the resonant frequency of the cantilever a bit. The usual way of detecting this is to drive the cantilever piezo with some frequency just above its resonance, and servo the tip-sample distance by forcing the amplitude to remain constant. Too low an amplitude means the tip is too close.

The tip has a tendency to stick to the surface because of the surface tension of the adsorbed water layer. Sticking of course stops the vibration entirely, causing a gigantic error signal, so the servo controller pulls the tip back and back until it lets go. Unfortunately, that plucks the cantilever very hard, making it ring strongly at its free resonant frequency. The servo thinks this means that the tip is way, way too far away, so it sends it crashing right into the sample again, and the cycle repeats until the tip is bashed into oblivion.

The solution is to measure the phase of the modulation instead of its depth—the phase of the phase, if you will. The modulation phase contains tuning information but is completely immune to this problem, because it can be tuned so that the free resonance is outside the passband, and the error signal is not enormously larger when the tip is stuck to the surface; if there's zero oscillation, there's zero output from the phase detector, which isn't too far from where it expects to be. You do need some amplitude sensitivity, because otherwise the tip will happily drag along the surface at zero oscillation amplitude, but

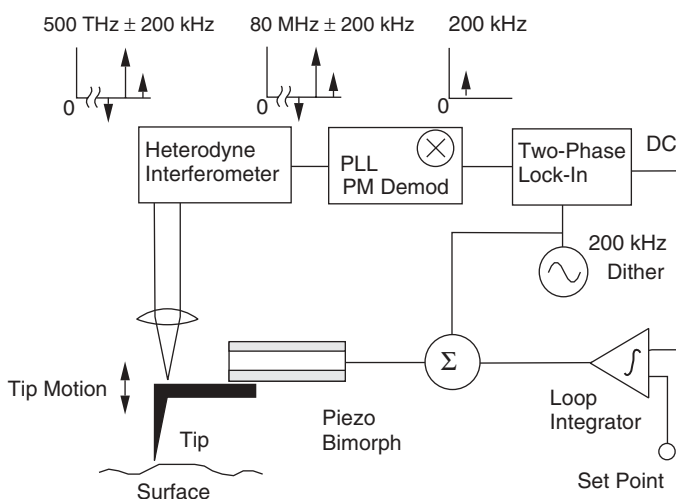


Figure 10.6. Attractive mode AFM.

that's just a matter of servoing slightly off the null of the phase detector. The increased resistance to tapping allows good measurements with the tip hovering closer to the sample, which translates to higher spatial resolution and sensitivity.

10.4.2 Multiple-Scale Measurements Extend Dynamic Range

Measure with a micrometer, mark with chalk, cut with an axe.

Anonymous

Remember the vernier dial? You still see them on inexpensive calipers, but that's about it nowadays. A vernier allows us to get three-figure accuracy from a two-figure eyeball. It works by spreading a small residual error out over many divisions, where it's easy to spot. We can do the same thing in optical measurements. The idea of this quotation actually works if you do it backwards; for example, good wide range position measurements can be made with an interferometer by counting fringes, or by subranging with a tunable source. If you have a 20 dB SNR, a straight measurement will get you an RMS accuracy of 7% in amplitude and 70 mrad in phase (Section 13.6.9), which is no great shakes. One solution is to use a two-wavelength interferometer, whose output is periodic at the beat length λ_B , where

$$\frac{1}{\lambda_B} = \frac{1}{\lambda_1} - \frac{1}{\lambda_2}. \quad (10.5)$$

Stepping λ_B by factors of 3 or so and measuring the phase residual at each step will give you the absolute distance to any desired accuracy, if you know the laser frequencies accurately enough. If you count fringes instead, that 20 dB is enough that you don't lose count under reasonable circumstances, so by counting fringes throughout the motion, you can measure long distances to an accuracy approaching $\lambda/10$ —once again, if you really know the laser frequency.

10.4.3 Fringes

Fringes and speckle corrupt all sorts of measurements, and especially bright-field laser measurements. With some work, these nuisances can be turned to good account; adding a piezoelectric mirror can make a phase-shifting interferometer, and a TV camera and video capture card can turn a HeNe laser into an electronic speckle pattern interferometry (ESPI) system that can measure surface deformations both in-plane and out-of-plane.

10.5 GETTING MORE SIGNAL PHOTONS

A bit of water and fertilizer can make the grass taller; an extra inch or two of height can make the difference between having a measurement and not having one. Be careful to keep the fertilizer out of your apparatus and especially your published data.

10.5.1 Don't Throw Photons Away

Thrift is a virtue in optics as in life. You can get more out of your photons by steering each of them to its appropriate bucket, rather than just throwing away the ones you're

not instantly interested in. In a flood-illuminated situation, where the light source is large in area or wavelength spread, you can't get all the photons into one resolution element. Use a staring detector that measures all the grass at once, rather than scanning one blade at a time.

It's okay to look at only one (band, pixel) combination at a time if your light source is doing the same (e.g., laser spectroscopy). There are lots of suggestions starting in Section 10.2 on how to reduce the background and get more photons.

10.5.2 Optimize the Geometry

The optimal measurement geometry depends on the interaction strength. Optimal sensitivity is reached at an absorption of order 1 (i.e., a transmission of order $1/e$). Some samples are so absorbing that we have to use attenuated TIR to look at their absorption spectra; we bounce our light off the hypotenuse of a prism, so that it is totally reflected from the interface, making the effective path length in the absorber very small.

Others are very weak, for example, the O₂ overtone absorption line at 760 nm, which is useful for measuring oxygen concentration with a diode laser. Those need a multiple-pass cell or a resonant cavity to provide enough interaction length.

10.5.3 Use TDI

Time delay integration (TDI) is a sometimes useful scheme for making 1D CCD detectors less bad at imaging areas. The idea is to make a narrow 2D array, maybe 20 pixels wide and 4096 long, and shift it in the frame direction at just the right speed that the motion of the image and of the photoelectrons is synchronized. Ideally, that gets us 20× more photons per pixel with no other worries except a longer integration time. Unfortunately, it isn't easy to make it work so neatly; the tolerances on the scan speed, clock rate, and the geometrical distortion of the lens are very tight, so the SNR improvement is less and the image degradation more than you might think. Altogether, TDI arrays are for people with big budgets and lots of patience.

10.5.4 Consider OMA Spectroscopy

The light from a grating is spread out linearly with wavelength, and so can be sampled with a 1D multielement detector such as a CCD or photodiode array, an arrangement called an optical multichannel analyzer (OMA). These have much better efficiency than monochromator types, but as there is no exit slit or second grating for protection against stray light, they are much poorer when the dynamic range gets large. Their typical 1024 element detector provides about 5 Å resolution across the visible; since silicon detectors are insensitive beyond 1 μm and IR arrays are expensive, insensitive, or both, nearly all OMAs are visible-light devices, with a detection range of 400–800 nm. Most OMAs operate at room temperature, which raises the dark current of the detector by orders of magnitude, and tend to have 12 bit ADCs, so that their dynamic range wouldn't improve much even if they were cooled. High performance OMA spectrometers with cooled detectors can provide shot noise limited spectral measurements, at least in dim light. Since we're still limited by the étendue of the entrance slit, that still isn't that great, but it's fine for lots of things.

10.5.5 Consider Slitless Spectroscopy

If we're doing spatially resolved spectral measurements (e.g., spectral reflectance to get film thickness), we can get more photons by removing the output slit of the monochromator and shining the dispersed light directly on the sample; provided that it's in focus, every point on the sample will be illuminated with a definite wavelength. Scanning the grating makes the pattern move across the sample. A camera looking at the sample will therefore get lots more photons, and the measurements can be figured out afterwards in software.

10.5.6 Consider Fourier Transform Infrared (FTIR) Spectroscopy

We saw in Section 1.3.8 that the power spectrum of a signal is the Fourier transform of its autocorrelation. A two-beam interferometer with nonzero path length produces a DC term, plus an AC interference term whose time average is the autocorrelation of the input radiation. Accordingly, measuring the autocorrelation at appropriately chosen path differences and taking the Fourier transform numerically gives us the power spectrum of the radiation.

This approach has two important benefits, which in spectroscopy are customarily known as the *Felgett advantage* (frequency multiplexing) and the *Jacquinot advantage* (étendue). The Felgett advantage comes from all wavelengths being detected at a duty cycle of 100%, which improves the shot noise—we don't block any photons with input and output slits of a spectrometer. The Jacquinot advantage comes from the fact that aperture angle enters only quadratically into the optical path length of the interferometer, and hence into the frequency resolution, whereas there is a linear trade-off between resolution and slit width in a monochromator system; thus the FTIR's detection NA can be widened enormously. It's as if the slit width were more like slit length; instead of a long narrow slit, you can use an aperture as wide as it is tall, and you can imagine how much brighter that is. The two produce an efficiency advantage on the order of 10^5 in a typical FTIR versus a Czerny–Turner, and 200 compared with an OMA type. There is also no grating order overlap to worry about, though you do have to move the interferometer mirrors very precisely, scanning is relatively slow, and the shot noise of bright components can corrupt the measurements of weak ones.

You get all this with a single-element detector, out there in the IR where even poor detectors cost a lot. It's a remarkable measurement idea altogether.

Aside: FTIR Errors. The simple account of FTIR operation given above ignores the infrared self-luminosity of the detector, mirrors, and beamsplitter. In the mid- and far-IR, FTIRs have to take account of these effects, which require careful calibration and a significant amount of thought.

10.5.7 Use Laser Scanning Measurements

With all that has gone before about widening our apertures and using staring sensors, it may seem contradictory to recommend a flying-spot laser scanner as a good way to get more signal, but nevertheless, it is. The reason is that although each point on the sample sees the beam only once per frame, a duty cycle of perhaps 10^{-6} , the beam is very very bright during that time, so its total dose per unit area is the same, and hence the counting statistics are too, for constant collection NA.

The other benefit is that the detector sees a very brightly illuminated sample point at all times, and the incoherent optical background signals get detected with that 10^{-6} duty cycle too, assuming of course that the baffles on the receive side are that selective, or that there is some heterodyne gain lifting the signal up. The detector side can have a high NA but a lowish étendue, because it is being descanned.

The very short dwell time per pixel also eliminates speckle (because only one sample point is being illuminated, it has nothing to interfere with to make speckles). There is no blur due to finite integration time, because the dwell time per pixel is 10^{-6} frame times. (Grass has to be blowing around pretty fast to cause much blur in a 50 ns integration time.) Furthermore, flying-spot measurements are naturally suited to signal averaging.

Very high contrast samples are hard to work with in a full field system, because of scattered light and possibly saturation, and things like Raman spectroscopy have such an enormous dynamic range that they are not possible without rejecting the unshifted beam with a sharp interference filter. Due to angle tuning, filtering is a lot easier in the collimated light you get after descanning. Selectivity is far better with scanning because the interfering area is not illuminated—no glints, no speckle—you name it, it's gone.

10.5.8 Modify the Sample

Sometimes you can modify the sample to make it easier to measure, either by preparing it differently, changing its size or concentration, or, as in this example, controlling its motions.

Example 10.4: Sub-Doppler Spectroscopy. Example 1.1 describes two-photon Doppler-free spectroscopy. It's an optical hack for eliminating Doppler broadening, which is limited to two-photon transitions. Other methods usually concentrate on making the molecules' radial velocities very small. This is done by supersonic cooling, which makes the random velocities small, or by the atomic fountain, where very cold gas rises and falls under gravity, with the measurement occurring at the peak of the motion. These are cryogenic, high vacuum techniques. Another technique is trapping, where molecules are held in a small volume for a long time. Ions can be trapped electromagnetically, and neutral species with *optical molasses*. In the optical molasses approach, the molecules are held at the intersection of three sets of counterpropagating beams, tuned just slightly above a strong resonance of the molecule. A molecule that is moving will be Doppler shifted closer to resonance and will absorb photons from the beams opposing its motion. It will reradiate into 4π steradians, so that the expectation value of the momentum change is just the momentum of the laser photon. This slows the particle down. There's a certain irreducible residual motion, since the random character of the photoemission means that the resulting impulse never quite cancels the particle's momentum.

Optical molasses is not a true trap, because there's no potential well, just optical friction. Particles can diffuse out of it, but it's a good way of making very slow-moving molecules.

10.5.9 Corral Those Photons

There are often geometric limits to how many photons we can easily recover. If we're using a small light source on a specular object 10 meters away, we'd need a truly huge

collecting system to achieve decent photon efficiency because most of the light is reflected away. A partial answer is retroreflective tape, which collects all the photons that hit it and bounces them more or less back into their source; the difference in signal power can easily be 60 or 70 dB (see Sections 7.8 and 11.8).

Aside: Video Techniques. Talking about staring measurements inevitably brings up video. Properly designed CCD detectors are excellent devices, but video cameras are another matter. Their automatic gain and black level controls corrupt our attempts at calibration, which makes intensity measurements difficult. (See Sections 3.9.1 and 3.9.14.)

Besides mapping intensity, cameras are used a lot in full field interferometric techniques (phase shifting, moiré, and holographic), and with *structured light* measurements. For example, consider a fan-shaped beam projected on a curved surface and looked at with a camera placed out of the plane of the fan beam; wiggles in the image of the line correspond to variations in height. With a camera as detector, even quite large variations in the surface (e.g., tubes of different diameters) can be accommodated without changing the setup, which is a very valuable attribute, but the attainable precision is poor.

In these situations, the efficiency and repeatability of the camera as a photon detector are of secondary concern, but the geometric distortion caused by the camera lens is often not properly accounted for. This distortion is a few percent for most lenses, and as much as 10% in a wide-angle, which makes a precise geometric measurement using a camera lens a nontrivial exercise. This is particularly so if the focus is being adjusted, or if the surface is going in and out of focus due to large curvature. If the camera is seeing the surface at an angle, the magnification will be a function of position; this is responsible for the familiar “keystone” effect with projectors.

The best way of using a camera for surface measurements is to keep the axis of the camera lens normal to the surface, to eliminate keystone and minimize defocus, and carefully calibrate the geometric distortion of the lens at different positions and different states of focus. A telecentric lens makes this dramatically easier, since to leading order the magnification is independent of focus; you get the equivalent of an isometric projection—see Section 9.3.7. Failing that, a careful 3D calibration is needed, which is sufficiently painful that you won’t want to do it twice.

10.6 REDUCING THE BACKGROUND FLUCTUATIONS

It really isn’t necessary to get rid of the background altogether, as long as its fluctuations aren’t limiting the measurement unnecessarily. Getting bright-field measurements to do as well as ideal dark-field ones with the same NA requires getting the background fluctuations down to the shot noise. Fortunately, that’s perfectly possible a lot of the time.

10.6.1 Beam Pointing Stabilization

Active-optics beam pointing stabilization systems are based on shining a sample of a beam into a quad cell and tipping a mirror or tuning AODs to keep it still. In a limited bandwidth, they can do a very good job. One limitation is the interaction of angular and position wobbles; focusing the beam on the quad cell with a lens will make the system

ignore wobbles in position but correct pointing errors. Operating near the null point makes the beam pointing insensitive to RIN, so that with enough control bandwidth, only shot noise and beam profile changes are left to cause errors.

10.6.2 Beam Intensity Stabilization

You can get feedback controlled attenuators that adjust the intensity of the laser beam to keep it constant, e.g., Noise Eaters. They're great when the beam intensity is what matters, for example, in dark-field measurements, laser heating, distribution of beams to several places, or pumping other processes. They can do a great job of stabilizing dark-field systems, because the shot noise limited SNR is much lower than that of the full beam (so it's easier), and because they get rid of the low frequency multiplicative noise, which is usually the main culprit in dark field.

Unfortunately, they are not able to get the noise down to the shot noise of the beam, which is why they're not a complete solution. There are two reasons why. First, they're feedback devices that rely on the loop gain to suppress noise (unlike noise cancelers and differential measurements, which rely on matching instead), and they don't have enough gain and bandwidth. The second reason is more fundamental and affects noise cancelers and ratiometric measurements as well; if we're using beam A as a reference to stabilize beam B , the shot noise of the two beams is uncorrelated. Thus by a standard propagation-of-error calculation (you remember—take partial derivatives and multiply by the corresponding rms noise, then take the RMS sum of all contributions), assuming we're well inside the bandwidth, the shot noise limited CNR_{out} of the stabilized beam obeys

$$\frac{1}{\text{CNR}_{\text{out}}} = \frac{1}{\text{CNR}_A} + \frac{1}{\text{CNR}_B}, \quad (10.6)$$

where CNR_A and CNR_B are the shot noise limited CNRs of the two beams. There is no way to stabilize a beam down close to the shot noise this way without wasting a whole lot of it; a 50% loss will get you to 3 dB above the shot noise.

10.6.3 Photocurrent Stabilization

For bright-field measurements, we have to cope with the shot noise of the beam anyway, so there's no need to stabilize the strength of the *beam*, just the photocurrent. Stabilizing the beam is one way to do that, but it's the hard way. An alternative is to use a differential or ratiometric measurement of the photocurrent, where a comparison current derived from a sample of the source beam is subtracted or divided out. Since both beams will see the same fractional noise, this will leave a quiet result. We'll leave subtraction for Section 10.8 and the detailed discussion of the principle to Section 18.6.3. This is a very powerful technique that should be considered in any laser measurement.

10.6.4 Ratiometric Measurements and Phase Errors

The ratiometric approach promises to correct for source intensity variations, and hence eliminate both additive and multiplicative source noise at a stroke. Examples of ratiometric techniques are analog division, deriving the ADC reference voltage from the source

power rather than from a voltage reference, or using the log ratio output of a noise canceler. We'll see in Sections 15.11 and 18.6.2 that divider ICs have very limited dynamic range and are slow and nonlinear, so they aren't as attractive as they look, though they are adequate for some things, especially slow measurements with quiet sources. They work OK for taking out 120 Hz from tungsten bulbs, for example.

For slow things like that, the trick of deriving the ADC reference voltage from the comparison photocurrent is also competitive; you do have to worry about matching the phase shifts in all the signal paths, though, or you may wind up making things worse rather than better—the more circuitry there is in the way, the harder the matching is to get. Since the phases have to be matched to 0.01 radian (0.6°) to get 40 dB noise suppression, and op amps have phase shifts of about f/f_c radians (where the -3 dB frequency f_c is not very well controlled), this gets hard very rapidly for signals above 1 kHz if each current has its own amplifier. Dividers have separate sections for numerator and denominator, so they suffer from this even if you wire the two in parallel for test purposes. It's possible to put in a phase tweak, but since it doesn't usually have the right frequency dependence, you can't get that far with it and it isn't especially stable with time and conditions; we're usually talking about subpicofarad strays, and those will change with time, temperature, and humidity.

The logarithmic output of the laser noise canceler doesn't have these problems, for two reasons: (1) the photodiodes are basically wired in parallel for AC, so the strays and phase shifts have little opportunity to differ; and (2) the balanced subtraction gets rid of the additive noise at all frequencies, irrespective of the feedback loop bandwidth, so all the loop has to cope with is the multiplicative noise, which is usually some orders of magnitude smaller and concentrated at baseband where the loop has an easy time. See Sections 10.8.6 and 18.6.3 for the details.

10.6.5 Changing the Physics

Sometimes we can solve a problem by changing the measurement physics. For example, consider a spectrometer for measuring impurity concentrations in ultrapure water. To get decent sensitivity, we'll need a long interaction length, which typically means a multipass cell, with mirrors bouncing the beam back and forth many times through different parts of the sample cell. However, even ultrapure water has a strong scattering background from Rayleigh scatter and bubbles, which depends on the temperature and pressure histories, dissolved gas concentration, hydrodynamic vibration, and other things that we don't regard as data. Since the scattered light misses the detector, it shows up as a spurious absorption signal and is difficult to correct for.

Now consider the same sample cell inside an integrating sphere, with source and detector shielded with baffles. The white walls of the sphere cause the light to make many passes through the sample before being absorbed, so the extra sensitivity of the multipass arrangement is preserved. However, the same walls homogenize all the light into π steradians every 15 cm or so, so the scattered photons are no more likely to miss the detector than the unscattered ones, and the scatter signal goes away almost entirely. The photon efficiency of such a measurement is of course much lower than in the multipass cell, where almost the entire beam goes into the detector, but a few inexpensive photons are a small price to pay for a stable and sensitive measurement.[†]

[†]Private communication from Mark Johnson.

10.7 OPTICALLY ZERO BACKGROUND MEASUREMENTS

An optically zero background measurement is one where the unaltered illumination beam doesn't get to the detector—in other words, we're talking about dark-field measurements of one sort or another. These are great if you're either in the photon-counting regime or the $>1 \mu\text{A}$ regime, but not so nice in between. The analogy is climbing the building to measure the grass; you probably get a better measurement, but your footing may be a bit precarious.

10.7.1 Dark Field

Ordinary dark-field microscopes aren't actually all that dark because of scatter, but they work pretty well visually. A *doubly-dark-field* measurement is one that uses dark-field illumination plus crossed polarizers, looking for specific kinds of problems such as diagonal scratches or particles that are diffuse scatterers.

10.7.2 Fringe-Based Devices

Two laser beams producing a fringe pattern within the sample space can do some very useful things. The simplest one is time-modulated scattering, as in *laser Doppler velocimetry* (LDV), which sends aerosol particles along $\Delta\mathbf{k} = \mathbf{k}_2 - \mathbf{k}_1$ and measures the resulting amplitude-modulated light pulses with one or two photomultipliers in a dark-field geometry; the frequency of the pulses gives the velocity, and (interestingly) the phase shift between different scattering directions tells something about the particle size. Using variable-pitch fringes to encode measurement data is also possible; if the LDV fringes were made from two spherical waves instead, it would be possible to work out where the particle went through the beam by the change of frequency with time.

A more complicated one is the *photochemical grating*, where the fringes pump some process in the sample, making a grating structure that can diffract a probe laser beam, with the diffracted beam being the signal. This is a bit like a thermal lens measurement, in that the grating can be made in the IR and the probe in the visible. The biggest advantages are that the signal beam winds up going off away from the other beams and that its amplitude is zero when there's nothing to detect. That makes it easy to detect very small signals. Fringe-based measurements are a spatial version of modulation mixing (see Sections 10.9.4 and 13.10.6).

10.7.3 Monochromator-Based Measurements

The classical way to do spectroscopy is a scanning Czerny–Turner spectrometer with a single-element detector, looking at the emitted or transmitted light from some sample. This is an enormous waste of good photons. Nobody in his right mind would do that today except with very bright light, or in very special circumstances, for example, Raman spectroscopy where the close proximity of the enormous unshifted line makes double and triple Czerny–Turners very helpful.

10.7.4 Fluorescence and Photon Counting

Fluorescence solves the source shot noise problem and ideally has the same sensitivity as a dark-field or coherent detection measurement with the same detection NA.

The improvements in practice come from being able to reject the incident light via highly effective optical filters, and being able to distinguish one sort of molecule from another—the ones we care about label themselves for us. Thus fluorescence measurements can be highly sensitive. Especially with the new development of negative electron affinity (NEA) photocathodes (see Section 3.6.1), a fluorescence detection instrument can detect more than 25% of all photons emitted by the fluor.

The bad news is that the quantum yield of the fluorescence process—the ratio of all fluorescent photons emitted to photons absorbed—is often very low, in the 10^{-5} range, although some highly efficient fluors such as rhodamine laser dyes are a lot better, a few percent to tens of percent. On top of this, real samples are usually optically thin at the pump wavelength, so only a small percentage of incident photons are absorbed. This means that photon-counting detection using PMTs must usually be used.

Laser dyes and other organic fluorophores tend to die at around 10^6 emitted photons per molecule, because the upper state of the fluorescence occasionally decays by electron transfer (i.e., chemical reaction).

Fluorescence methods are best used where the desired species can be made to fluoresce efficiently, usually by being labeled with a specially chosen marker such as fluorescein. Thus they are used extensively in biochemistry. Other uses are atomic fluorescence measurements, where an atom or molecule in the gas phase is pumped with one wavelength and detected at another. By chopping the pump beam, the desired signal can be made to appear at AC, with the background left at DC. Of course, this assumes that the pump has no detectable energy at the probe wavelength, and that there is no residual fluorescence of other components in the system at the probe wavelength. Colored glass filters and plastic or glass components exposed to pump radiation must be scrutinized especially carefully.

Putting filters in the wrong order can cause enough filter fluorescence to get through that the measurement becomes impossible—see Section 5.5.1.

10.7.5 Nonlinear Measurements

A nonlinear generation measurement produces no output signal with no sample, which makes it a dark-field measurement. In fact, a lot of the time they produce no output even with a sample, but you can make them if you try hard. Examples are second harmonic generation (SHG) and Raman spectroscopy. Raman is a special case, since separating out the unshifted light is a big job—it's very bright, and very nearby.

10.7.6 Nonoptical Detection

Some kinds of samples are kind enough to work as detectors for us; examples are photoionization spectroscopy, photoacoustic measurements, and some ultrafast ones such as picosecond electro-optical sampling, where the output is a short pulse on a transmission line. These count as zero background only because we think of the whole output as the signal.

10.7.7 Active Fringe Surfing

You can get a more-or-less dark-field measurement by actively stabilizing an interferometer on a dark fringe. You can do this by dithering one arm slightly with a liquid crystal

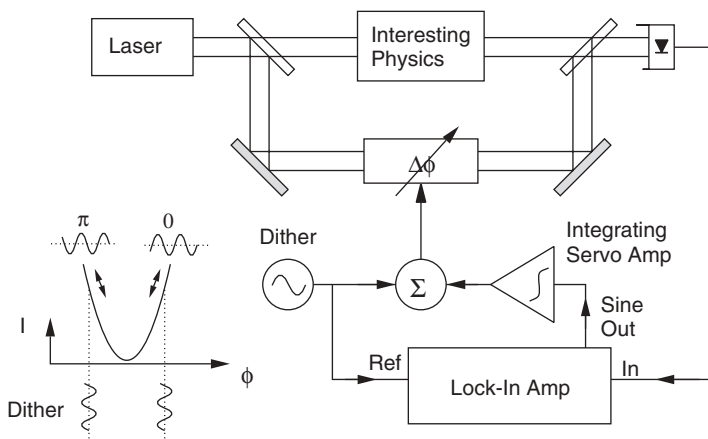


Figure 10.7. Fringe surfing: the phase of the fundamental signal changes 180° going through the null, so a lock-in can keep us right on the null.

phase shifter and using a lock-in amplifier to sense when the fundamental signal goes through 0, as shown in Figure 10.7. The fundamental signal goes as the first derivative of the fringe slope, so it goes through zero and changes sign at the null; thus a lock-in produces a properly bipolar error signal, and a simple integrator will stabilize the operating point. You wire it up, close the loop, and *presto*, you're sitting on a dark fringe. Using the control voltage as the output, signals well inside the feedback loop bandwidth will be detected, and those far outside will be rejected. Using the lock-in output, it's the other way round. Either way, additive laser noise (including shot noise) will be greatly reduced, because most of the light is going somewhere else.

Everyone loves fringe surfing (probably because most of us have invented it independently at some point, see Section 8.16.3). It's deceptively easy; you see, the surfing part is well behaved—it's the dark fringe itself that is noisy and poorly defined. When you think about it, it's really pretty silly to servo on a dark fringe, because the SNR there is zero (not 0 dB, *zero*). It's also a tough thing to do in a real instrument, because it's mechanically flaky and you'll lose lock if the phase shifter goes out of range, which it's going to, probably at an inconvenient time. Phase shifters are nonlinear, so even if it comes into lock somewhere else, it isn't simple to stitch together data from the different operating points (see Section 8.16.3).

The other things to remember are that the coherent background you get from working off the null actually improves the SNR a great deal, as we saw in Section 10.3.5, and that the laser intensity noise can be got rid of with a laser noise canceler. Thus if you must surf fringes, it's much better to do it in bright field (for more details, see Sections 10.8.3 and 10.8.6).

10.7.8 Polarization Tricks

High school and undergraduate optics labs are full of crossed-polarizer experiments, where some sample such as cellophane tape or a plastic model of some architectural structure is placed between crossed polarizers, and the colors of the output are photographed. This is a famous old trick for finding stress concentrations in 2D models.

Dithering the polarization with a liquid crystal retarder or a Pockels cell allows the same sorts of tricks as in the previous section, and it's more useful since a lot of the time the background light isn't coherent with the signal, especially if we're using a thermal source with appreciable NA.

10.7.9 Optical Time Gating

In a pulsed measurement, the background can often be eliminated by time-gating the detection side appropriately, for example, an OTDR, where the huge pulses from the ends of the fiber and from the connectors will saturate the detector and so mask nearby signals. An APD or MCP detector can be turned on and off very rapidly, so this isn't hard to do.

10.8 ELECTRONICALLY ZERO BACKGROUND MEASUREMENTS

There are lots of methods for reducing the influence of noise in measurements, many of which are explained in Chapters 13 and 17, and most of which boil down to filtering in one form or another. Here we'll concentrate on how to keep that gunk out of the data in the first place. These methods get rid of the excess noise of the background but are powerless to reduce the shot noise. They are thus quite competitive with optical background reduction approaches for coherent bright field, but distinctly inferior for dark field or when there's an incoherent background. The analogy here is looking from a neighboring building or a helicopter; if the two buildings are swaying in the wind, the measurement isn't as stable, although much better than from ground level, because the angular size of the grass tuft is much bigger.

10.8.1 Polarization Flopping

In polarizing measurements such as Nomarski, a ferroelectric liquid crystal half-wave plate between the Nomarski prisms will invert the image contrast when its state changes (see Example 7.2). Using a CCD detector or instrument-grade camera (*not* an auto-gain, auto-black-level video camera) and flopping the phase on alternate frames allows separating the amplitude and phase information; summing alternate frames gives amplitude, and subtracting them and dividing by their sum gives $\cos \phi$. You need more phases than that for a real phase measuring system (at least three phases spaced around the unit circle), but just separating out amplitude and phase is often good enough.

10.8.2 Electronic Time Gating

If the pulse interference can't be gated out optically, there are other things we can do, for example, an IF noise blanker, which looks for noise pulses above some threshold, and cuts off the IF gain when one arrives; this is how some car radios deal with ignition noise, for example.

If the signal pulses have a low duty cycle, then turning off the receiver in between reduces the noise power we detect by a factor of the duty cycle; of course, the bandwidth has to be wide enough to reproduce the pulse, but we're normally stuck with that anyway. The ping-pong gated integrator idea from Section 15.5.6 is one example; even though

the noise is being integrated 100% of the time, if there's no signal pulse, we just throw the result away, so it doesn't reduce the measurement SNR.

10.8.3 Nulling Measurements

Fringe surfing doesn't have to be around a dark fringe; since an interferometer has two output ports, we can servo the phase of one arm to keep the two output powers exactly equal. Although the shot noise of the two arms of course adds, this isn't a serious limitation because of the Bright Field Rule of Section 10.3.6.

It is also possible to use this idea to null out other effects. The long-distance particle counter of Figure 10.8 is based on extinction (Figure 10.9) and uses an AOD at a pupil to make a laser beam wiggle back and forth sinusoidally. Because the AOD is at the pupil (where the beam comes to a focus in this case), when the beam is recollimated by the next lens, its motion becomes a pure translation, with no angular scan (it is intended for use in a 30 m long by 15 cm high belt oven full of hydrogen at 1000 °C). A particle will be in the beam part of the time, and out of it part of the time; thus it will cause a differential extinction signal at the scan frequency. Dust on the optics, minor misalignment, and varying Fresnel reflections cause the output to have a large background component at the fundamental, which will mask the small (10^{-6}) transient extinction signals of the particles. It can be nulled out with a lock-in sensing the fundamental, connected to an integrating servo amplifier that controls the center frequency of the scan. If there's

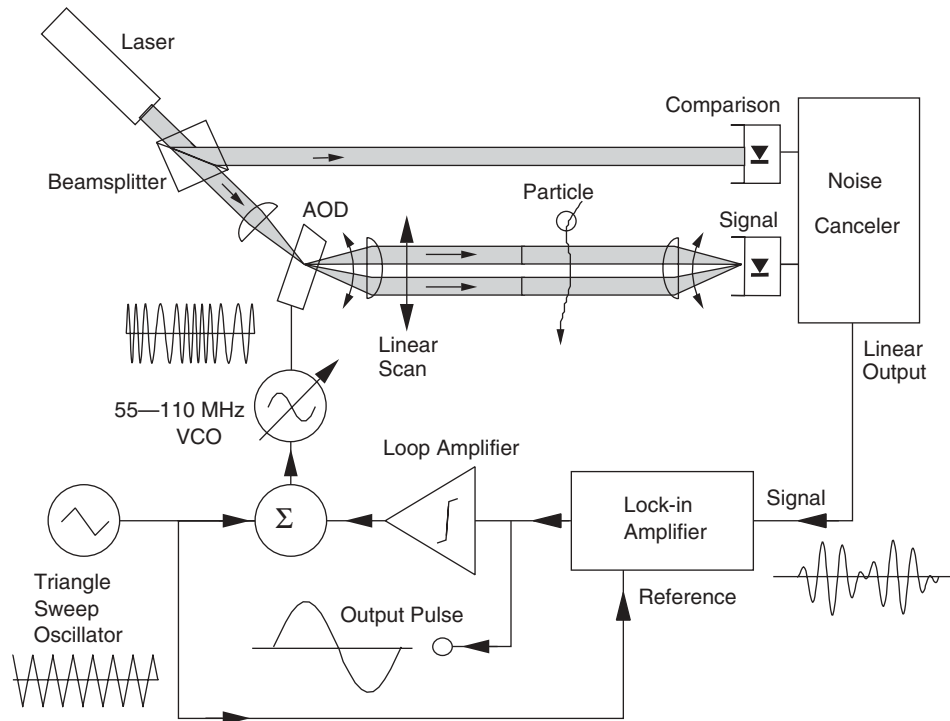


Figure 10.8. This acousto-optically scanned long-distance particle counter uses the parabolic shape of the diffraction efficiency curve of the AOD to null out any static background at the fundamental (e.g., due to vignetting or dust).




Figure 10.9. Operating principle of the extinction system of Figure 10.8: the sign of the fundamental depends on which side of the diffraction peak it's on, so nulling gets rid of it.

more extinction on the positive half-cycle, the center frequency will slide slightly down the low frequency skirt of the diffraction efficiency peak, causing just enough variation to cancel the background. The lock-in's output is a highpass filtered version of the extinction signal from the particles. This is dramatically better than surfing a dark fringe, because the SNR is so much higher. As usual, analog techniques are much superior in nulling measurements—simpler, faster, quieter, and more accurate. Null first and digitize second.

10.8.4 Differential Measurements

Interferometers have differential outputs, as we've seen, and there are lots of other cases, too. Dual-beam spectrometers, where one beam passes through the sample and one doesn't, are another example. Subtracting two carefully balanced photocurrents ideally cancels the source excess noise exactly, at least if the noise in the two beams hasn't been decorrelated by vignetting or etalon fringes (which turn FM noise into AM). In the Nomarski case, the normalized difference signal corresponds to the phase, and this is true of other interferometers as well, when they're working near null, and the total phase excursion is small. As we've seen elsewhere, the problem is how to keep them in balance.

10.8.5 Other Linear Combinations

Mere differences don't exhaust the usefulness of linear combinations. An interferometer has output beams, $a(1 + \cos \phi)/2$ and $a(1 - \cos \phi)/2$, so that their sum is a and their normalized difference is $\cos \phi$. Operating near $\phi = \pm\pi/2$ gives an electronically zero background phase measurement, even if we don't divide out the amplitude, and is linear over a range of $\pm 5^\circ$ or so, which is frequently lots, especially in common-path

interferometers where the phase drift is small. The beam intensities will need a nulling adjustment.

10.8.6 Laser Noise Cancelers

As we've already noted in Section 10.6, the laser noise canceler is capable of making bright-field laser measurements equivalent to ideal dark-field ones in practice, as they are in theory. A good example of this is tunable diode laser absorption spectroscopy, as shown in Figure 10.10.[†] The diode laser is tunable over a wave number or two by changing its bias current, but in the process its output power changes by a factor of 2, which makes gigantic additive and multiplicative backgrounds. An extremely simple optical system, consisting of a Wollaston prism, one polarizer, a sample cell, and a noise canceler, can make absorption measurements down to the 10^{-6} range even in the face of huge variations like this.

The circuit details are in Section 18.6.3, but the main point is that the automatically balanced subtraction gets rid of the additive noise, and the log ratio output takes this additive-noise-free signal and gets rid of the low frequency multiplicative noise as well, so that the peak heights are independent of the laser power, and does it without the serious noise and accuracy limitations of analog dividers, or the bandwidth, phase shift, and dynamic range problems of digital methods based on numerically dividing the two detected signals. Etalon fringes can be subtracted out by admitting air into the sample cell, which pressure-broadens the real absorption into oblivion, leaving only the etalon effects, which can then be subtracted out. The before and after pictures are shown in Figure 10.11. The two just-discernible bumps on a huge sloping background have turned into a well-resolved spectrum with ppm sensitivity and a flat baseline.




Figure 10.10. Tunable diode laser spectroscopy with a laser noise canceler.

[†]Kurt L. Haller and Philip C. D. Hobbs, Double beam laser absorption spectroscopy: shot-noise limited performance at baseband with a novel electronic noise canceler. *Proc. SPIE* **1435**, 298–309 (1991). (Available at <http://electrooptical.net/www/canceler/iodine.pdf>.)




Figure 10.11. Before and after pictures of the signal from the tunable diode laser spectrometer: (a) signal-averaged photocurrent, showing the huge background, (b) detail of (a), and (c) the same spectral region as (b), seen by the log ratio output of the noise canceler; the largest peak has about 1% absorption. Other peaks were observed at 12 ppm absorption with good SNR.

Unfortunately, noise cancelers don't work as well with fiber systems, because of fiber's nasty way of producing huge amounts of coherence fluctuation noise through etalon fringes and double Rayleigh scatter (see Sections 2.5.3 and 8.5.13). Different fibers, or even different polarizations in the same fiber, have different fringe phases, so the noise they produce is uncorrelated and hence uncancelable. Once again, this is a beam problem, not a circuit problem—the noise canceler gives you a really great measurement of the uncorrelated part of the noise. This probably isn't what you want, but it's as much as any technique can do with such badly corrupted beams.

10.9 LABELING SIGNAL PHOTONS

Painting the grass fluorescent orange is useful in two ways; we can wait until night time, shine a black light at it, and look for its fluorescence through an orange filter (a dark-field technique), or we can do it in daylight and do some sort of spectral differencing

measurement, where grey things like concrete cancel out but orange things don't (a photon-labeling technique).

Labeling good photons may let us subtract off the background photocurrent, but not of course its shot noise—it's a lot like the electronic zero background approaches, which isn't too surprising when you think about it. Like them, photon labeling is second best unless we're doing a shot noise limited bright-field measurement. There are lots of ways to label photons, so if you can, try to use ones that preserve your SNR by not wasting most of the power in the carrier.

10.9.1 Chopping

The naïve approach to laser-based AC measurements is to chop the laser beam and use a lock-in to look at the amplitude of the chopped signal. It does avoid electronic drift and $1/f$ noise, but that is very rarely the problem. It doesn't help optical stability one bit, because all the low frequency laser junk just gets translated up to the chopping frequency; AC measurements help only when they treat the signal and noise differently. If the reason this doesn't work is unclear to you, you may wish to read Chapter 13 carefully before proceeding.

If you can chop the right thing, though, this can be a good technique. It's rarely the best available, because the frequency shifts you can get are modest, the photon efficiency poor, and mechanical jitter high. Acousto-optic and photoelastic modulators make better choppers than choppers do, at least with lasers.

Chopping the right thing means chopping the sample. One example is a pump/probe measurement, where a chopped beam pumps the sample and another beam measures the changes so produced. Another would be scanning on and off the sample, or moving the sample very rapidly in and out of the beam, e.g. with kilohertz microfluidics. Chopping isn't just for laser experiments; infrared astronomy would be almost impossible without it.

Example 10.5: Chopping Secondary IR Telescopes. Infrared astronomers have a tough job. From a 270 K planet with a 100 K sky temperature, using lousy detectors, they have to see the IR emission from a protostellar nebula a kiloparsec away or a primeval galaxy on the rim of the universe. These signals are far below the sky brightness level, and that sky brightness changes with time and atmospheric conditions (the situation is redressed in part by being able to do it during the day as well as at night—the λ^{-4} dependence of scattered intensity chops the Rayleigh scatter down pretty far in the IR). The solution they adopt is the *chopping secondary*. The telescope's secondary mirror flips back and forth very slightly at a rep rate of a hertz or so. The detector accordingly sees the source alternating with a patch of sky a few arc seconds away, whose background level is very nearly the same. The sky brightness signal is thus DC, and the source signal is AC. Using two detectors works even better, since no signal photons are lost, and the background gradient can be removed too.

One really important caveat is to make sure that you have enough actuator and detector bandwidth that the chopped signal has a chance to settle down on each cycle. Otherwise you'll find yourself measuring the slew artifacts rather than the signal.

10.9.2 Scanning

Some classes of measurement rely on making the sample look different at different times, so that the signal alone is modulated in a distinctive way. The simplest one is scanning (see Section 7.9). Scanning can be done in x or λ , making the signal periodic at the scan frequency; signal averaging will select the signal and leave the low frequency noise and instability behind (Section 13.10.4). If it's possible to keep the optical system constant and scan the sample, either by moving it on an accurate translation stage or changing some other sample property, the background is more nearly constant, and hence easier to get rid of. Any scanning artifacts (e.g., changing etalon fringes, vignetting, or laser power variations) will usually be indistinguishable from signal, so we have to go further.

Continuous scanning produces an ordinary sampled measurement (see Section 17.4.3), but a stepped scan requires that the electrical bandwidth is wide enough so that the signal settles accurately between samples. A triangle-wave scan is better than a sawtooth, because settling artifacts show up clearly when the trace and retrace don't line up with each other. A combination of scanning and chopping can work well too, especially if you have a rapidly moving baseline; if you have random access to your sample points, instead of scanning $p_0, p_0 + \Delta, p_0 + 2\Delta, \dots, p_0 + (N - 1)\Delta$, you can go $p_0, p_0 + \Delta, p_0 + 2\Delta, p_0 + 3\Delta, p_0, p_0 + 3\Delta, p_0 + 4\Delta, p_0 + 5\Delta, p_0, \dots$. That way, every frequency is measured twice at different times in the sampling window, and there are a whole bunch of baseline correction points; discrepancies due to settling funnies become obvious, and the baseline error can often be fixed very accurately. Do keep the bandwidth wide enough for all that jumping around, though.

10.9.3 AC Measurements

The most common way of labeling photons is by making them arrive at some modulation frequency sufficiently far from DC that the worst junk can be filtered out. Heterodyne measurements are great at this, but modulation techniques are widely used too, for example, FM derivative spectroscopy, in which a small, very fast dither i_{dither} is superimposed on the diode laser bias current I_{bias} . (See Section 13.1 for more on AC versus DC.)

Example 10.6: FM Derivative Spectroscopy. By Taylor's theorem, if the tuning sensitivity is K and the total photon efficiency of the system (including sample absorption) is $A(\omega)$, the detected photocurrent is

$$I_{\text{photo}} = BPA \approx B \left(\sum_{j=0}^{\infty} \frac{(i_{\text{dither}})^j}{j!} \frac{d^j P}{di^j} \right) \left(\sum_{j=0}^{\infty} \frac{(i_{\text{dither}})^j}{j!} K^j \frac{d^j A}{d\omega^j} \right). \quad (10.7)$$

The double sum is approximate because the laser power and tuning are not uniquely defined functions of I_{bias} , but depend on temperature and history as well.

Detecting the second harmonic gives terms proportional to $d^2 A/d\omega^2$, $d^2 P/dI^2$ (the second derivative of laser power vs. bias current), and the cross term of their first derivatives. Since the laser power is reasonably linear in I_{bias} above threshold, its second derivative is small, so its additive noise is strongly suppressed. Differentiating sharp spectral features makes them grow, so the cross term is suppressed as well.

The lineshapes you get from this sort of measurement aren't what you'd measure in an ordinary spectrometer, which is a drawback. For measurements of the concentrations of known species, it works pretty well, and the optical apparatus is about as simple as it gets: a collimated diode laser, sample cell, and detector.

FM spectroscopy is an AC measurement done by modulating the beam, but it is very different from naïve chopping. In a chopping measurement, the beam switches from full on to full off regardless, but in the FM spectroscopy case there's almost no second derivative signal when there's no sample absorption peak.

10.9.4 Modulation Mixing

The FM spectroscopy example leads to another possibility, namely two-tone spectroscopy. Modulating the laser with two frequencies at once will cause mixing products that depend only on the cross terms in an expanded version of (10.7), and so are more nearly zero background. A high-frequency modulation and a lower-frequency one can do measurements far from the baseband noise and scan artifacts, with a nearly zero background characteristic. The downside of this is poorer shot-noise-limited SNR, because we're starting with two relatively weak sidebands rather than the main beam, and of course lots of spurious signals due to all the possible sidebands that can mix together.

The Doppler-free two-photon spectroscopy experiment of Example 1.1 is another example; the sum and difference frequencies come only from molecules that have absorbed one photon from each beam, and hence exhibit no Doppler broadening. Another one is the attractive mode AFM of Example 10.3, where the two modulations are the cantilever vibrations and the heterodyne interferometer's carrier frequency, and the resulting mixing appears as PM sidebands. Photoacoustic measurements with optical readout can be done this way too. There's more theoretical detail on modulation mixing in Section 13.10.6.

10.9.5 AC interference

One good way to make good fringes different from bad ones is to make them move in a rapid and predictable fashion, by using two beams of slightly different frequency. The frequency difference can be produced with an acousto-optic (Bragg) cell, a two-frequency laser such as a Zeeman-split HeNe, or by rapid frequency modulation of a diode laser, together with a path delay in one arm of the interferometer. Slower fringe motion can be obtained using mechanical means, such as a corner reflector spinning about its three-fold symmetry axis.

There is a trade-off of optical versus electronic complexity. Some people like to use multiple Bragg cells, in order to get a fixed IF directly from the photodiode. This is OK if you have a drawer full of Bragg cells and are just using it in the lab, but in an instrument it wastes photons, costs a lot, and isn't really needed by anybody who can build a radio. It's generally better to use more analog electronic signal processing and less optics; this preserves dynamic range while reducing complexity and cost.

10.9.6 Labeling Modulation Phase

If your measurement physics has an intrinsically slow response time (e.g., thermal effects), you can phase shift the signal with respect to the background by chopping the

excitation light. Choose a modulation frequency where the real signal is phase-shifted by $20\text{--}45^\circ$ from the chopped background (above 45° the amplitude falloff starts hurting your SNR). Null the background by tweaking the lock-in's phase. A nice stable setup will give you 40 dB (electrical) background rejection this way.

10.9.7 Labeling Arrival Time

There are other time-sensitive measurements besides OTDR. For example, you can make good range measurements by putting a linear current ramp on a diode laser in a Michelson interferometer and looking at the frequency of the output signal. If the tuning rate is $\dot{\omega}$ and the round-trip delay is Δt , the output frequency is $\dot{\omega} \Delta t$. You can do a good job with this, if you can make your diode stable enough.

10.9.8 Labeling Time Dependence

Pulsed lasers are rotten spectroscopic sources, because of their pulse-to-pulse variation in power, which changes irregularly over time. One way of dealing with it is the *cavity ring-down* approach, a kind of optical ju-jitsu; let the pulse have whatever energy it wants, and watch it decay as it rattles around in a resonator containing the sample. Regardless of what the pulse energy is, the attenuation per pass gives the sample absorption.

You use a laser whose pulse width is shorter than the cavity round-trip time, so that the pulses don't interfere with themselves, and look at the leakage out the end mirrors. A logarithmic detector (usually a DLVA) will produce an output whose slope is proportional to the fractional loss on each pass.

10.9.9 Labeling Wavelength

Spectroscopic measurements are of course labeled with wavelength, but as a background rejection device we need something like a two-wavelength colorimetric chemical measurement. There are a number of reagents whose color changes in the presence of some other chemical, for example, copper sulfate and ammonium hydroxide. Measuring the absorption of the mixture at two wavelengths can normalize out minor concentration errors and source brightness, leaving the ratio of the absorbances at the two wavelengths.

10.9.10 Labeling Coherence

Light whose coherence function is narrow (e.g., LED light or thermal light) can be used in interferometers for absolute distance measurements, because the fringe visibility goes through such a big peak near zero path difference. A sinusoidally modulated diode laser can do the same thing; its fringe visibility goes to 0 at the nulls of the Bessel function for its phase modulation (see Section 13.3). The modulation index m is the peak *relative* phase deviation in radians, which depends on the transit time as well as the FM spectrum itself.

10.9.11 Labeling Coincidence

If you're down in the mud scrabbling for photons, and the dark current of your photomultiplier is a serious problem, coincidence detection is a lifesaver. We usually think of it

in nuclear measurements, e.g. two-gamma transitions in nuclei, or scintillation counting. However, it is equally applicable to other measurements, such as surface scatter from a particle adhering to a silicon wafer. Multiple detectors looking at a large solid angle are one way to get all the photons, and if you run a coincidence or majority-rules scheme to weed out noise pulses, you can avoid the limitations of dark pulses in PMTs. (If you're using Geiger-mode APDs, remember they emit light when they avalanche—see Section 3.6.4.)

10.9.12 Labeling Position

Even in a complicated, multiple-scattering sort of sample, it is possible to restrict the measurement to a small region. While this may waste photons, those weren't going to help anyway. An example is an ordinary bright-field microscope with a translucent sample, where out-of-focus planes scatter light into the detector that makes the image fuzzy and ambiguous. Good baffles and spatial filters are a good way to do this, for example, the confocal microscope, where light has to make it back through a pinhole before reaching the detector; the light that would have blurred the image misses the pinhole.

10.9.13 Labeling Polarization

Sometimes we can manoeuvre the background signal into a polarization we can do without, and then get rid of it with an analyzer. If you have a sample that you can modulate, e.g. a magneto-optic film that responds strongly to an applied AC field, this allows the same sort of electronic zero background measurement as sitting on the lock-in's null.

Example 10.7: Patterned Thin Film Thickness Measurements. Let's say you have a dielectric stack with patterned metal at the bottom—it could be a silicon wafer or a magnetic recording head—and you're doing a broad-area thickness measurement based on interference between layers as a function of λ . The patterned metal doesn't obey the nice plane-parallel film assumptions we usually make in film thickness measurements, so the a priori approach won't work. If you care about the interference between the wafer surface and the top of the metal, but not about the thickness of the dielectric, you can use p polarization at Brewster's angle, which will pretty well get rid of it.

It's more interesting if you want the dielectric thickness: in that case, you might try flopping between s and p with a liquid crystal variable retarder. The fringes in the p -polarized light will contain mainly the metal-to-substrate fringes, and the s will have all three pairs. By knowing the metal-substrate distance, sorting out the dielectric thickness from the three sets of fringes isn't too hard.

10.10 CLOSURE

There is a broad class of measurements based on *closure*, first used as a way of correcting for atmospheric phase shifts in very long baseline interferometry (VLBI) astronomical measurements. The basic idea is that if you have N sources, you can get $N(N - 1)/2$ independent pairwise comparisons, which gives an overdetermined system when $N > 3$. Overdetermined systems can be solved with least squares by using singular value decomposition (see Section 17.9).

If you need to measure the short-term stability of a laser, the classical method is to build two sources, lock them loosely together, and look at the beat note between them. That gives you a signal you can measure on a spectrum analyzer, and it's reasonable to expect their short-term instabilities to be uncorrelated, provided you account for microphonics, temperature gradients, power supply variations, and stray pickup. If you're more concerned with drift, temperature stability, and so forth, you can lock one laser to an atomic transition (e.g., an iodine line) and beat it against your unknown.

The closure approach would be to build three lasers and look at the pairwise correlations. From measurements of pairwise phase deviations $\Delta\phi_{12}(t)$, $\Delta\phi_{23}(t)$, and $\Delta\phi_{31}(t)$, you can solve for $\Delta\phi_1 - \Delta\phi_3$, the individual deviations from the mean phase. (You can't get the mean phase itself—you're never going to be able to measure the carrier frequency from the relative phase fluctuations, for instance.) The nice thing about this is that you can assign an individual phase deviation to each source. In the laser case, that means you don't have to assume that the lasers have the same noise, as you do in the two-channel measurement. Similarly, if the measurements themselves have noise (which they will), adding more sources and using SVD will improve the estimates further. The two-channel correlation noise measurement of Section 17.11.6 uses a similar idea to get $N(N-1)/2$ separate measurements from N voltmeters, allowing noise measurements $\sim 20 \log N - 6$ dB below the noise floor; with closure, you can measure signal as well as noise.

Designing Electro-Optical Systems

Anything worth doing is worth doing well.

—English proverb

Anything not worth doing is not worth doing well.

—Anonymous

Anything worth doing is worth doing badly.

—G. K. Chesterton[†].

11.1 INTRODUCTION

Designing an instrument isn't just figuring out what should go where, although that's part of it. An even more important design task is to get the concept and motivation right—to think about the work in a way that maximizes your chances of success. This chapter starts off with the thinking, and then deals with the nuts and bolts. (You should approach it the same way.) By the end, you should have a framework for understanding how your instrument fits with your style, goals, and resources, and an orientation to how to proceed with the detailed design. Systems are so diverse that striving for complete generality would make it impossible to say anything concrete; thus we'll work through some examples instead.

11.2 DO YOU REALLY WANT TO DO THIS?

Before starting, do understand that every original or groundbreaking development project will nearly fail, and most of them nearly fail several times. This can be profoundly discouraging if you're not prepared for it; nonetheless, if the measurement physics is right and the technical risks have been dealt with appropriately, the project will usually

[†]“What's wrong with the world,” P + 4, Ch. 14. Elsewhere we says, “There are some things we expect a man to do for himself, even if he does them badly.” *Orthodoxy*, <http://www.gutenberg.org/files/16769/16769.txt>.

TABLE 11.1. Some Common Ways Projects Fail

Design Methodology Errors	Designing the Wrong Thing
Not making a photon budget	Pyramid building
Ignoring constraints (e.g., cost, size, use)	Creeping featurism
Runaway schedule and budget optimism	Runaway complexity
Not reducing technical risk fast enough	
Nobody having a clear picture of the whole	
Execution Failures	Moral Hazards
Not building what you designed	Fickle customers
Not verifying the design as you go	Fickle managers
System integration failure	Ignoring problems
Running out of time	Not confessing when in trouble
Running out of money	Not keeping everyone honest
	Loss of enthusiasm
	Failure of nerve

succeed. Overall, the success rate for well-conceived projects seems to be around 80%, and the failures are usually due to discouragement or delays (Table 11.1 has other ways to fail, for ready reference). You will spend time staring at the ceiling at 1 a.m., so make sure you have some answers to those wee-hours questions: “Why am I doing this?” “How sure am I that the physics is right?” “What am I missing?” “Does the prototype really match the design?” And especially, “Why doesn’t it work?”

11.2.1 Collegiality

These questions are much harder to answer if you are working alone. The principal reason is that you’ve been over the ground many, many times and have cut a rut so deep that you can’t see over the edge to where the new ideas lie. A secondary one is that personal and private failure is staring you in the face. Discouragement and fear sap creativity and enthusiasm faster than anything else, leading to little progress, leading to more discouragement and fear, until someone puts you out of your misery. This danger looks trivial to anyone who hasn’t faced it. The wisdom here is: don’t attempt a really difficult instrument project alone. That doesn’t mean you can’t be an engineering team of one, but it does mean that you need colleagues around you who are generous with their time and are willing to work to help you succeed, not just make encouraging noises. The author once had to take up an entire month of a friend’s time getting a complicated, ultrasensitive interferometric sensor to work properly, nailing down a decibel here, a partial vignetting there, and a half-wave of coma in another place. We worked elbow to elbow all month long, and although he didn’t know the system all that well at the beginning, that project would have been doomed without him.[†]

Once again: collegiality is an absolute psychological necessity for almost everyone, and the lack of it leads to severe stress and loss of enthusiasm even for good projects. If you don’t have it, don’t embark on a long and difficult project—stick to stuff with

[†]The colleague was Dr. Marc Taubenblatt, and the system was the ISICL sensor of Example 1.12.

shorter term payoffs, either data you can put on your wall, things you can sell, or a talk you can give.

11.2.2 Collegiality and Team Productivity

A person who is encouraged by the admiration and interest of colleagues and management will come in on weekends just to see how the experimental run is going, and will think about his work whenever he's not thinking about something else; his whole effort will be bent to it. It is easy to underestimate this, but the author's experience is that his productivity can easily change by a factor of 3 depending on collegiality alone.

Remember this: design teams are a bit like roses. Good ones can produce amazing results when they're properly tended—pruned, sure, but also watered, cared for, and given a place in the sunlight. Neglected ones get bedraggled very quickly. If you have people working for you, remember that they're the rosebushes and you're the gardener. Your glory comes from theirs.

Your enthusiasm and confidence are your most precious resources.

11.2.3 Choosing Projects

The 6 stages of a project:

1. Enthusiasm,
2. Disillusionment,
3. Panic,
4. Search For The Guilty,
5. Punishment Of The Innocent,
6. Praise And Honours For The Non-Participants.

—Anonymous

Is this project a smart risk? Will I be motivated by enthusiasm or by fear? Fear is a bit like amphetamines—you can stay up all night working, all right, but your integrated output will be lower, and your enjoyment will ultimately vanish. What you should look for in a project is fun, good leadership, enthusiasm, and the feeling that what you're doing is not only valuable but *seen* to be valuable, seen to be a smart risk—that is, if it fails due to bad luck, you won't suffer too much. Having management commitment that will survive failure is both rare and precious. Think about this soberly before starting a project.

11.2.4 Procedural Advice

Take Play Seriously. It is commonly said that people do their most original work before the age of 30, and this is often so. It's commonly held that the reason for this is rapid aging, but this is not so; like composers, the best instrument designers continue to improve through late middle age. It's just that most over-30 people have had their creativity flogged out of them, which is a huge waste. There's not much you and I can do about this except refuse flatly to have it happen in our vicinity, but it's worth at least naming the antidote: *play*. Play is not a waste of time that should have been productively used; it's where all the creative ideas come from, and a world of furious crank-turning

can't replace it. The seed corn sitting in the barn appears wastefully unproductive, but without it, we won't eat next year. The author learned what he knows about instrument building by playing, not by being strapped to a schedule like a silent-movie heroine on her log heading for the buzz-saw. Working in the real world, we have to make real-world moves, but driving ourselves so hard that we lose our sense of fun is a disaster. Overworked engineers are the ones who burn out and get obsolete. Students have time to play, and that's why they learn fast and dream large—it isn't youth so much as *time*. Spend at least one day a week on play and head maintenance—reading journals, messing around on side projects, going for walks to think about measurement ideas, inventing things on the white board with a colleague. You will get more done and not less, and you won't burn out.

Take Counsel of the Devil. One very important step is to spend time in your customer's and (especially) your competitor's shoes. Most engineering groups have a certain hothouse mutual admiration, which is healthy and valuable, but leads us to rate the competitor's ingenuity too low and the customer's discernment too high. Have a Blue and Red team match, where you pretend to be your competitor, and try to find defects in your own design or ways to avoid your patents. This can be bruising, but it's a lot more pleasant than having your competitor do it for real. You learn a lot, and besides, you wind up with better patents and better products.

For a successful technology, reality must take precedence over public relations, for nature cannot be fooled.

—Richard P. Feynman[†]

Resist Overpromising. Managers in instrument companies know that schedules are usually too optimistic and have some idea of how difficult a given project actually is. On the other hand, if you're building custom instruments for some other type of organization (e.g., a large computer company), your customer is probably a lot less technical than you are and is almost certainly not an instrument designer. His expectations depend not only on what you tell him, but on what he infers from your words and your manner, as well as what he wants to believe. You have to manage his perceptions of how long it will take and what you'll deliver. Of course, your estimates of time and difficulty will be as realistic as you can make them, but nevertheless, the temptation to sound too enthusiastic is almost irresistible; if you didn't think your gizmo was wonderful you wouldn't be building it. You may well create the fatal impression in your customer's mind that the work is practically done when you get your first data plot.

This can lead to a paradox—you succeeded brilliantly, but completely disappointed your customer, because (a) he thought it was easy, and (b) you didn't leave enough slack, or look as though you were struggling hard enough. Make sure you keep him up-to-date on all your technical work, and what you think of your rate of progress.

Keep Some System Margin in Your Back Pocket. If you're building a product, you're managing several things all at once: the project, your career, and the expectations of your customers. Your design spec, be it a photon budget, a link budget, or a resolution

[†]Richard P. Feynman, Appendix F—Personal observations on the reliability of the Shuttle. *Report of the Presidential Commission on the Space Shuttle Challenger Accident*.

specification, is your best shot at what you ought to be able to do. That means that any surprises you encounter along the way are going to reduce performance, not increase it, which is quite natural—we aim at the highest possible point, and never quite reach it. People not directly involved in the project will normally interpret the system design as being the guaranteed specification, and will interpret any performance reduction as failure—*your* failure. This is extremely painful. Thus it is very important to keep two sets of books, and don't show the customer the real one. This is not dishonest, as in money laundering and tax evasion, but is a recognition that you and your customer have different understandings of how a system design functions, *and that yours is right and his is wrong*. He will insist on regarding it as a firm promise, whereas to you, it's an image of perfection. Hide at least 3 dB worth of buried treasure in your public version, so that when the inevitable snags occur, you have some margin to make up for it. If there's any left when you're done, use it to make yourself look like a hero. You will be.

Have People to Cover Your Back. The near-certainty of near-failure means that we all need defenders. In the bad patches, there will be people who stand back and throw rocks, saying that the project is doomed and should be killed off. Some organizations have a worse case of this than others, but the remedy is the same: don't poke the alligators, and cultivate allies by spreading the credit around, taking advice, helping out, eating lunch together, and letting them have some ownership of the system—that is, tell everybody how much the project owes to everybody, and *make it true*. If it's theirs too, they will dig in and defend it during the bad times. The most important people in this category are your management and your customer.

Confess When Your Project Is on the Skids. If you expect people to defend you when you're in difficulty, it is only fair that you tell them when it happens. Follow the airplane rule: "When lost, climb and confess." Expect the same from your colleagues, and be willing to press a bit; since designers are normally very eager to discuss their latest idea, a colleague who is evasive is very likely to be hiding a failure.

Define What Constitutes Success. Having thought about the consequences of failure, spend a little time on what is likely to happen if you succeed. It's important to think out what you want to happen then, and to reach agreement with everyone involved as to what success looks like. For a student, success usually means getting enough data to write a couple of papers and graduate, which requires a lot less engineering than producing a design to be used by thousands of people.

If you're a product designer, you might intend to go back to the lab for the next brilliant scheme, but wind up wedded to this one for your whole career. Being lead designer of a really new product in a small company will oftentimes lead to this. On the other hand, a successful development project in an area that is not your organization's core business is liable to be sold off, with you attached. If you're lucky, being sold off may make you rich, but usually it's more like being sold down the river—the job and the salary stay the same, and the pressure and insecurity get much worse. Is going that way okay with you? If not, how will you disengage, to go on to your next project? You need to be fast on your feet; leave the project about six months before the sale to avoid being caught in the undertow.

If the project is likely to remain with the organization, how will the system be manufactured? Sometimes the designer moves to manufacturing to become product manager

for the system, and may not get to design another one for some time (if ever). Other times, the design is transferred to the manufacturing department, which is responsible for replicating it. This is a famous source of friction and finger pointing: the production yield is too low—whose fault is it?

Build bridges to the manufacturing people early, and make sure they have some ownership by bringing them into the design process. Explaining how it works and what it's for may seem elementary and obvious, but they probably have no idea, and most of us like to be treated like colleagues instead of cogs in the mechanism. You'll get a lot less grief from people who feel part-ownership of the system, and the transfer will work better, too. This of course presupposes that you know who will be manufacturing it.

Know Your Organization. Large companies tend to churn managers, so the expected life of your project may extend across two or even three different managers at all levels. This is because reorganizing gives the illusion of progress and is so much easier than attacking the actual problems. A new manager often feels pressure to make his mark quickly—so like a lion taking over a pride, the first thing he's liable to do is eat all the cubs. You're pretty safe if you're just incrementing an existing product, but watch out if you're doing something the new guy doesn't understand and isn't excited about. In that case, he doesn't care about success but risks getting blamed for failure, so you see the temptation. This is one big reason that large companies have trouble producing really new things over long periods; the exceptions, such as 3M, Corning, and IBM, have cultures that encourage patience.

To keep your project alive, make sure you give lots of sales pitches to managers, even ones slightly outside your chain of command, to build their enthusiasm and give them something to use to convince the new guy to let you do what you're doing. Make sure that when you do get a new manager, you make time to sit down with him to explain why what you're doing is exciting and good for the organization—managers, like most other people, want to do the right thing if it isn't too risky, so concentrate on that. Don't get so concerned with management-tending that you lose your balance, though. Management jobs attract people who are comfortable scheduling their time to the nearest nanosecond, and who equate busyness with productivity—an attitude that leads to a clean desk and zero creativity, which won't help you. Sometimes you just have to take your lumps and move on.

Don't Build a Pyramid. Everyone seems to build one pyramid per career. A pyramid is an ambitious system that one person really cares about and that winds up working well, but then just sits in the desert because nobody else cares the same way. This happens usually just after leaving graduate school.

There are two kinds of pyramids: yours and your boss's. Sometimes you get an idea into your head and sell it hard. Other times, you fall for a sales pitch about how you're the key guy in the department, how this just *has* to get done, how it will result in showers of glory, that you're the only one who can possibly do it, how you of course can't let the group down, and so on.

It's exactly the same sales pitch that officers and NCOs give infantrymen before a battle. Being unlucky in the infantry is a lot worse than in engineering, but the soldier's odds are usually better. Don't fall for a romantic but hopelessly understaffed project, no matter how big a hotshot you are; and if it's your idea, make sure you think it out very carefully. Remember, if you're in love with it, own it, and argue for it, nobody will want to tell you it's a waste of time.

Understand the Pecking Order. Without instruments people, most of science would be impossible; a major fraction of what appears to be progress in basic science is really the progress of the instrument-building art. In our sublunary world, credit is seldom bestowed really accurately, and most of us are quite used to that. Besides, gizmo building is such fun that it's worth trading off quite a bit of glory for. There are lots of interesting problems in the world, though, so you might as well work where you're appreciated. In biomedical instruments, for instance, if you aren't an MD you might as well be part of the furniture. Astronomical instruments folk are also treated like spare telescope parts. Working for a small company where your expertise is similar to that of the August Founder is usually uncomfortable unless you're the Anointed Protégé, which has its own troubles. In general, status-conscious places are miserable for everyone, and the more, the worse.

Don't Fight "Good Enough." A common error, related to pyramid building, is to build a new alternative to an existing technology that is already good enough and is in wide use. A large-scale recent example is magneto-optical storage. It was introduced in competition with magnetic storage, pushed hard for 10 years, and lost (at least for fixed-disc applications), because its then-great advantage in storage density was not good enough to overcome its bulky, massive read/write head designs, which forced the platters to be stacked too far apart and slowed down the track-to-track seek times. Magnetic storage didn't stand still, either; better media, the magnetoresistive (MR, and then giant-MR) read head, extremely small head-disc gaps (a few hundred angstroms), and improvements in servo design allowed track widths to shrink and linear recording densities to increase enormously in that 10 years. (Magnetic is now far denser.)

A new technology in a busy field needs a sheltered niche in which to grow and mature, and (what is usually forgotten) for users to become familiar with it. Mainstream technologies have huge development efforts behind them, consisting of people who have years of experience with the technology and who are backed with lots of money. Even if your idea is 10 times better, you're going to get steamrolled unless your development effort is comparable in scale.

Agree on Specifications Before Starting. Specifications are often written by people who have no idea what they're doing, but think that they have. The first task on the list should always be "specification resolution," with a customer sign-off required. That way, when they discover that the specification isn't exactly what they want, the responsibility (and financial liability) clearly lies with the customer. This is especially important when you're an independent contractor.

11.3 VERY BASIC MARKETING

11.3.1 Who or What Is Your Customer?

Before picking up a single lens, make sure you know who your customer is. For most of us, there is more than one: first, the managers, bean counters, and contract monitors we have to convince, and second, the real end user of the system.

For each class of customer, we need to know the answer to a few rude questions (it's possible to fish for the answers politely, fortunately). What do they want? Do they know what they need? Experience suggests that they often don't, but always think they do;

they are the ones with the checkbooks, though, and so can't just be ignored even if they are provably wrong. Unfortunately, they often assert their power by demanding long lists of "checklist features," which they'll never use but which your competition has. Will your market support the amount of engineering required?

11.3.2 Making a Business Case: Internal

A grubby necessity of the instrument-building trade is having to convince people with money that your latest brainstorm deserves to be made into a product, will save the company a big chunk of money, or (a much harder sell) has enormous prestige value. All trades are like that at some level, but in other businesses you're usually dealing with somebody who has some understanding of what you're doing. Unless your business manager is a technical person of some talent, he probably won't know how your instrument works, and so will have only your say-so (and perhaps someone else's evaluation) that it does. This is naturally anxiety producing for the manager, so you'll need to be prepared with a truly compelling business case. It must at least contain the following elements: a detailed design concept (what are you building?), marketing information (who's the customer, what's it worth to him, how much of that can you recover, and over what time span?), a fairly detailed plan for how to do it, including estimates of time and resources required, a levelheaded assessment of the technical and business risks (including anticipated work-arounds and how you'll go about reducing the risk quickly), and an estimate of the opportunity cost (i.e., the value of what you would otherwise do). Assuming you get taken seriously, you'll need to be able to answer questions on all these things. Doing a bit of thinking and research about how the market and the competition is likely to develop between now and when you're done will help too.

The development plan is especially fraught, because you'll often wind up being forced to invent on a schedule—a notoriously difficult task. It's worth having a couple of good ideas tested and confirmed before you make the business case, which will require stealing a bit of time from your current project to work on them. Don't hesitate to do this, because it makes you much more valuable to your organization as well as to yourself. It also helps you to avoid overpromising (see above).

Consider keeping a couple of good unannounced results (including pretty pictures) in your desk drawer, ready to pull them out on some rainy day to encourage yourself and others. There are lots of times in development projects (especially ones with lots of electronics and software) where you're nearly finished before you have much to demonstrate. It's good to have something new to show when you make yet another sales pitch.

Projects usually run in two stages: a feasibility study followed—if all goes well—by a full scale development effort. There's a sales pitch before each stage, for which you'll need some solid support: a few plots for the first sales pitch, and a working demo for the second.

11.3.3 Making a Business Case: External

Your customer often won't know what he needs, but he'll probably know what his problem is, and how much it's costing him. Thus the easiest kind of case to make goes something like this: "Here's a pile of damaged items P , which cost Y dollars to make, and could have been sold for Z dollars if they hadn't been damaged by Q (contamination,

frost, broken bottles. . .). This happens about V times per year, and could have been avoided by using my instrument, which costs W dollars installed. Thus the instrument will pay for itself in $t = W/(VZ)$ years (where $t < 2$.)” This goes down very well, provided that you really can point to the pile of damaged goods, and that you’re pitching to the person who owns the pile.

It is important in making this case to ask yourself what else he could do, if he didn’t use your instrument. If something simpler can do nearly as good a job, it might make sense to do that. In a bottling plant, an audio system listening for the sound of breaking glass (lots of high frequencies) might do the same job as a machine vision system, at a $10\times$ lower cost. You’d have to make your case on the basis of quality control, which is much harder (though not impossible).

The “What else would he do?” question also means that the payback period has to be shorter for instruments than for most other investments, because there’s an opportunity cost to not waiting for a possible cheaper solution. That 2 year number is about right—much longer than that and the sale gets much harder; much shorter, and you’re not getting enough of the benefit, and should charge more money.

11.3.4 Figuring the Price

The usual rule of thumb is that the selling price of an instrument is about three times the parts cost, and more if it’s going to need a lot of after-sales support. If this sounds high to you, it isn’t really: the cost of your gizmos includes not only the parts and the facilities needed, but the people who build them, the people who fix them, the people who look after the payroll, and (last but not least) the people in Italian shoes, eating expensive lunches, who persuade people to buy them. If you can’t sell it for $3\times$ the parts cost, redesign it to use less hardware; be ready to shed the less important features. (Note that this is the real cost of the parts when you buy production volumes, not the onesie-twosie price.)

11.3.5 Budget for Market Creation

That factor of 3 is for an instrument that people already know how to use. A really new instrument often requires educating customers about how it can help them. This market-creation activity is slow and sometimes painful. Don’t underestimate how difficult it is; if you are aware that it’ll take a while, and don’t disguise the fact, you’ll have a much easier time of it when no tidal wave of orders arrives.

One piece of marketing wisdom: kind words are much easier to get than purchase orders. If your prospective customers are saying they’d like to buy your gizmo, make sure you know exactly what practical problem it will help them solve, and whether there’s a genuine business case to be made for it. They’ll have to make that case before the PO gets generated.

11.3.6 Budget for After-Sales Support

The education task doesn’t end when the sale is made. If you’re building a new class of instrument, one where the customer probably hasn’t used one before, you’re going to get a lot of phone calls afterwards. You’ll have to pay people to answer them—good people, not phone fodder—so factor that into the price. If you’re in the in situ instruments

business, the cost of sale and after-sales support goes way up because of the complexity of getting your instrument into your customer's system.

11.4 CLASSES OF MEASUREMENT

What exactly do you want to measure? Do you have a clear idea of your measurement physics, or are you burrowing around in spectral, video, or time series data to find something that correlates with what you're trying to measure? Or just looking for pretty pictures? This is a somewhat rude question to ask, but the answer really determines what your design strategy should be.

11.4.1 Know Your Measurement Physics

In deciding on the feasibility of a measurement, there must be a balance between theory and experiment. Ideally, theory shows that the measurement is possible, and how big the signal should be; experiment shows whether the theory is right (often leading to improved theories) and that the apparatus can be built. A combination of theory and data, demonstrating good agreement, shows that the measurement is at least possible; the actual instrument design is less important, because there are usually a few different good ways to go about it. How well your measurement idea lives up to that ideal is the principal determinant of your strategy.

11.4.2 Crunchy Measurements

Measurements come in three textures: crunchy, squishy, and in between. Crunchy measurements are those whose measurement physics is well understood, and where the parameter being measured is identical with what you want to know, e.g., optical pyrometry for the temperature of a closed furnace, where you're measuring cavity radiation, whose properties are known in advance, and where calibration on one system is adequate for making measurements on another.

Crunchy measurements are always the best. The better you understand the measurement physics, the more robust you can make the measurement with a given amount of testing.

11.4.3 In-Between Measurements

There are lots of in-between measurements, where we have at least an arm-waving idea of why it works, and can logically connect the measurement data with what we want to know. An example of this is a fiber-bundle spectrometer looking at thin film interference fringes to measure film thickness in semiconductors, when the structure underlying the film is poorly controlled. These measurements are based on correlations between some feature of the data and the parameter you want to measure, and are as valid as your knowledge of the origin of the correlation: the better you know it, the better your measurement and the better off you are.

Making in-between measurements robust means a significant amount of experimental work in realistic conditions to establish whether they are really telling you what you need to know, and if so, how to extract the information you care about from the more

or less closely related experimental data. This must be followed up by looking at a whole bunch of pathological special cases, and putting in tweaks and hacks to handle them. Measurements without much theory tend to break the first time an unanticipated situation arrives, so you have to do a lot of anticipating. An example is downstream particle counting, where an optical sensor in the roughing line[†] of a reactive-ion etch tool for semiconductors (say) looks for particles in the turbo pump exhaust. You care about the particles on the silicon wafer, not down the pump, but there may be a correlation you can use, since we expect a dirty chamber to have dirtier exhaust.

11.4.4 Squishy Measurements

Squishy measurements are based on a correlation between something you can measure and something you care about, where you don't really know why the correlation exists, for example, multispectral reflectance measurements of agricultural produce to indicate freshness. This is a tough business, for two reasons: first, the correlation is often unstable with time, and second, there is an enormous amount of *nuisance data*, that is, complicated changes in the background signal, which you have to be sure of rejecting (see *ground truth* in Section 11.7). Fresh tomatoes look pretty different from fresh lettuce, and Iceberg looks different from Romaine; wilting due to drying out looks different from mold, which looks different from wilting due to frost.[‡]

A sufficiently squishy measurement cannot be made reliable by any amount of testing and correlation whatsoever: as Deming said, “you can't test quality into a product,” even if your product is a measurement technique.

Keep your measurements crunchy if you can, and if you can't, expect a lot of 3 a.m. phone calls. The moral of the story is: *good data and good theory matter more than easy data gathering*.

11.4.5 Pretty Pictures

People are busy, and most of us are intellectually lazy at least part of the time, including grant committees, managers, and thesis advisors. Almost every designer has to sell his instrument to someone at some stage. Pretty pictures help a lot—a measurement whose data generates no pretty pictures will get less attention and be a harder sell than a more pedestrian one with dazzling output. If you share the popular disdain for pretty pictures measurements, think again.

This far into this book, no one should mistake the meaning of this: it's the data that matter, the pictures are for the sales pitch.[§] The pitch is important, even for a graduate student; if your thesis talk has data whose significance leaps off the screen, your defense will be a lot easier, as will your job search afterwards. For others, you know that the opportunity to build instruments for a living makes frequent sales pitches (to customers and your own management) a necessity. Have pretty pictures whenever possible, and use some ingenuity on them.

[†]A high vacuum system typically has a *roughing* pump (exhausting to atmosphere) to get the chamber down to the millitorr range, and then a high vacuum pump such as a diffusion or turbo pump for lower pressures. The two are arranged in series with the roughing line in between.

[‡]This isn't a silly example: the author was once approached with a serious proposal to use multispectral imaging of this sort for quality control in fast-food hamburger restaurants.

[§]Sometimes the pictures help you spot patterns in the data, but those aren't just pretty pictures.

11.4.6 Pretty Pictures Measurements

As all technical people know, the credibility of measurement data goes down as the number of colors on the chart goes up. Real data should look as different from marketing presentations as possible. Nevertheless, there is a legitimate class of measurements whose output is exclusively pretty pictures.

Pretty pictures measurements include using UV imaging to reveal counterfeit documents, proof-of-concept experiments to show that your measurement idea can actually work (e.g., the first scanning tunneling microscope picture), or writing the IBM logo in novel ways in order to get free advertising on TV. If they use existing equipment, they can be thrown together quickly, and usually should be. Good physical measurements leave the pretty picture stage behind very rapidly.

In measurements whose frank aim is pretty pictures, it is appropriate to use somewhat stronger postprocessing than you would if you were after quantitative data. For example, grey scale pictures have limited dynamic range; in turning scanning measurements into a grey scale image, we often *plane subtract*: that is, subtract out a function $f = ax + by$ for some a, b so as to flatten out the background and avoid wasting dynamic range on artifacts. In a pure pretty pictures measurement, you might take out a quadratic term in the background as well, to make it come out a visually pleasing uniform dark grey, or use false color to make it stand out better.

11.5 TECHNICAL TASTE

Genius is 1% inspiration and 99% perspiration.

—Thomas A. Edison

Edison correctly identified the *inspiration/perspiration ratio* (IPR) as a key metric of project quality. The fact that his quoted ratio is only 1% shows that he was doing it completely wrong, of course.[†] Technical taste is a feeling for the right way to do things that develops with experience and with being steeped in a really competent and imaginative technical culture. Good taste helps you maximize the IPR; that is, the idea is to use neat ideas to minimize grunt work. Keep the IPR as high as possible, by all means, but do make sure that the neat ideas are neat because they work really well, not because they're unusual or just cool.

If you've got a tingling feeling in the back of your mouth, suggesting that there has to be an easier way, *pay attention to it*. It is one of the few indications you will get that your latest gizmo still has a lot of bric-a-brac that can be removed. There is no substitute for the process, but a few representative points can be made.

Know What Your System Should Look Like. The ideal instrument for fixed use is the size of a paperback book, with an Ethernet jack and an on-board web server for control and display. For hand-held use, it should be the size of a deck of

[†]Nikola Tesla worked for Edison for a while and was amused and frustrated at his boss's insistence on doing many, many trials without theory to guide them. Edison had excellent intuition, was a marvel of energy, worked in relatively simple technology, and eventually had a whole lot of lab assistants. Not many of us are in that position.

cards and communicate via USB or possibly an IR link. Some instruments need to be bigger, but large size costs money, and many customers won't want some big feature-less rack cluttering up the place. Decent quality packaging and cabling are obscenely expensive.

Example 11.1: Cost-Reducing an Angular Displacement Sensor. A device for sensing the position of an object in 3D originally consisted of several IR LEDs mounted close together, illuminating the object, with a lens plus a two-axis lateral effect cell to sense its angular position, as shown in Figure 11.1 at the top. The total photocurrent gave its radial position. The photocurrent was dominated by ambient light, and the cost of the sensor was dominated by the lateral effect cell (\$100). A quad cell wouldn't work since the angular sensitivity would depend strongly on the state of focus.

Some days after the tingling feeling started, the idea of a sundial came to mind, leading to the shadow mask detector design of Figure 3.3. Eventually (bottom of Figure 11.1), the system consisted of a small PC board, with three pairs of \$1 photodiodes on it, detected differentially. A slot cut in the lid of the plastic box provided the shadow mask, and a Velostat (carbon-loaded plastic foam used for shipping ICs) gasket compressed between the board and the lid prevented the source from illuminating the detectors. The signal level and reliability increased, and the manufacturing cost went down by six




Figure 11.1. Angular displacement sensors.

times (there's more on this in Example 3.3). (Neither system is limited by diffraction—why?)

Don't Measure Anything You Don't Care About. One sign of a good measurement idea is that you're only measuring what you care about, and don't need to correct for a whole lot of large extraneous effects (keeping track of the temperature is okay). Most eye trackers used with external displays (i.e., not helmet-mounted) measure the angle of gaze with respect to the head, using a glint of light off the cornea. In order to determine where on the screen you're looking, they have to measure the position and orientation of your head in six axes, to high accuracy, as well as the gaze angle. This makes things difficult and expensive. Having to measure eight things when you care about only two is a losing proposition.

Avoid Obese Postprocessing. The trade-off between hardware and software depends on the task and the relative costs, but also on the expertise and inclinations of the people involved. In our software-heavy era, some people feel a strong temptation to try to use lots of postprocessing to squeeze a good measurement out of an inadequate sensor, which is an invitation to failure. An example is trying to use a video camera on top of a computer monitor to do accurate eye tracking, by piling on frame averagers, median filters, feature extractors, pupil-centroid-finders, saccadic-motion-blur-eliminators, and so on. Your measurement gets squishier and squishier, and so less and less reliable, for no good reason. As is repeated elsewhere, *postprocessing is not a substitute for good data.*

Have a Fallback Position. In building electro-optical instruments of any complexity, it's important to have a backup plan in case each new item fails to work. A typical project will get in trouble between two and five times before eventually succeeding.

It isn't necessary to have a recovery plan if all of your new bits fail to work; we assume that you've reduced the technical risk before going all out on the project in the first place, so that the scariest new technology is at least possible. You can't recover from total technical failure, but usually you can think of a work-around for each piece ahead of time. Pieces that don't admit work-arounds have to be nailed hard, fast. That's where potential show-stoppers lie.

Wear a Belt and Suspenders. For high performance applications, don't rely on just one layer of defense; you won't have thought of everything that can go wrong, so add as many useful extra safeguards and sanity checks as you can without overburdening the system with electro-optical bureaucracy.

The ISICL sensor of Example 1.12 uses 10 layers of protection against laser noise and time-varying speckles due to back wall and window reflections: a very low noise diode laser driver (to make the scattered light quiet), millikelvin temperature stabilization (to avoid mode hops), automatic laser current adjustment (ditto), coherent detection (to get the high gain and selectivity), a laser noise canceler (to get to the shot noise), highpass filtering (to reject the incoherent speckles and plasma fluctuations), a CFAR servo on each frequency band (which allows continuous checking of the false alarm statistics), and burst rejection in both hardware and software (to prevent buffer overrun if we hit

a glint). None of these is expensive or elaborate, and the result is a reliable, shot noise limited measurement in very difficult conditions. *Defense in depth need not be expensive, but it does need thought.*

Avoid Underengineered Complexity. You can get into a lot of trouble by assuming that if the prototype works the first time, that means it's easy to do again. Make sure you know *why* it works.

You may not even get that far. A common error, especially among professors and managers, who don't have to do it themselves, is that a measurement made up of well-understood and widely used elements may be impossible due to sheer complexity. Remember the stories of the old vacuum tube computers? Although they were made of the same radio tubes that brought music to everyone's kitchen table, there were so many of them that a tube would fail every few hours, requiring laborious reloading of programs, checking of results for corruption, and so on. Even when the components are reliable, getting them all to work simultaneously can be very hard indeed.

It is a mistake to build instruments that require more than two new hard things to work at once; it's like fighting a two-front war. An interferometric readout atomic force microscope is not too bad to build and get working, but a two-wavelength combination near-field optical microscope, AFM, lateral force microscope with tunneling capability is too complicated, unless you have a shipping product (that already does most of it) to build on.

Know When It's Time for a Clean Sheet of Paper. As comforting as it may be to have a product to build upon, at some point this bird in the hand becomes an albatross around your neck. Sometimes you just have to start over. Roughly speaking, look for this possibility whenever your legacy hardware is less than 75% used in the final design, or when it reduces your performance by a factor of 2.

Beware of Signal Processing Fads. There has been a lot of snake oil peddled in the signal processing world over the years. The latest and greatest technique will allegedly pull clean signals out of ugly backgrounds so effectively that all previous methods are immediately obsolete. Image processors will take blurred pictures of ugly people and turn them into sharp pictures of pretty people—you know the tale. Nowadays the medicine shows concentrate on techniques like genetic algorithms, neural networks, maximum likelihood estimators, and so on. Many of these techniques are genuinely useful, but none is revolutionary—they're just tools, like Crescent wrenches.

When you're fixing a car, you need a set of combination wrenches, a socket set, some pliers and screwdrivers, and maybe a meter or two. If you need a special tool, like a gear puller, you go out and buy it specially. Doing the work yourself is partly an excuse to collect a nice set of tools with the money you saved.

Signal processing is like that too. If you know how to build and use amplifiers, filters, mixers, and A/D converters, you can make a good analog signal processor. Put that together with some simple postprocessing, and you're most of the way there. If you need something special like a phase-locked loop or a Kalman (adaptive) filter, you can look up how to make one. Besides solving the problem, you've acquired a useful skill. Do things a bit differently when you have an opportunity—that's a good excuse for collecting signal processing tricks. Deep and broad expertise can be acquired this way, providing you don't treat the things you look up as mere recipes to be followed. Spend a little time figuring out what's going on in each one, and you will be richly rewarded.

As a rule of thumb, it takes half as long to become really familiar with a new technique as it does to debug one you half-understand.

Beware of Optical Fads. There are a lot more electronic and computer people than optical people, so there are more fads in the electronic and computer businesses. However, there are some optical ones to avoid as well. A few years back there was a big vogue for “nondiffracting beams,” which allegedly exhibited greatly increased depth of focus for a fixed wavelength and numerical aperture. The trick was to build up the beam from plane wave components that all had the same value of k_z , so that their relative phases would be constant along the axis of the beam, and so they would never go out of focus.

Such a beam is typically generated by sending a collimated beam through an *axicon*, a cone-shaped prism pointing along the beam axis, that refracts the beam into an annular cone. The pupil function of the beam is a ring δ -function, $\delta(\rho - a)$, where $\rho = (u^2 + v^2)^{1/2}$, and the Fourier transform of that is a Bessel function, $2\pi a J_0(2\pi ar/\lambda)$.

It’s perfectly true that the 3 dB intensity radius of such a beam is very narrow and is not a strong function of defocus; what was glossed over, however, is that only a small fraction of the beam energy was inside the 3 dB intensity radius, because J_0 falls off asymptotically as $r^{-1/2}$, so its energy density goes as r^{-1} and the total beam power is infinite—asymptotically, each individual diffraction ring contributes 28% *more* power than the central spot. In instrument terms, that means that even if we chop off the rings at some point (as we must), most of the signal comes from outside the central spot. The apparently improved resolution/depth of focus trade-off is illusory, just another instance of specsmanship.[†]

There are other examples: diode laser feedback measurements, bright-field intracavity measurements, and some types of fiber sensor.

11.6 INSTRUMENT DESIGN

Nothing is more dangerous than an idea, when it is the only one that you have.

—Émile Chartier[‡]

Okay, so we know what the aim of our measurement is, and what crunchiness class it’s in. Now what do we do?

Know the Problem. One of the most basic laws of engineering is that you can’t solve what you don’t adequately understand. Do a conceptual design from end to end, first, as we did in Chapter 1, and play around with it to find the best trade-offs. Once you have a rough cut at the system design, you can identify the hard parts.

For example, if you want to build an imaging system, you need to know the resolution requirements in advance; these will of course depend on other system parameters, such as camera to subject distance, field of view, and so on. Use rules of thumb (see Table 11.2) to spread the pain out equitably among subsystems and among designers. This will normally lead to the lowest cost system using a given technology mix. Be

[†]One can argue that lithographic applications can benefit, because with the high contrast of photoresist, you can adjust the exposure to make the central spot print but not the rings. Since the rings of J_0 alternate in amplitude, that means you’d need a phase-shift mask, and the author is still dubious.

[‡]Quoted in Miller and Friedman, *Photonics Rules of Thumb*. McGraw-Hill, New York, 1996, p.187.

TABLE 11.2. Reasonable Limits for Common Design Parameters

Field angle: $<30^\circ$ half-angle
Pixel rates: <20 Mpel for scanning systems
Bandwidth: $10 \text{ Hz} \leq \text{BW} \leq 100 \text{ MHz}$
Mirror flatness: $1/10$ wave @633 nm
Etolon fringes vs. bandwidth and spatial coherence: 1% p-p, slope 15%/GHz collimated
Polygon alignment: 15 arc seconds
Mechanical backlash: $5 \mu\text{m}$
Motor speeds: <3000 rpm best
Eye-safe lasers: 1 mW visible, $100 \mu\text{W}$ IR
Diffraction limit: $\lambda/4$ rms wavefront error \rightarrow 0.8 Strehl ratio
Diode laser feedback tolerance: 3×10^{-7} in power
Diode laser current tunability: 1 cm^{-1} @ 800 nm, 0.5 cm^{-1} @ 650 nm between mode jumps
Diode laser temperature tunability: 10 cm^{-1}
CW laser noise: 0.1% for gas lasers, 1% for N_2 , excimer, 50 dB above shot noise for small YAGs
Diode laser linewidth: 50 MHz for a single-frequency diode laser, 500 MHz for VCSEL
Laser pointing stability: 10 arc seconds
Diode laser lifetime: 50,000 hours at 25°C
Silicon PIN diode stability: $10^{-3}/^\circ\text{C}$ with windows, $10^{-5}/^\circ\text{C}$ windowless
PMT/APD gain stability: 1% without too much difficulty, 0.1% if you really work at it
Departures from Gaussian noise: outliers often begin near 4σ
Electronic component stability: $100 \text{ ppm}/^\circ\text{C}$
Highest reasonable impedance: $ Z < 60 \text{ k}\Omega/f(\text{MHz})$
Coefficient of thermal expansion: $10^{-5}/\text{K}$ metal and glass, $10^{-4}/\text{K}$ plastic
Temperature coefficients of index: $10^{-5}/\text{K}$ glass, $10^{-4}/\text{K}$ plastic
Thermoelectric cooler temperature drop: 40°C for 1 stage, 90°C for four stages
Small heat sinks: $5^\circ\text{C}/\text{W}$ without fan, $1^\circ\text{C}/\text{W}$ with fan
Tungsten bulb life vs. temperature: $\text{life} \propto \exp(10,500/T)$
Surface leakage vs. humidity: $10^{14} \Omega/\text{square}$ dry, $10^9 \Omega/\text{square}$ at 95% humidity
Inductor Q : <80

aware that a shift of emphasis may give an offbeat, but robust and cheap solution; for example, an imaging spectrophotometer system might use a lens with bad lateral color but good longitudinal color; the lateral color is a change in magnification with wavelength, which can be removed by using several spectral bands and interpolating between pixels to change the images to a common scale. The trade-off there would be more CPU power and possibly more pixels required versus better resolution at fixed focus and a cheaper lens.

Mess Around with the Tools. Leave lots of time at the beginning of the project for playing. Diddling around with the new laser, drawing Lissajous figures with the scanning system, and looking at the trees outside the window with your new IR viewer are all things that may seem unproductive, but are actually important for a couple of reasons: first, seeing what cool junk we have to work with builds enthusiasm, and second, messing about this way builds technical taste faster than any other activity except debugging. Whimsy is a highly practical thing in a designer.[†]

[†]See Jim Williams, The zoo circuit, in Jim Willams, ed., *Analog Circuit Design: Art, Science, and Personalities*. Butterworth-Heinemann, Woburn, MA, 1991.

Understand the Sources of SNR Limitations.

Good, fast and cheap: pick any two.

—Anonymous

Increased measurement accuracy is a perennial goal. A new scheme that yields better accuracy is like money in the bank—it can be traded for improved speed, reduced cost, or a quieter life, so even if your accuracy is OK for now, keep your eyes peeled for improvements.

To improve a measurement, you have to know what limits it. In Chapters 1, 3, and 18, we talk about shot noise, which is the noise of the light itself, and Johnson noise, the noise contributed by the front end amplifier, as well as some less common types, for example, generation–recombination noise in photoconductors. In Sections 3.10.1 and 19.10.12, we talk about source noise. Although we talk a lot about shot noise limited measurements, what really matters is the SNR, understood broadly as the ratio of signal power to true noise plus spurious signals.

Other kinds of noise crop up all the time. The simplest ones are additive noise from the background signal, for example, shot noise and 100 or 120 Hz spurs from room lights.

Many systems (e.g., disc drives, video microscopes, thin film measurements with pattern below, structured light, and speckle interferometry) are not limited by true noise, but rather by motion or random variations in the medium or in the object being measured. These sources of uncertainty produce complicated data sets, with a lot of structure that is not relevant. Understanding how to pull a decent measurement out of this sort of thing requires some basic physics and a lot of head scratching and trial and error. Make sure you find out what the limiting noise source is going to be before you begin serious design.

Look for Other Constraints. Lots of measurement problems are easy in principle but hard under the circumstances. The canonical example is the World War II radar proximity fuze, which was able to survive being fired from a cannon (remember, this was before the transistor—it used collapsible vacuum tubes). A more common constraint is an *in situ* sensor, which usually has to be small and noninvasive, and must accept all sorts of background light, noise, vibration, and probably a nonideal measurement geometry as well.

These constraints will largely determine not only the packaging but the measurement principle as well; an *in situ* IR spectrometer probably won't be a moving-mirror FTIR, and an interferometer attached to a vacuum chamber with a turbopump or cryopump won't be working at DC.

Write a Specification. Any complicated system needs a detailed system-level specification. The system spec is not merely a list of operational requirements, for example, signal-to-noise ratio, tuning range, and gas mileage, but includes detailed specs of every interface in the system, including, for example, the dynamic range of the detector, the backplane of the card cage, the API and user interface of the software, and the purpose and signal levels of every interconnection between subsystems. This is especially necessary if more than two people are going to work on the system, because otherwise you'll have complete chaos, but even if you're working alone it forces you to think all that stuff out in advance.

TABLE 11.3. System Properties

<i>Passive Features</i>
<i>Layout:</i> aperture, field of view, resolution, etalon fringes, stray light
<i>Wavelength range:</i> wideband, single line, monochromatic
<i>Polarization:</i> <i>s</i> or <i>p</i> ; ellipticity, purity, birefringence, effects of coatings and mirrors, topological phase
<i>Fidelity:</i> image quality, aberrations, CTF, OTF, Strehl ratio, pupil functions, illumination functions
<i>Efficiency:</i> étendue, radiometry, photometry, materials, components
<i>Alignment:</i> mechanical design, adjustments, verniers, locking, pointing
<i>Component Quality:</i> quality specification, testing, vendor reliability
<i>Coatings:</i> efficiency, durability, cost, yield, test environment, wavefront degradation, blooming, absorption, thermal effects
<i>Stops and Baffles:</i> flare, type (field, aperture, hybrid), ghosts
<i>Filters:</i> wavelength, bandwidth, tuning, collimation, stability with time, temperature, humidity
<i>Detectors:</i> photon budget, type, quantum efficiency, gain, background, capacitance, noise, cooling
<i>Front End:</i> bandwidth and noise vs. capacitance and signal level
<i>Signal Processing:</i> SNR, spurs, $1/f$ noise, trade-offs between subsystems, frequency plan, bandwidth, pulse response, overload, data rate
<i>All Components:</i> cost, availability, second sources
 <i>Active Features (Includes All of the Above)</i>
<i>Stimulus Geometry:</i> size, 3D shape, divergence, beam control
<i>Choice of Source:</i> thermal, laser, flash
<i>Other Source Properties:</i> power, intensity stability, spectral stability, pointing stability, intensity noise, FM noise, lifetime
<i>Coherence:</i> spatial and temporal source properties, mode structure, speckle, coherence fluctuations
<i>Stability:</i> stray light, laser isolation, etalon fringes, drift, mode partition noise, coherence fluctuations

Example 11.2: Laser Shutoff. In a laser-based instrument controlled by a computer, you'll need to make sure that the laser turns off if the computer crashes or loses power. If you've thought about the software in advance, you'll realize that all the exception handlers and exit list processing in the world won't do this adequately, and that a very little additional hardware (e.g., a monostable that turns off the laser unless it's reset once a second by a software-generated heartbeat pulse) will make a robust design.

Include All Relevant Parameters. A partial list of system properties to think about, loosely modeled on Smart,[†] is presented in Table 11.3.

Solve Stupid Problems by Overkill. Once in a while you'll encounter a really stupid problem that stops you in your tracks. An AM detector is a good example; few people who have never designed one know how hard it can be to get it both fast and linear. If it's a small part of a big system, bring out the big guns: use a \$100 commercial module

[†]Anthony E. Smart, Folk wisdom in optical design. *Appl Opt.* 33, 8130–8132 December 1, 1994.

that's guaranteed to work, or use three cheaper ones with three different amplifier gains, and digitize them all—one's sure to be in its sweet spot that way.

The number of warts your system grows in this way is nearly insignificant compared with the schedule delays caused by not doing it. If they really bother you, get rid of them in an engineering change after the product is shipping.

Make Trade-offs Early. Some trade-offs affect the basic strategy of your measurement. For example, a sensitivity/cost trade-off will show that a 3 dB increase in sensitivity in a shot noise limited system will require $2\times$ the laser power or $\sqrt{2}\times$ the collector diameter—which is $2.8\times$ in weight and probably $2\times$ to $3\times$ in cost. This may be able to be saved by improving the signal processing, or coatings, or using a different measurement approach (e.g., TDI).

Identify Show-Stoppers Early. It's worth making a list of all your assumptions, no matter how stupid and obvious, and what would happen if they weren't true—in any complex measurement, at least one will probably be wrong (see Section 19.4).

Once you have your design specification done, hold a beer check. Get all your technical friends to gather round a big table for a couple of hours as you lay out the specs, and offer them a bounty of one beer for every mistake they find. Don't be bashful about it—this is an excellent use of everyone's time. Everybody learns something, and on average, the amount of waste and delay saved is much greater than six people times the two hours it takes.

Do Your Tool Building Early. There are usually tool-building tasks to be done, such as writing a data display program that will produce histograms, maps, and time series data for your customer. That program is probably extremely useful for hardware debugging too; it's much more informative to have graphs and reports than just screens full of numbers (or worse, files full of binary data). By one-third of the way through the schedule, you'll be feeling deadline pressure already, so the temptation is to do testing and characterization work using highly manual methods such as sitting there with a clipboard watching something that software could do better and faster. This temptation gets stronger rapidly as the deadline approaches, so if any tool building is going to occur, it has to start *right away*, even at the risk of building the wrong thing.

The amount of time those tools will save you during development is easy to underestimate; if you can spot what's wrong immediately, you can fix it quickly. Speeding up the rate at which you can iterate is the real key to getting your product to market efficiently, and that needs good tools.

Know What Limits You're Pushing and Why. As a corollary, you need to know what things are hard. Measurement complexity is one part of it, but things get difficult as you approach the limits of current engineering practice, as well. Table 11.2 has reasonable limits on a number of common parameters. These are not hard limits, and experts can certainly exceed most of them, but try to stick within them if you can; life gets more difficult out there. If you really need higher performance than this, budget extra time and money to make sure that those hard things can be done reliably.

Use Standard Parts. Design your optical system with standard parts (especially standard lenses) wherever possible. Make sure you know which ones are really standard (i.e.,

kept continuously in stock) and which are just there to decorate the catalog. A lot of catalogs have *diner syndrome*; there are 200 things on the menu but only 15 are available fresh.

Make Good Drawings. You may know exactly how to build your system, but unless you want to build every one yourself, and have an infallible memory, make engineering drawings good enough that a technician of average skill can easily read and understand them. Especially don't scrimp on the interfaces: drawing schematics of backplanes is very tedious but very necessary.

Annotate the drawings well, too; notes about what is normal and what to watch out for in testing and assembly are enormously useful in preventing people from sweeping blunders under the rug, intentionally or unintentionally. Some examples are what the signal and bias voltages should be, what the bandwidth and data rate are just here, what the purpose of this lens is, how pure the polarization should be here, or where on the base plate a nice reflection that can be used for focusing happens to hit. Bob Pease quotes Milligan's Law (named for an acquaintance of his): "If You Notice Anything Funny, Record Amount of Funny." Having the right value on the drawing makes this easy.

Some organizations discourage this, because the notes may become out of date with engineering changes, so that they may be misleading. This is a rotten excuse; the same is true *a fortiori* about comments in source code, and nobody is arguing for leaving those out. You just have to have the discipline to maintain the notes as well as the rest of the design.

It Isn't Finished Until the Test Stand's Done. A measurement consists of data plus ground truth; if there is no *a priori* way of knowing your data are good, you can't assume that they are. A cardinal principle of life, in engineering, mathematics, and cookery, is this: *If it isn't tested, it's broken*. A corollary is that your design isn't done until the test stand's done. Build good ones; for example, an integrating sphere may be expensive, but how many days' schedule slip are you risking, having to cobble something together that won't work as well?

Don't leave the test stand until the end; very often a minor change in the instrument, such as a well-thought-out test plug, will let you put the test software right inside the instrument where it belongs. That way, when a unit is assembled, it gets stuffed into the jig and turned on; a green LED says it passed, and a red LED says it failed. Calibration data are stored in EEPROM right on board, along with revision level, serial number, and other parameters that will come in handy later.

11.7 GUIDING PRINCIPLES

There's a whole lot in Chapter 10 about what measurement principles are good for which jobs, but it's worth some discussion here about how to decide on your own. The main goals are to get the best possible data with the most robust and cost-effective instrument, in a reasonable time; this requires maximizing the IPR.

Trust Freshman Physics. If you have some simple approximation that says the gizmo should be perfect, it'll usually be very good. There are lots and lots of examples, from the

massively parallel Fourier transforming action of a lens, to the way bipolar junction transistor (BJT) differential pairs split current with perfect linearity. We'll use a mechanical example.

Consider an active vibration isolation table, where a small optical breadboard is mounted on a few actuators (typically piezoelectric) to make it insensitive to vibrations coupled through the supporting structure. You might want to be able to null out a 100 nm amplitude vibration up to 10 kHz bandwidth, with a breadboard weighing 15 kg. Wiggling a large mass fast takes a lot of power even for such a small displacement: the mechanical power required to jiggle a mass m is $P = \mathbf{F} \cdot \mathbf{V}$, and $\mathbf{F} = m d\mathbf{V}/dt$. For a sinusoidal disturbance of amplitude s , the peak power is

$$P = \frac{1}{2} m s^2 \omega^3 \quad (11.1)$$

or 12.4 W, which for a piezo is really whomping. Besides, the breadboard will certainly resonate way below that, and who knows how much stuff will be loaded on it?

In fact, it isn't necessarily as impossible as all that, because what we're trying to do is keep the breadboard in an inertial frame of reference, and as we know from freshman physics, that requires zero force. Down at DC, the table has to follow the Earth, so very low frequency forces have to be transmitted, but not high frequency ones. Thus to isolate at high frequencies, all the piezos have to do is to move *themselves* at that rate, to keep the table from pushing on the breadboard. Not only is this much easier, but if it is done properly (e.g., mounting the piezos on the table instead of the breadboard), there should be no force on the breadboard to excite its resonances. This won't help forces applied directly to the breadboard itself, for example, from moving stages or vibrating fans, but it sure makes conducted noise easier to handle, and what's more, the attenuation needn't be a strong function of how much mass you put on the table, since it is being kept still.

Where it will make a difference is in the transfer function of the piezo actuators, which will make feedback loops impossible to design. Put an accelerometer under each piezo, to measure the forcing function a known time before it reaches the piezo, and use feedforward (see Section 15.12.7). Use an adaptive digital filter that adjusts its tap weights to minimize the residual vibration.

Believe Your Photon Budget. Your photon budget is not a lucky rabbit's foot, to be kept in your pocket, but a map of your way home in unfamiliar territory: it's vital to keep looking at it, and tracing your way as you go. You can waste a lot of time lost in some pretty unpleasant places if you don't make careful, *numerical* measurements of signal strength, noise floor, pulse width, and so on. We've belabored this and belabored it some more, but it's astonishing how many designers don't do it, even when their livelihood depends on the results.

Reduce the Background. Reducing the background signal is a very important way of improving the sensitivity and stability of your measurement. Doing measurements at AC (e.g., chopping, heterodyning, time-resolved pulses) will get you out of the $1/f$ noise region, providing that your modulation scheme doesn't move the background up in frequency along with your signal, or add significant artifacts of its own (e.g., beats between signal and chopping frequencies). We talked about that extensively in Section 10.2.

Don't Yearn for Greener Pastures. Electro-optical system design nearly always makes us do some things we're not good at or we're not sure will work. There's an insidious tendency for the problems in unfamiliar fields to look small. Software people sometimes seem to think that the way to do any optical thing is to stick a video camera on a frame grabber. Optics people, conversely, will sometimes try to shift off onto software what is best achieved with a pot, for example, fixed gains and offsets on a high speed data stream, or will rely on calibration to remove a nonlinearity that should have been designed out in the electronics. Get some good advice, backed up with more than just professional opinion (i.e., prejudice) before adopting a design that relies on a technology you don't understand thoroughly.

Try to stick to what you have a comparative advantage in. When faced with two nearly equally good ways of doing a measurement, choose the one that is a better fit to your skills. If you're a gifted electronic designer, start off building stuff with lots of electronics and only enough of everything else to achieve a good signal-to-noise ratio, and branch out into more complicated optics as your expertise broadens.

Model It. Get into the habit of modeling what you're planning to build. This doesn't mean building full 3D electromagnetic models of everything, or even finite-element models of the response of the system to vibration. It does mean calculating how the aberrations of the system propagate to the business end, what the SNR is going to be, how the system performance varies with filter bandwidths, scan speeds, dark current, and so on. Go through your signal processing strategy in exact detail, for example, how you're going to separate closely spaced peaks, and how the time/bandwidth behavior of your filters will affect the measurement. You can build confidence this way, and it will help you to understand any anomalies that surface later. Finding problems early on gives you time to work around them before they become large and hairy.

Note that the model doesn't have to involve computers; math programs are convenient and popular, but an analog model or a simple physical model can be very helpful (e.g., the pad stack capacitance model of Section 16.2.6), and of course for generating lots of intuition in a short time there's no substitute for grinding out an analytical model manually. Some combination of these is usually best, but don't fall into the trap of relying solely on computer modeling, lest you become a one-trick pony.

Get Ground Truth. In the remote sensing business, you have a satellite that can take immense amounts of data, for example, hyperspectral images of the entire Earth at 10 m resolution, with coverage once a week at the same sun-time each day. This is a wonderful thing—imagine all the information in those pictures. When you get them to look at, though, it's just this oddly colored map of surface reflectance averaged in some way over that 10 m diameter area. What in the world is really down there? Someone has to travel there and look, and look, and look, to correlate the imagery with what's there on the ground—*ground truth*, it's called, and the usefulness of your measurement depends utterly on the uniqueness of the mapping between it and the imagery. Satellite imagery isn't much use looking for deforestation if forest and mud have the same spectral signature with your apparatus. The squishier your measurement gets, the more ground truth you need, and the luckier you have to be to find a unique mapping. This is actually a very general problem, from microscopy to astronomy.

Keep It Simple. People buy or use things because they make life easier enough to make the pain of buying or using worthwhile. If you are the only person who can use a system,

you will be the only person who does. Instruments have to be good enough to ignore. For example, if your instrument uses expendable items such as tungsten bulbs or small amounts of inert gas, consider having online spares that can be activated automatically, allowing maintenance to be done at some convenient time rather than *right now*. Try to achieve graceful degradation, where a stressed or damaged system functions as well as it can until replaced or repaired, and gives an intelligent indication of what state it's in, in real time.

Consider Moving the Goal Posts. One good way to improve the IPR is redefining the problem to make it easy. There are no hard or fast rules, but when you find yourself murmuring “This would be a lot easier with more light,” or something of that sort, always ask “Well, so why can't I have more?”

Consider scanning or time gating to improve SNR in Johnson noise or background limited systems, or use retroreflective tape to improve signal levels. Noise cancelers or feedback stabilizers can dramatically reduce laser noise. Single-sideband mixers are 3 dB quieter than ordinary ones, since they get rid of the noise at the image frequency. Homodyne detection (IF at DC) has the same advantage, but is phase sensitive.

A tunable diode laser gas spectrometer with etalon fringe problems can be improved by two orders of magnitude by taking one measurement with the sample, subtracting a second one taken after admitting room air to the sample cell; the sample absorption will be pressure-broadened into nonexistence, but the etalon fringes will be closely similar.[†]

Build the Test Fixture into the Instrument. Nowadays most of the effort in designing test fixtures goes into their software, and the instrument can hold that just as well as the test stand. Building the testing software in costs little and ensures that there are no problems with incompatible versions of the instrument and tester. Instruments that can calibrate themselves online have hugely less maintenance cost, can correct for drifts and measurement errors in near-real time, and can tell you reliably about their own health.

Keep Gross Errors Obvious. Keep zero signal on-scale. Bring out the DC photocurrent so you can see how much light you've got. Put in a pilot light. Add a viewer. Make failures obvious and alignment easy, and you've got an instrument whose maintenance is painless and hence cheap. It's very frustrating for your customers to have to send an instrument back to the manufacturer. The author once had a shiny new shear-plate measuring interferometer whose Nitinol-flexure actuator broke. The distributor insisted on having it sent back, whereupon they installed the new part grossly misaligned. (Would you believe 15°?) It took the best part of a month, and the author had to realign the instrument anyway. If your customer is competent, and your instrument is simple, just ship him the part and the manual page. If he screws it up, it's fair to make him pay for the replacement; if he doesn't, he's happy and you're happy.

11.8 DESIGN FOR ALIGNMENT

Systems whose alignment is easy and stable don't appear by accident. As you're laying out the system, figure out the alignment algorithm and how you'll test it. Make sure

[†]At least they will if the cell is short enough; for larger cells, you can do several sweeps as the cell bleeds up to atmosphere, and correct for the etalon fringes' phase shift with air pressure.

each subsystem can be aligned and verified independently. Sections 12.8 and 12.9 have a fair amount on alignment, and you should know something about it before beginning a design.

Use Corner Cube Interferometers. The best alignments are the ones you don't have to do. Using corner cubes instead of flat mirrors in interferometers makes alignment the optician's problem, and a stock of corner cubes is like a stock of perfectly aligned interferometers. Use them whenever you can afford to; they do cost money and photons. Watch out for the polarization shift in TIR cubes.

Use the Poor Man's Corner Cube: Retroreflective Tape. As we saw in Section 7.8, retroreflective tape is an excellent solution to alignment worries in many systems. It is commonly used in fiber-fed bulk optics sensors, where it greatly simplifies the problem of how to get the light back into a single-mode fiber. You can use it to enormously increase the sample space accessible from a small sensor and illuminator, at least for optically thin paths where diffuse scatter doesn't dominate.

Put in a Viewer—You Can't Align What You Can't See. As a corollary, try to have as many dials as you have knobs. Aligning a system with too many interacting adjustments and not enough independent readouts is very frustrating indeed. An example is doing initial alignment of a multimirror laser, where nothing at all happens until you get very close, and then every adjustment seems to make it worse. For such a system, an auxiliary alignment system such as a HeNe laser plus a set of cross hairs or iris diaphragms and a viewing system is a great help.

Instruments requiring fine alignment should have a microscope trinocular head[†] designed into them; since these expect parallel light coming in, it's very simple—pick a convenient pupil (i.e., the Fourier transform plane of whatever you're trying to align), and just put in a beamsplitter to shine some of the returned light into the trinocular head, which you mount somewhere convenient. It looks a bit odd having microscope heads sticking out of your setup, but it surely makes alignment easy. Try it just once and you'll do it forever. In some cases, you may need to put in a field lens (see Section 12.3.14).

Note: Don't do this in a laser system without great care—be sure that there's enough attenuation so that the laser light cannot exceed 100 μW CW per eyepiece, even in fault conditions, for example, a tilted sample that sends the main beam right toward the eyepieces. Visible light that bright will make your eyes hurt, and IR brighter than that may damage them. Filter out all the invisible radiation just to be on the safe side.

Use Verniers. If you need a fine adjustment, put in a vernier, such as a full wave plate which can be wobbled to get exact correction of small polarization shifts, a couple of wedges that can be rotated for fine beam steering, or a glass plate glued to a single ball bearing pivot[‡] for x - y beam position adjustment. Don't put a full range adjustment on an expensive mount and expect to do as well. Make sure you consider the effects of such a change, however; most full wave plates based on birefringence don't work well in white-light systems, pivoting plates will produce fringes, and rotating wedges make the pupil move.

[†]The part that the eyepieces slide into.

[‡]As used in those ugly pen-and-pencil desk sets for the-person-who-has-everything.

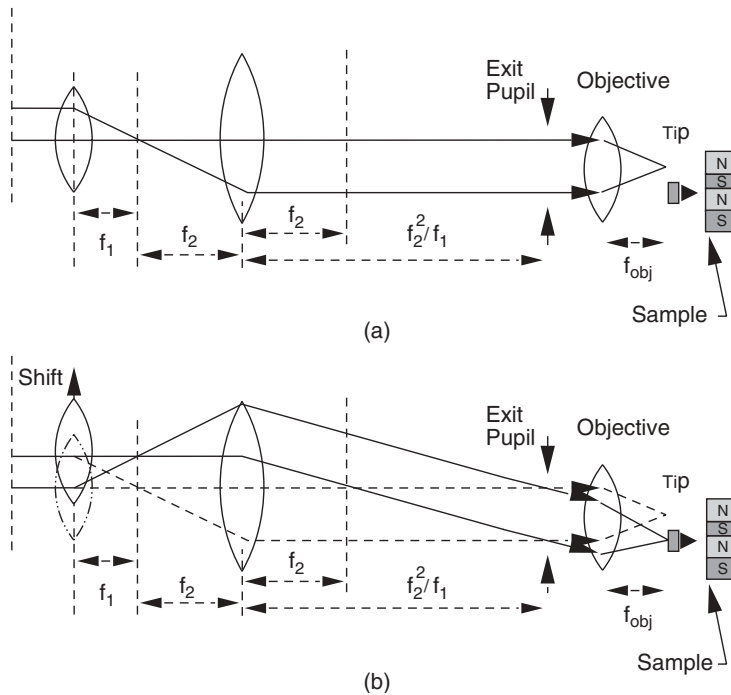


Figure 11.2. Shifting lens A moves the focused spot without misaligning the interferometer, making tip changing a lot easier.

Adjust the Right Thing. It is frequently possible to tailor the adjustment to the purpose. Consider a magnetic force microscope using cantilever tips made from fine nickel wire, electropolished down to $5\text{--}10\ \mu\text{m}$ diameter. If we want to use a heterodyne Mach–Zehnder interferometer to detect the vibration of the cantilever, the beam must be focused to a spot less than $2\ \mu\text{m}$ in diameter with a microscope objective. The cantilever has to be changed often and is not easy to position that accurately, but we can't stand to realign the interferometer every time.

One good solution is the auxiliary two-lens system of Figure 11.2. Shifting lens 1 sideways moves the beam around in angle in the pupil of the objective, which causes the focused spot to move telecentrically (i.e., the cone of light doesn't tilt as it moves). Moving lens 1 in and out a bit gives a fine focus adjustment. Because the cone is always normal to the cantilever, light reflecting from the tip retraces its path regardless of where exactly the cantilever winds up.[†] Thus moving the beam around doesn't misalign the interferometer. All you have to do is get the beam onto the cantilever and focus to get a good signal, and as long as there's a viewer, that takes about 10 seconds.

Watch Out for Temperature Gradients. Gradients are much worse than uniform variations with time. They cause things to *bend*, so that the error builds up and builds up with distance. Objects bend, and the temperature coefficients of length and refractive index mean that beams bend too—from the Schlieren effect and from the thermal

[†]There's a small amount of vignetting, but not enough to notice.

distortion of the figures of lenses and mirrors. See the chapter problems (available on the Web at <http://electrooptical.net/www/beos2e/problems2.pdf>) for more.

Usually Follow the Leader. There are a lot of nonobvious potholes in optical instruments. If something in someone else’s instrument looks much too complicated for what it does, it probably is; but it might not be, and the reason for it may be very subtle. Do not be too quick to label other designers idiots, but find out why they did it that way. For example, a Savart plate has unequal phase delay in the two beams, unlike a Wollaston. A nonzero path difference interferometer is vulnerable to frequency noise in the laser; even a small fraction of a coherence length will produce a large increase in the noise of a bright-field measurement if you’re not careful (see Section 19.1.1).

Don’t Always Follow the Leader. A good design hack can overcome some pretty ugly limitations. Two-photon Doppler-free spectroscopy is an excellent example (see Example 1.1), and there are lots of others.

11.9 TURNING A PROTOTYPE INTO A PRODUCT

With the emphasis placed on prototyping elsewhere in this book, it is easy to assume that a working prototype translates into a working instrument by itself. This is of course not true, but by avoiding the most common mistakes, you can improve your odds enormously.

11.9.1 Be Very Careful Of “Minor” Optical Design Changes

He [Ferdinand Porsche] is a very amiable man but let me give you this advice. You must shut him up in a cage with seven locks and let him design his engine inside it. Let him hand you the blueprints through the bars. But for heaven’s sake don’t ever let him see the drawings or the engine ever again. Otherwise he’ll ruin you.

—Ernst Heinkel, *He 1000*

There really isn’t any such thing as a minor optical design change, because each optical element does so many things, and we only know some of them. This means that a mildly redesigned optical system has to be tested as though it were an entirely new optical system, that is, breadboarded and run through the whole range of tests (see Section 19.2.1). Failure to understand this can be very expensive; the liquid-borne particle sensor of Section 19.1.1 was derailed by replacing two elements, a HeNe laser and a Nomarski prism. Ironically, it would have worked with either of the two changes, just not both, and a breadboard would have shown the problem before the trouble got started.

11.9.2 Don’t Design in Etalon Fringes

In turning a prototype into a limited production product, there is a temptation to take all the optical mount designs and turn them into one big complicated custom mount that can be made in one step on a CNC milling machine. This saves money by reducing the amount of handling and assembly, but if it is done naively, all the incidental minor misalignments that protected your prototype from etalon fringes will be zeroed out; the resulting fringes will blow you right out of the water—just like Microbench, only worse. Not only that, but you’re at the mercy of the tolerances of the components you’re using;

for example, lenses often come with a 5% tolerance on their focal length. Make provision for strategically chosen adjustments and baffles, and watch those stray beams.

11.9.3 Handle Demo Karma Gracefully

Besides pretty pictures, your sales job will include demonstrating your system in action (unless you make sensors for explosions). Prototypes are a bit flaky, always. Demonstrations of prototypes are dangerous; the stakes may be high and preparation time limited, which makes for jitters. Jittery people break things, so the death rate of prototypes is ironically highest during demos, leading to the phenomenon known as *demo karma*, where you break the prototype at 5 p.m. the day before. Even working all night to fix it will usually not gain you enough merit to get it going in time.

Resist the temptation to cut corners in preparing for demos. You know, hot plugging boards and cables, wiping lenses on your shirt instead of cleaning them properly, stuffing loose bits in for test purposes without securing them, generally doing things in a hurry. Your system will be more likely to survive if you get it in shape a few days in advance, in a lab with spares for all critical parts, and then spend the three days before the demo polishing up your talk, cleaning your desk, or writing a paper. Unfortunately, this requires more willpower than a pastry chef on a diet. It remains an unsolved problem.

Building Optical Systems

Here is Edward Bear, coming downstairs now, bump, bump, bump on the back of his head, behind Christopher Robin. It is, as far as he knows, the only way of coming downstairs, but sometimes he feels there really is another way, if only he could stop bumping for a moment and think of it.

—A. A. Milne, *Winnie The Pooh*

12.1 INTRODUCTION

Once you have a measurement principle, a photon budget, and a signal processing strategy for your measurement (at least in rough), you'll want to try it out, good and hard. Don't chicken out on this—the best failures are the ones that happen early. The first prototype serves to test your assumptions, not only about the physics involved but also the suitability of your strategy when faced with real-world components and situations. Don't tiptoe around the prototype; if the instrument is a big success, it'll have to work well without you around, and knowing what vulnerabilities you have to overcome in the real design is extremely valuable. Bang on it.

Plan your experimental approach to eliminate the most egregious technical risks first—the ones that can reduce your project to rubble. Things such as whether any affordable diode laser can reliably be made to simultaneously current-tune, frequency-modulate, and remain stable. Or whether you really can get the nonuniformity, readout noise, and dark current of your CCD detector down low enough without cryogenic cooling. Or whether a mass produced molded optical system can have its aberrations custom-adjusted in a stable way by indenting a TIR surface with a ball bearing on an XYZ stage.

After the first prototype has set your mind at rest, modify it or build another one to demonstrate end-to-end function and (a bit later) performance. Don't break it in the process—it is important to have a system functioning at some level all the time. This helps with sales pitches but, more importantly, keeps you from going too far astray if you make a design mistake. We've seen that an electro-optical system will go beyond our range of competence often, so this sort of belt-and-suspenders approach is vital to prevent disasters.

This chapter concentrates on how to go about assembling and aligning prototypes and limited-production optical systems in instruments. High volume optomechanical engineering is a specialized field, and if you're planning on building more than a dozen units, it is very worthwhile to consult someone who's good at it—it's not an amateur's job.

12.2 BUILD WHAT YOU DESIGNED

There is something about building a prototype system on an optical table or breadboard that sometimes seems to disable the higher functions of a designer's brain. After spending three days on the detailed Fourier optics of the system, he might accept a 10% vignetting of a Gaussian beam by a poorly placed prism; or fail to notice strong scatter brightening up the vicinity of his photodiodes; or carefully adjust and cement a diode laser collimator without making sure the beam axis is really horizontal.[†]

Since optical theory works pretty well, these examples of catastrophic synapse failure have the unpleasant consequences we'd expect. The first imperative in building an instrument is to really build what you designed. If the beam isn't clean, find out why. If it gets vignetted, move it. If there's a lot of stray light around, you've probably got a dirty element, a wayward beam hitting a ground glass surface, or a poor quality component with a lot of surface or internal scatter. It's in there, and it's doing you no good. Find it.

Debugging an optical system is an onion problem: you peel off a layer, cry, and peel off the next. A sufficiently hard debugging problem equals a failure. Being compulsive about tracking down little unexpected things you notice will be richly rewarded: your expertise with the system will grow more rapidly, and it will have fewer flaky problems. Debugging therefore gets easier quadratically.

12.3 ASSEMBLING LAB SYSTEMS

Most lab systems are built out of a mixture of purchased mounts, custom-machined pieces, and random chunks of metal held together with glue, double-sticky tape, and a few bolts. As a prototype gets modified more and more, the nice mounts deteriorate and there are more and more glued things, which makes the prototype get flakier and flakier. This is the normal life cycle. The key is not to break too many expensive optical parts when the glue fails under normal wear and tear.

12.3.1 Build Horizontally

Even if your instrument will be used with the optical path predominantly vertical (e.g., a microscope), build the prototype horizontally. It'll be much more stable and much easier to assemble and align. Also—crucially—it won't need so many custom-machined parts. All laser systems are interferometers, usually unintentionally. There will almost certainly be stray fringes there that will make life even more difficult if you build a floppy system.

12.3.2 Use Metal

Plastic is often a lot easier to work with than metal, being attractive, flexible, easily cut, and corrosion resistant. These attributes make it great for boxes, thin shims, baffles, and shrouds. Unfortunately, plastics have horrible dimensional instability. Your average plastic has $10\times$ the CTE of ordinary metal, spreads heat 500 times more slowly (so that temperature changes lead to $\sim 5000\times$ more warping). It also creeps under load and absorbs enough moisture to make it swell by a percent or so. Some plastics that are tops

[†]The author is not immune: that's where the examples came from.

in other ways (e.g., PTFE) have the worst thermal and mechanical properties. If you use plastics for optical mounts, be sure you understand, quantify, and accommodate these effects. Soft sheet aluminum (1100-T0) is almost as easy to work with and has none of these problems.

12.3.3 Scribble on the Optical Table

We often need to remove and replace some mount or stage, to an accuracy of better than 1 mm. This is difficult on optical tables and breadboards, which usually use 1/4-20 bolts in plain holes. To avoid problems, do one of two things: build a corner out of two post flanges for the mount or stage to snuggle up into, or mark the optical table by running a scribe or sharp-tip felt pen around the perimeter of the mount. You can easily do 0.05 mm this way, and especially with mag bases or universal mounts (the ones that can twist and translate), it saves a great deal of irritation.

12.3.4 Mounts

All the manufacturers' catalogs are full of beautiful optical mounts, translation stages, and so on; you can get six-axis adjustable stages with differential-screw micrometers and amazing specifications, if you're rich enough. You can also get snookered by specsmanship. A six-axis stage is about as stiff as overcooked broccoli, and a so-called kinematic two-tilt mirror mount exhibits hysteresis, wobble, nonorthogonality, and serious stick-slip energy storage[†] (see Section 12.9.2). Your best bet is to use fixed mounts for most things and keep the adjustable ones for where they're really needed. Whatever you do, don't believe translation stage manufacturers when they claim 0.1 μm stability and repeatability; they may have seen that in the factory, but for us it's more like 0.5 μm on a good day, with a new stage and a light load. Bang the stage just once, and the resolution becomes 2 μm , forever.

The basic trouble with most mounts is that they have a huge adjustment range and attempt to get precision and stability by using a fine thread. This is okay for some things, but if you think you can get a stable setup this way, think again.

You can get flexible shafts (made of wire rope) that connect the adjustment screws to handwheels bolted to the table, and these certainly reduce the amount of incidental motion due to hand pressure on the mount. They do nothing to reduce the mount's inherent floppiness, however; for that, you need a good design using verniers instead of coarse tweaks with fine screws.

12.3.5 Use Microbench for Complicated Systems

Individual elements in individual mounts can be placed anywhere you like. That's the good news and also the bad news. If you need several elements to be aligned correctly (e.g., a beam expander), use Microbench[‡] or the equivalent—systems based on mounts that thread like beads on centerless-ground steel rods. Their performance is quite good, and they allow easy adjustment of focusing. A hidden benefit is that it becomes natural to use fixed mounts for almost everything.

[†]The phenomenon is familiar from earthquakes.

[‡]Manufactured by Spindler & Hoyer.

There are a few problems with Microbench, too. The enforced alignment encourages etalon fringes and makes it difficult to use beam paths at odd angles. Components mounted on the sides of the rails are not as well located, because they rely on screws in plain holes (which are not stabilized by the rods).

12.3.6 Machine a Base Plate

Systems using noncollimated beams, narrow-angle components (e.g., Brewster windows or Pockels cells), or off-axis mirrors are hard to align from scratch with mounts bolted to an optical table. It gets a lot easier if you can make a base plate on a milling machine. Choose enough strategically placed adjustments and nail all the rest still. A base plate can easily have shallow slots and corners milled into it to make rough alignment automatic and fine alignment easy. This does require calculating the optical layout, but that is not normally too onerous. If the system is going to require a lot of fiddling with, you can use dowel pins instead of slots and corners, but that needs precise work—tolerances of $0.0001''$ ($2.5\ \mu\text{m}$) are typical.

12.3.7 Use Irises

Aligning beams gets a lot easier if you have a few iris diaphragms sprinkled around. An iris on each end of the beam path makes it simple to get the beam centered and allows you to tighten the irises down to get more precise location. This is especially useful with micro-optic components in Microbench setups, where just finding where you are can be a huge pain.

12.3.8 Getting the Right Height

A perennial hassle with optical mounts is that they're never the right height. If you're using sliding posts, that isn't the end of the world, but whenever we need to align one subassembly with another, height is always a problem. About 15% of the work of assembling the optical system is getting it the right height, and the easy methods (e.g., lab jacks and z translation stages) aren't very stable or rigid. One certainly needs a decent stock of shims, but they're never quite right, and the amount of cutting and trying is astronomical. What to do?

Use a pad of sticky notes. They come apart and go back together again, unlike shim stock, and they're a pretty repeatable 0.1 mm thick each; the thickness of the glue is negligible. There's no other shim material that makes it so easy to adjust beam positions in $100\ \mu\text{m}$ increments, and you can use a paper punch to make the mounting holes. For thick shims (say, over 0.5 mm), you'll usually need something more stable, but a couple of metal shims and a few residual stickies work fine. (If the mild sleaziness of this trick bothers you, trim the paper with a razor blade after assembly so nobody sees what you did.)

Aluminum soda pop cans are a bit thinner, $70\text{--}100\ \mu\text{m}$ thick depending on what height on the can you measure. They are made of very hard aluminum, so they work well in places where sticky notes are unsuitable. Their large thickness nonuniformity and their tendency to pucker makes pop can shims unsuitable for stacking. The wedge-shaped nonuniformity can be a help, since you can slide the shim a little for fine adjustments.

12.3.9 Light-Tightness

A dark-field system, especially a photon counting or high-accuracy photometric one, has got to be light-tight. This does not happen without thought. A butt joint is never light-tight—at least two 90° bends between close-fitting black surfaces is required, as when two tubes are butted together inside a sleeve, and more is better. A really light-tight system performs just the same in bright sunlight as in darkness.

12.3.10 Chop Up 35 mm SLR Cameras

A 35 mm SLR camera is a surprisingly convenient and economical optical toolkit for many things. Its two best features are a through-the-lens viewfinder and a film plane, where you can just butt your image receiver (e.g., a fiber plate, CCD, spatial filter, or quad cell) and be sure that it's focused where the viewfinder says it is. We do have to overcome our natural revulsion at hacking beautiful things, but they're so cheap nowadays that we can buy one to look at and one to sacrifice (the manufacturer is in it for the money anyway).

Modern cameras usually don't have mirror lockup levers, and many don't even have T or B shutter settings, which allow the shutter to be kept open continuously. A strategically placed piece of music wire will allow you to raise and lower the mirror manually, and a milling machine or a Dremel tool will remove the shutter in a couple of minutes.

12.3.11 Try to Use at Least One Screw per Component

It's usually the small mounts stuck on the corner of some translation stage with Crazy Glue that get knocked off when your hand slips.[†] Optical mounts are coated with Tufam or some other smooth and durable anodized finish, which glue does not stick to very well. Don't hesitate to drill holes in commercial mounts; they are comparatively cheap and will survive most kinds of abuse apart from banging them and brinelling the bearings. Use a small cordless hand drill and hand taps to put threaded holes in odd places on mounts, then use the glue to keep the mount from rotating. Setups built this way survive better, and the mounts and stages actually accumulate less damage, since you don't have to scrape the anodizing off to glue them.

12.3.12 Detector Alignment Needs Thought

Aligning dimly visible laser beams on small detectors is enough to drive you nuts. The surface is specular, so you can't go by the reflection unless you have a strong visible beam. The detectors are snuggled right up into the optical system, so you can't get a good view of them anyway. The best defense is to use the accuracy of the milling machine to do your work for you. Position the detector accurately in its module, position the beam accurately, then drop the module in (preferably located using roll pins or taper pins). You're done.

A good method for measuring the beam position is to mount an eyepiece reticle (available inexpensively, e.g., from Edmund) in the same type of mount you're using for the detector, and point a video camera plus macro lens at it. Put a thin coat of white

[†]Written an hour after destroying a \$400 diffraction grating this way.

spray paint on the ruled side of the reticle to make the beams show up and eliminate parallax error. Shine the beams on the painted surface and look from the opposite side, where there is a nice bright dot with a scale superimposed. A frame grabber will make a permanent record for your lab book.

It's harder with multiple detectors, for example a differential system with two photodiodes and a Wollaston prism. One good solution is to put the detector die right in the plane of the top of the box, or embed it in a small grounded plate whose top surface is coplanar with the die. The benefit of this is that you can slide the detector module off to one side by some known amount, adjust the beam on a target on the box where it's nice and visible, and then slide the detector back again. You can draw the target under a binocular microscope with an eyepiece reticle. It doesn't get any easier than that.

A viewer that can be arranged to look at the photodiodes near normal incidence (and with adequate magnification) is even better, but normally much harder to arrange; if there's space, a clean microscope cover glass at an angle, with a camera looking at the reflection of the diodes, can be a godsend; the cover glass is too thin ($100\ \mu\text{m}$) to offset the beams by much. Watch out for the fringes, though, and don't leave it in there when you're done. Another alternative is a mockup detector module with a piece of ground glass with two dots of opaque where the centers of the diodes lie. In either case, you'll need to put an external alignment reference on the module, e.g. a couple of roll pins or (for bigger detectors) countersunk mounting screws, which are mildly self-locating. It is really worth the investment of design and machining time to get the detector alignment right.

Aside: Photographic Paper. In the visible and UV, you can measure the positions of beams in inaccessible spots by using a chunk of photographic printing paper. You don't develop it or anything; just stick it in and turn on the beam for a moment. The resulting black spots are stable under room lights for at least several minutes, easily long enough for you to do your measurement with calipers or a measuring microscope. This is a great way to measure the separation of two spots for two detectors, or to measure the position of the beam in the clear aperture when you can't get at it to use the white-painted reticle trick.

12.3.13 Do-It-Yourself Spatial Filters

Making decent pinholes yourself is easy if you have an excimer or femtosecond laser handy to drill them. If you aren't so fortunate, you can use a trick: cut several small sheets of aluminum foil, make them into a stack, and stick a sewing needle partway through the stack—through the top layer, but not all the way out the other side. When you take the stack apart, you'll have a series of smaller and smaller holes; the smallest one will probably be useful. Measure the diameter by shining a (collimated) laser beam on it, finding the angular diameter of the first diffraction null and using (9.36). You can also do it with an evaporator and some polystyrene latex spheres: apply a (very dilute) aerosol of the spheres to a glass surface, and then put down $25\ \text{\AA}$ of chromium and $\sim 250\ \text{\AA}$ of gold or aluminum by directional evaporation, and wash off the spheres. This gives you very accurate diameters and clean edges, but zero control over pinhole position.

You can easily make good slits in the lab from single-edged razor blades. Their edges have a 1 mm wide bevel, which is cut at a big enough angle to avoid reflection back into a laser, and they're very straight. Utility knife blades (the big trapezoidal ones) have

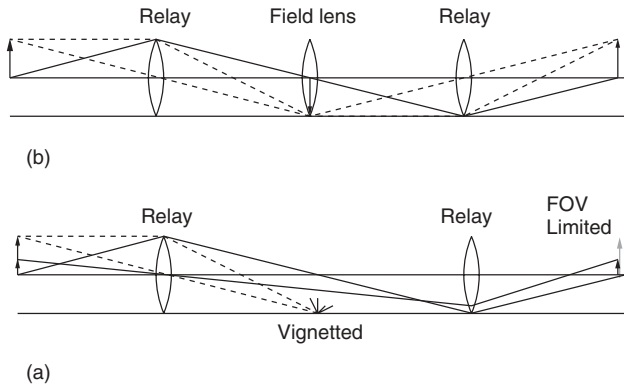


Figure 12.1. A field lens placed at an intermediate image doesn't change the image position but increases the field of view by corralling rays that would otherwise be vignetted. (a) Double relay lens has limited FOV. (b) Adding a field lens roughly doubles the FOV diameter without requiring a larger diameter system.

wider bevels for broader beams, and shaving razor blades are very smooth, although fiddly to use since you have to disassemble the plastic cartridge.

The easy way to make slit assemblies is to take a piece of plastic magnetic strip (as used in refrigerator magnets, with glue on one side), and drill a hole in it for the beam. Arrange the razor blades on the strip, adjust them until they're just right (perhaps with a shim), and then put on a drop of glue to secure them. This trick is great for tight spaces, and since you can stack several blades, it's easy to make weird shapes. Silicon makes atomically sharp cleaves in the [111] direction, so you can make good sharp edges by cleaving.

It's often convenient to use a slit with a slight V in it, made by crossing a pair of razor blades at a very small angle. This is useful with round beams, where the V allows adjustment of the slit width by moving it up and down, but doesn't affect the filtering operation much. You obviously wouldn't do this in a grating spectrometer, where the whole length of the slit is used at once.

12.3.14 Field Lenses

Optical detectors are not phase sensitive, so we can put an arbitrary phase screen at an image without damaging it. If we put a lens there, it doesn't change the image intensity distribution, but it does redirect the rays, as shown in Figure 12.1. This helps preserve the field of view of optical systems without requiring absurdly large diameter elements. The same idea, redirecting rays without altering the image, is used in Section 5.7.9 to improve rear projection screens.

12.4 ALIGNMENT AND TESTING

Testing optical components to verify their performance is a big field, and there isn't space to do it justice here. Instrument builders usually buy their lenses and mirrors, so exact measurement of figure errors is the vendor's problem. What we need to be able to do is

to verify that a particular component is doing its job, align a complex optical system, and make sure our beam or image quality is what we expect. Malacara and colleagues give a very extensive discussion of many kinds of optical tests (which portions of this section follow), plus good references, and Yoder talks a lot about how to design mounts and align optics within a complex lens. The old-but-good MIL-HDBK-141 (Optical Design) covers much of the same ground and is a free download. (See the Appendix for these and other references.) We won't go beyond qualitative tests of the sort that are useful in debugging and troubleshooting, and instead concentrate on alignment.

12.4.1 Things to Count On

Considerations like the following underlie all the simple alignment schemes we'll consider here:

1. Light in a uniform medium travels in straight lines.
2. Lines parallel to a given line are parallel to each other.
3. Making a beam retrace its path exactly guarantees that it was normally incident (the *autocollimation* condition).
4. Rays passing through the center of a lens are undeviated in angle (although they are displaced slightly due to thickness of the lens).
5. A spherical surface rotated eccentrically will make a beam wobble around.
6. An accurate spherical surface is easy to make, and once it exists, it can be located very accurately by shoving it up against a well-turned circular mounting ring. As long as the optical surface stays in contact with the ring all round, the center of the sphere will lie on the axis of the circle.
7. A lens element has two surfaces. For any two nonconcentric spheres, there is exactly one line that joins their centers, and this is the lens axis. That means that within our tolerances for thickness variation in the finished lens, we can polish the two surfaces more or less independently and still wind up with the right lens when we're all done.
8. A beam that hits the vertex of an accurate corner cube is autocollimated.

12.4.2 Clamping

One consequence of facts 6 and 7 is that a lens barrel or mounting cell of high enough precision can automatically center a lens element, by forcing both surfaces into intimate contact with two circular rings lying in parallel planes, as shown in Figure 12.2. Since the two circles are coaxial, the centers of both spherical surfaces are forced to lie on the axis too. Yoder says that this produces adequate centration as long as

$$Z = \left| \frac{Y_1}{2R_1} - \frac{Y_2}{2R_2} \right| > 0.07, \quad (12.1)$$

where Y_i is the radius of the i th clamp ring (remember that $R_2 < 0$ for a biconvex lens). Thus a lens element can usually be adequately centered by clamping it on its optical surfaces.

It is not altogether trivial to get the two rings aligned sufficiently well, especially if one is located only by its threaded edge; a high precision running fit of one ring into the




Figure 12.2. Clamping between two parallel circular rings enforces centration.

other works better. If you secure the ring with a threaded section, make sure the threads are loose, or they may bind due to overconstraint.

12.4.3 Soft Lenses

For soft or friable lens materials (e.g., ZnSe or NaCl), consider using an indium gasket. Indium is a very soft metal that has reasonable thermal conductivity and is available in wire form. Cut a length of indium wire slightly less than the circumference of the lens, place it in the fixed side of the mount, and put the lens in against it. A spring-loaded locating ring will be needed on the other side. Over several hours to a few days, the indium will deform to fit the surface of the lens, so that it will be held in without serious stress or deformation. This is especially helpful in high power CO₂ laser setups, where the high thermal conductivity of the indium and the excellent thermal joint between it and the lens improve the lens cooling dramatically. The bad news is that the centering action is much less since the lens has some freedom to rotate out of the plane of the mounting ring on the indium side.

12.4.4 Dimensional Stability

Clamping is a good example of a dimensionally stable mounting method; as long as thermal expansion doesn't cause the lens to lose contact with the ring, it'll stay still. When using glued joints, shrinkage causes the parts to move as the glue sets, and stability is limited by the thermal expansion of the glue. This can be dealt with by keeping the glue layer very thin, of course, but if there's a CTE mismatch, thin glue layers tend to delaminate due to extreme shear stress. Alternatively, you can load the glue with glass spheres of uniform diameter and apply a small preload while it cures. This forces the spheres into intimate contact with both surfaces (so the spacing will be set by the glass rather than the glue) while leaving the glue thick enough to handle the thermal strain. Choose glue with enough shrinkage that it will remain in tension over temperature—with 1% shrinkage and $10^{-4}/\text{K}$ CTE, it'll take 100°C change to make the glass spheres lose

contact. With thermosets such as epoxy, a postcure bake will help as well—5 minute epoxy is horribly gooey and unstable, but an hour at 100°C improves it amazingly.

12.4.5 Beam Quality

In laser-based systems, we are constantly concerned with beam quality—but what do we mean by that, exactly? There are various measures in common use, such as beam divergence, M^2 , Strehl ratio, RMS phase ripple, and so forth, and you’ll see them all quoted. All these are static or time-averaged properties, which take no account of noise, let alone the spatial variation of the noise[†] (see Section 2.13). We’re primarily interested in phase errors, which are what we mean by aberrations, but intensity nonuniformity is also important. The main thing is that if you want to use the fancy measurements, you need a fancy measuring interferometer that will give them to you, and if you don’t have one, they’re no use to you, except in deciding which lasers not to buy. Beam divergence is a partial exception, since you can measure that on the back wall of the lab for sufficiently small beams, but it isn’t as useful since it depends sensitively on beam diameter and intensity profile.

As a practical matter, the two best lab tests of beam quality are a measuring interferometer for quantitative work,[‡] and a combination of a dithered shear plate device (Section 12.6.4) for collimation plus a Foucault knife edge test for aberrations. For very small diameter beams (less than 3 mm in the visible), we can just shine the beam on the back wall of the lab and have a look.

Aside: Siegman’s M^2 . Many high powered lasers produce very ugly beams. Multistriple diode lasers, unstable resonators, and high pulse power beams don’t have a nice single transverse mode structure, and because their different modes have different frequencies, they can have deceptively smooth near-field patterns that look a lot like Gaussians but have far worse divergence. In these cases, Siegman’s M^2 beam propagation factor gives the best single-number summary of the beam propagation behavior. M^2 is the product of the second moments of the near-field pattern in (x, y) space and the far-field pattern in (u, v) space, normalized to 1 for a TEM₀₀ Gaussian beam. Taking the second moment (i.e., multiplying by $x^2 + y^2$ and integrating over the xy plane) vastly overemphasizes small artifacts in the wings of the beam, so M^2 values near 1 aren’t as useful as the Strehl ratio, which is why we haven’t used M^2 much in this book. On the other hand, uglier beams can easily have M^2 of 10 or more, in which case M^2 is informative and Strehl is useless.

12.4.6 Image Quality

A laser beam has a single definite path, which normally doesn’t change much except in preobjective scan systems. Imaging systems are expected to form images of any object presented to them, usually over a wide range of distance, field angle, and aperture, which makes image quality a much more elusive beast than good beam focusing.

The best measure of image quality depends on circumstances, but the most generally useful metric is the modulation transfer function as a function of position and orientation.

[†]By now, this should make you suspicious, but that’s not the main point here.

[‡]The author fondly remembers an old Wyko Ladite that was hard to align but gave beautiful results. It got sent to the warehouse because no other machine could read its data.

As we saw in Section 9.3.13, the MTF is the modulus squared of the OTF, normalized to unity at 0 frequency (which automatically corrects for vignetting and power losses). The usual quick method for checking this uses a radial bar pattern with a circle drawn on it to show the nominal resolution limit. Moving this around the field gives a good idea of the limiting resolution versus position. For a more accurate version, we use a set of sinusoidal amplitude gratings of different pitches and measure their contrast as a function of position, azimuth, and frequency.

12.5 OPTICAL ASSEMBLY AND ALIGNMENT PHILOSOPHY

There are two general classes of alignment test, based on beam steering and interference fringes. Beam steering tests are much easier but require lots of room.

12.5.1 Stability

The first requirement of any adjustment procedure is that it be *stable*. As the iteration proceeds, the corrections at each pass must be steadily decreasing. Stability can be achieved by performing measurements at each stage, or by choosing a design and a procedure such that the alignment is inherently stable. (Towing a trailer is inherently stable, but backing one up requires constant corrections to avoid failure.)

In cases where the adjustment is not iterative, for example, progressively adding components to an optical bench setup, we require in addition that our alignment criterion remain valid at all times during the assembly; it is not much help to put a HeNe alignment beam down the optical axis if the first lens is made of germanium, or to allow errors in placing the earlier elements to steer the beam completely out of the system by the time the later elements are added; the assembly procedure as well as the alignment procedure must be stable.

Since most optical adjustments are nonlinear in character, the easiest way to ensure stability is to make sure that the interaction between controls is small, and then proceed in the right order.

12.5.2 Orthogonality

Two adjustments are said to be *orthogonal* if changing one leaves the other unaffected; for example, x and y on a good translation stage, or translations in the object and translations in the pupil (i.e., position and angle). Orthogonality is the Holy Grail of the manufacturing department: set each screw once and forget it. An example is the beam steering system of Figure 11.2. It isn't always easy to achieve, but really work at it before giving up; a stable alignment strategy will prevent errors from growing, yet if the controls interact a lot your alignment algorithm may never converge. Consider aiming a beam at a spot on the wall using nonorthogonal mirror adjustments: if moving the spot 10 cm to the right needs 1-turn adjustment of one screw, then any accompanying up and down motion should require no more than 1/10 turn of the other screw to correct it. It gets a lot worse when there are fewer readouts than adjustments; and it's worst when there's only one, as in maximizing the output power of a laser.

A general rule for such nonlinear 1D adjustments is that an interaction of 1/10 turn between two knobs is pretty manageable, but 3/10 requires an expert and 5/10 is impossible to do by hand unless the system is nearly aligned to start with. This rule applies

in both directions, so that it does not help to put one adjustment on an extremely fine thread. It's a bit easier in two or more dimensions, but the point remains: avoid interacting controls if at all possible.

12.5.3 Use Serendipitous Information

As we get more experienced with a system, we develop powerful intuition about it—we get faster at adjusting it, and we know what's wrong when it breaks. One of the main sources of this intuition is noticing the information that can be gained from unintended places. For instance, focusing a laser microscope gets a lot easier when there's an incidental beam somewhere whose diameter changes with focus; put a card there, and trace the beam outline on it when the focus is perfect. Forever afterwards, rough focusing will be a piece of cake. Diffraction patterns in stray return beams are similarly useful for finding features on a substrate. The colors of the white-light fringes in an interferometer are another good pointer. Be on the lookout for these in your systems; they'll save you lots of time and frustration.

12.6 COLLIMATING BEAMS

Collimated beams are very easy to work with, which makes them popular. Collimation is particularly easy to test for, so you can be sure when you're there. You can change the spacing between elements freely (at least those with no optical power), put in polarizing cubes, and so on, without causing any change in the system aberrations—at most, a slight change in magnification. They have several nasty features, however, such as maximal sensitivity to etalon fringes and increased danger of eye damage if the beam is powerful.

How to collimate a beam depends on what it looks like to begin with. A beam from an extended source such as a light bulb cannot be collimated well, because its spatial coherence is not good enough. Here we specialize to laser beams, and so assume that the beam is monochromatic and fully coherent.

A collimated laser beam is one whose phase fronts are perfectly planar, so any transverse slice of the beam will have a constant envelope phase. A simple calculus of variations argument shows that this minimizes the equivalent width of the far-field pattern for a given intensity distribution. Thus collimation can be measured in three general ways: directly, by looking at the behavior of the beam as it propagates; in the far field, by measuring its optical Fourier transform; or in the wave picture, by measuring the phase fronts with an interferometer. Which of these is best depends mainly on the size of the beam.

12.6.1 Direct Collimation

A beam N wavelengths across stays collimated for N beam diameters. Direct collimation is done in the obvious way: stick a mirror in the beam to shine it on the back wall of the lab. Walk back and forth, using an iris diaphragm (best) or a sheet of paper with a circle drawn on it to gauge when the distant beam is the same diameter as that exiting the collimating lens. With a wall 3 meters away, this procedure will provide diffraction-limited results for visible beams of <2 mm diameter, but gets rapidly less accurate as the beam gets broader.

You can mount a corner cube or a mirror on the wall to get twice the path length with much less travel. Note that this is susceptible to strong $1/f$ noise from people occasionally walking through the beam. When you do this, make sure you avoid sending the beam where it may encounter someone's head. Optical tables are usually at just the right height for escaping beams to hit the eyes of someone sitting down—like you, for instance. If you need to do it a lot, a piece of 3 inch plastic drain pipe with a corner cube at the other end can stretch all along one wall of the lab with perfect safety. (Leave an opening at the far end so you can align it without peering down the tube.)

12.6.2 Fizeau Wedges

A Fizeau wedge is a very narrow triangular air gap between two optical flats with a thin spacer on one side only. Under narrowband illumination, it produces a system of parallel fringes that are easily visible. Fizeau wedges are good for gross collimation, because a collimated beam produces straight fringes of equal spacing, but for fine work, or in the presence of aberrations, Fizeau wedges are nearly useless. The strong asymmetry of the pattern, and the necessity of spotting, say, a $1/10$ cycle irregularity in spacing or warping of a fringe, make it impossible to do by eye. Using a CCD camera to look at the fringes would have a much better chance, since the lithographically defined sampling grid of the CCD will show up deviations from the predicted peak positions very clearly. (Cast the fringes directly on the CCD, or at least calibrate the geometric distortion of the lens if you do this.)

12.6.3 Shear Plates

A shear plate is a sort of Fourier transform of a Fizeau wedge: instead of interfering two waves offset in angle, it uses two offset in space, for example, by reflection from the front and rear surfaces of a tilted flat plate; the phase of the fringes is thus a finite difference of the phase front of the incoming beam. A collimated beam produces a uniform intensity that goes up and down as the path difference varies, slight defocus changes that to a system of parallel lines, and more complicated errors to a more complicated pattern. Shear plates are insensitive to phase curvature orthogonal to the shear direction, so you need two plates for a general beam.

12.6.4 Collimeter

The late lamented Blue Sky Collimeter was a scanned shear plate interferometer. They don't make it any more, but the function isn't that hard to reproduce. The relative phase of the sheared beams is scanned by turning the plate slightly, using a flexure powered by a bit of Nitinol wire (a shape-memory alloy, sold as "Muscle Wire" by robot hobby suppliers). This makes the fringes move back and forth at about 0.2 Hz, which makes them easy to see, and easy to tell apart from the envelope of the beam intensity. When the beam blinks on and off rather than exhibiting spatial fringes, it is collimated to high accuracy ($\lambda/50$ or even better). These devices are most useful with good quality beams, where it is obvious what to adjust to fix the problem. Heavily aberrated beams (e.g., astigmatic diode lasers) produce complicated fringe patterns with odd motions, which makes adjustment a more subtle business. A slower, wider range motion would be better, and if you're doing many, you really need a proper measuring interferometer.

12.7 FOCUSING

12.7.1 Autocollimation

The autocollimation condition is useful for focusing as well as for alignment. By symmetry, if there's a surface at the focus of a beam, aligned normal to the beam axis, the beam will retrace its path. Not only will the two beams be coaxial, but if the incident beam is collimated somewhere, the reflected one will be too. This reduces the focusing problem to collimation of the reflected beam, but good accuracy requires the incoming beam to be collimated very well, since any collimation error turns into a focus error. This method is good for quick coarse focusing if you know there's no vignetting, since the two beam diameters will be identical in focus.

12.7.2 Direct Viewing

Even if you don't use the 35 mm camera trick, a viewer can make focusing a snap. We talked about this in Section 11.8, but it's worth reemphasizing. Use a microscope trinocular head, a couple of eyepieces, and a beamsplitter to form an image of your focal plane, put a target there, and adjust the focus of the beam until the spot on the target reaches its minimum diameter. If the viewer has a higher NA than the beam, you can spot beam defects this way too. You can do this in reflection, right through the objective lens and back again, so that no additional high-NA optics are required. Be sure the beam power at the eyepieces is below $10 \mu\text{W}$.

12.7.3 Foucault Knife Edge

The Foucault knife edge test is a good way of locating an aerial focus. As Figure 12.3 shows, it's very simple; put a razor blade on a translation stage and move it across the beam while looking at the transmitted light. If the razor is beyond the focus, the shadow moves the same way as the blade, and if it's inside the focus, it moves the other way (this is obvious from geometric optics). What may be a bit less obvious is that when the blade is right at focus, the transmitted beam just fades out and blinks off, with no obvious shadow motion. This is not so for an aberrated beam, and in fact the knife edge test is a sensitive qualitative test for aberrations. Near focus, an astigmatic beam either gets chopped off from the inside out, or the cut edge goes left to right, swings round to up and down, and then goes to right to left, depending on how astigmatic it is. If the knife edge pattern is funky-looking, you've got at least a wave of aberration. For smallish aberrations, the pattern gets steadily better looking as the aberration decreases, but if you have more than a couple of waves of astigmatism or coma, it takes a trained eye to know what to adjust.

12.7.4 Intensity-Based Chopping Tests

There are beam scanning instruments available that use a slit in a rotating drum to produce a 1D beam profile. Alternatively, a rotating chopper will give you the integral of the 1D beam profile (don't use a vibrating chopper, because the blade velocity is poorly controlled). These are great for finding the beam waist, but not much use for diagnostics; unless the beam is invisible, or the NA is high, a viewer is better and faster. Their major




Figure 12.3. Knife edge test.

advantage over a direct viewer or a CCD camera is that they produce a trace for an oscilloscope, and they can work in the IR and UV. The other thing is that a chopper can be used as a reasonably accurate focus finder for good quality beams; with a split detector or quad cell in the beam, the output will be nearly 0 V all the time if the chopper is at focus, a positive pulse waveform for one side of focus, and a negative pulse for the other. It isn't perfect because the chopper blade is comparatively thick, so that the beam can bounce off the side as well as the top.

12.7.5 Diffraction Focusing

The focused image is the Fourier transform of the pupil distribution, so one good way of making sure your imaging detector (e.g., CCD) is at the focus of your optical system is to put a strong diffractor at the pupil and adjust the detector position until the Fraunhofer pattern is in sharp focus. The usual method is to put a wire or a piece of tape across your pupil, so that the Fraunhofer pattern exhibits a sharp line running normal to the wire. Adjust the CCD position until the pattern is a single sharp line (not two closely spaced lines), and you're in focus. This works best with long-focus systems focused on infinity, where the pupil position is not especially critical (it was first developed for astronomical telescopes).

12.7.6 Speckle Focusing

Unless it's very tightly focused, a laser beam bouncing off a rough sample (e.g., in a microscope) produces speckles in the scattered light. These are usually thought of as a noisy nuisance but (like other kinds of noise) can often be turned to advantage, as in

speckle interferometry. One thing they're good for is focusing. If you shine the speckles on a card while moving the sample from side to side, they move opposite ways for samples on opposite sides of focus, and when the sample is in focus, they just blink on and off like sunlight through a tree's canopy.

12.7.7 Focusing Imagers

One good and very fast way to focus an image sensor is to use a point source at the correct distance, bright enough to cause blooming. Adjust the focus to maximize the length of the blooming tail in the image. This works because blooming is caused by charge overflowing from the CCD well, and the more light is concentrated on a small area, the more charge will overflow to cause blooming. This of course assumes that the CCD fill factor is high, or that the image of the point source is at least a few pixels in size, because otherwise if you have an unresolved point source that happens to be centered on an insensitive area, the blooming will *decrease* near focus instead.

If you have a repeatable, calibrated focus actuator, another approach is to take images at known focus positions, compute the contrast, run a low-order polynomial through the data points, and interpolate to get the position of best focus. The contrast can be computed as the sum over the image of the RMS values of the first finite differences in x and y . There are lots of variations on this; for an astronomical imager, you can look at the diameter of star images as a function of focus.

12.7.8 Standards

Film camera manufacturers have standardized the distance from the mounting flange to the film plane; each manufacturer has its own standard, unfortunately, but at least this has to be constant for all lenses in their lines. (Why?) If you have the matching camera, the film rails give you a fixed reference plane, and you can measure the focal position with calipers. That way, you can use dead reckoning for focus, at least with a sufficiently slow lens.

12.7.9 Subapertures

You can put a mask on top of your lens, consisting of two holes near opposite edges of the pupil, and adjust the focus until the focused spots coincide. This is a bit like a rangefinder camera and works very well. The hole diameters should be about one-quarter of the clear aperture so that the spots aren't too big.

12.8 ALIGNING BEAMS WITH OTHER BEAMS

It is frequently necessary to align two beams accurately with one another, propagating either in the same direction or opposite directions. For the copropagating case, we may want to wavelength-division multiplex a fiber link, or to pass two beams of slightly different frequency down exactly the same optical path so that any vibration or scan motion causes only common-mode errors. Counterpropagating examples include lining up your beam with a guide beam provided as an alignment aid, as in the Burleigh WaveMeter, or two-photon Doppler-free spectroscopy, where the resolution improvement

requires the beams to be well aligned. How difficult this is really depends on whether it is enough for the beams to be aligned with each other, or whether they must in addition be aligned accurately with the rest of the optical system.

Check the overlap with a translucent screen, such as an index card or sticky note; by putting it in the beams, and then blocking and unblocking the brighter of the two, you can see the overlap very conveniently in both the co- and counterpropagating cases.

12.8.1 Copropagating Beams

Getting two beams to propagate together requires some sort of beam combiner, which is a beamsplitter run backwards, as shown in Figure 12.4a. Put the combiner on a two-axis tilt stage for one adjustment, and use a mirror on a two-axis tilt mount for the other. For applications where back-reflections are a serious concern, use a Wollaston prism for




Figure 12.4. Aligning beams to other beams: three techniques.

the beam combiner.[†] Use the mirror to superimpose the beams on the combining surface and then tilt the combiner to align them in angle. Commercial mirror mounts and prism tables are relatively poor for this, because the adjustments are too tweaky, so consider using a stiff flexure for the combiner instead.

This is of course not a complete solution to the problem, because we often want to align two beams of very different wavelength, for example, a HeNe guide beam and a CO₂ cutting laser. You may have to use something a bit more special for the beam combiner, such as a hot or cold mirror, which passes one wavelength while reflecting the other. Watch out for subtle chromatic effects downstream of your combiner, especially polarization funnies.

Narrow beams can be aligned well by eye. A rule of thumb is that two beams of diameter d can be made coaxial to within the diffraction limit by eye on the first try, as long as the distance L between the two mirrors obeys

$$L > \frac{d^2}{5\lambda}. \quad (12.2)$$

This estimate takes account of the visual difficulty of aligning the centers of two smooth luminous spots, so variations among people limit its accuracy to $\pm 20\%$. You can make do with less space if the beams have strong features (e.g., diffraction rings or sharp edges) that can be used as visual cues. The main things to avoid are putting one beam steering element on a four-axis stage, and using two adjustments that are nearly degenerate, e.g. two tilting mirrors spaced very close together.

If the beams are too large for this, you can use a lens to Fourier transform the beams. Examine the focused spots with a magnifier or a low power microscope to check that the focused spots are really superimposed. Because the focused spot gets smaller as the beam gets broader, this technique takes care of the broad-beam case very well.

12.8.2 Constrained Beam-to-Beam Alignment

A slightly more difficult case is where the rest of the optical system provides more than one axis, which the beams must follow (Figure 12.4b shows the counterpropagating case), where, e.g. a laser and wavemeter each defines an axis, and we don't want to mess with the system alignment just to use the wavemeter. The figure shows one angular and one xy adjustment, which is not orthogonal but does converge rapidly because the xy alignment doesn't mess up the angular alignment.

If the beam alignment is critical, combine the beams as early as possible and pipe the results down the system so that subsequent pointing drifts affect both beams identically. An example is a common-path interferometer, where errors in the relative position of the beams cause an enormous degradation of the vibration immunity of the system. An orthogonal approach is to use the beam alignment jig of Figure 11.2, translating the beam and the jig together to make the beams overlap at the exit pupil of the jig, and then work the angular adjustment until they are perfectly aligned.

[†]It isn't so easy to use a Wollaston, because you can't just tilt it to change the alignment. You'll need another mirror.

12.8.3 Counterpropagating Beams

Aligning two counterpropagating beams is the easiest of all if you have lots of space, and the hardest otherwise. The best way to do it is to make both of them adjustable in angle, e.g. by bouncing them off mirrors on two axis mounts. Adjust Mirror 1 to overlap the beams on Mirror 2, then use Mirror 2 to overlap them on Mirror 1. (Figure 12.4c shows the same idea with fiber collimators.) These two adjustments are orthogonal, so the two beams should now be in alignment. If it is necessary to iterate, the mounts can be adjusted independently without interacting, which makes the alignment fast and convenient. Sometimes one beam cannot be moved, e.g. the one defining the optical axis of the system. Then the second best thing is to use one translation stage and one tilt stage. Adjusting the translator to overlap the beams on the tilting mirror and then adjusting the tilt to align the beams in angle is a workable two-step procedure with no interactions. If you need to iterate, it is no longer orthogonal. Every time you touch the translation stage, you have to readjust the tilt. The iterations are always stable, because the tilt does not misalign the translation, so that the iterations must terminate.

Both of these schemes work best when there's lots of room for angular misalignments to turn into spatial errors. If you have a single-mode fiber collimator with a beam coming out, a 3 meter spacing with a 3 mm HeNe beam will allow you to send another beam back into that fiber in one try. It gets harder from there down.

If there isn't enough space for the easy way, you have to use the harder ones, which rely on making the beams copropagating and using the tests of the previous section. You can do this in two ways: first, if it's okay to block the beam during alignment, you can use a beamsplitter and an accurate corner cube retroreflector, as in Figure 12.5. Second, if you need the beams unblocked during the test, you can put in a weak reflector (e.g., a microscope slide) and adjust it for autocollimation of the first beam. Leaving the first one fixed, align the second using the constrained copropagating technique. The first technique gives you the autocollimation free, which is considerably more convenient.

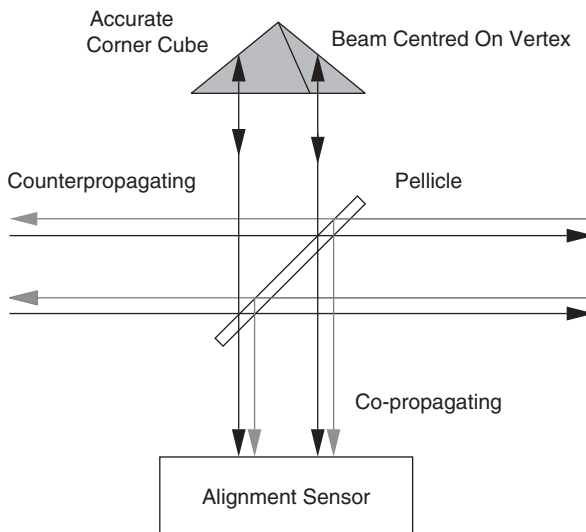


Figure 12.5. Alignment aid: corner cube and beamsplitter.

12.9 ADVANCED TWEAKING

12.9.1 Interferometers and Back-Reflections

Lasers are amazingly sensitive to back-reflections into their cavities. Michelson and Fabry–Perot interferometers are very bad for this, because they rely on back-reflection in order to operate. Mach–Zehnders are much better in this regard. Before aligning your system so that it just exactly blows your laser out of the water, it's important to think about what the results of perfect alignment will be. A lot of problems are reduced by strategic minor misalignment.

12.9.2 Backlash and Stick–Slip

Anyone who has ever successfully aligned a complicated laser system has educated fingers. There may be 20 knobs to twist, and the correct alignment setting is in there somewhere, but where? The problem is made worse by the usual tendency of the adjustments to exhibit twin evils: *backlash* and *stick–slip*. Backlash is like a loose tap handle in the shower: the same water temperature corresponds to different knob positions, depending on whether you're going clockwise or counterclockwise. Stick–slip is what keeps all the bolts in your car from falling out, but it's less useful in adjustment screws: when you twist the knob, it doesn't move smoothly, but goes in jerks, usually shaking all the other screws out of adjustment as it goes. Surfaces stick and slip because the coefficient of static friction is higher than the coefficient of dynamic friction. Energy stored in elastic deformation can be released as motion, causing instability over time—the San Andreas fault is a well-known example, but it happens in translation stages too.

The net effect is that even when you've located the region where the sweet spot lies, getting the alignment just right is an iterative procedure requiring patience.

Example 12.1: Aligning a Regenerative Amplifier. At this writing, the author's fanciest toy is a tunable picosecond laser source, producing 20 ps pulses continuously tunable from 420 nm to 10 μm , with a small hole around 710 nm. It's pumped by a frequency-tripled, flashlamp-pumped Nd:YAG laser that has both active (Pockels cell) and passive (dye) mode locking. It's a thing of great beauty but requires a lot of laser jockeying to keep it working. Just the pump laser has three separate cavities. The most frequent adjustment is the alignment of the regenerative preamplifier section, which takes 5 μJ pulses and amplifies them to the 500 μJ that the power amplifier section expects.

The mirror adjustment has two stages per axis: ordinary ball screws for coarse setting, and a vernier consisting of a fine screw acting on an inclined plane, forcing the anvil opposite the screws in and out slightly. This vernier is a poor design that has a lot of stick–slip energy storage, requiring a lot of babying. To align a system like that, you have to do all your tweaking in the screw-tightening direction. The reason is that the screw is so much stiffer than the return spring that the same energy release corresponds to a far smaller motion, and hence much less instability. During tweaking, this also makes the X axis adjustment stay still while you're tweaking Y , which is a huge help.

Tweaking in one direction requires a bit of memory and judgment: for each axis in turn, slack the screw off until you're well out of alignment in that direction. Slowly tighten it until it starts to lase, and watch the energy meter to see what the highest value is. You have to tune past it to find this out, so on the second pass, tweak it near the peak position and then very slowly tighten until it reaches peak output, and stop. You

need to remove your hand from the knob after each small motion, so this may take a few tries. On this particular laser system, the backlash is wider than the whole range in which lasing occurs, so it really matters.

12.9.3 Adding Verniers

If you have a system with a vernier (coarse–fine) adjustment, you may find that there are places where the coarse adjustment isn't stable enough for the vernier to be used. In that case, move the coarse adjustment somewhere nearby and try again. Verniers are far superior to just putting a fine screw on a coarse adjustment—if you're designing the system, make the vernier separate, for example, a very weak lens that you can shift around.

12.9.4 Cavities with Obstructions

When you're aligning an optical cavity containing obstructions such as pinholes, it is sometimes better to align the mirrors without the pinhole, then replace the pinhole. An obstruction produces such strong slopes on the merit function (e.g., power output) that the weaker ones due to the cavity alignment are often obscured. If you can't take out the pinhole, make sure you measure the sensitivity of the merit function versus pinhole position: align the mirrors, measure, misalign the pinhole slightly upwards, realign, measure, move the pinhole slightly down from its original position, realign, measure. Then you'll know whether to move the pinhole up or down to approach the optimum. Repeat in the horizontal plane and iterate to convergence. This is much more time-consuming than just taking out the pinhole. Mark the table so it's easy to put the pinhole back where you got it.

12.9.5 Aligning Two-Beam Interferometers

There are two kinds of interferometers: two-beam ones, such as Michelson and Mach–Zehnder types (see Section 1.6), and multibeam ones such as Fabry–Perot and Fizeau types. It is seldom enough to get the beams propagating in the same direction. They must also be in the same state of collimation and (often) have zero path difference between them. To do this, use the fringes between the two beams as a long-path shear plate.

Zero path difference is easily found using the white-light fringes. Thermal sources have coherence functions with a single peak (see Section 2.5.4), so fringes appear only near zero path difference. Make sure you cancel out the dispersion as in Figure 1.10: use a compensation plate in a Michelson and put the beamsplitters the right way round in a Mach–Zehnder. Otherwise, the different paths in glass produce dispersion that will wash out the white-light fringes and generally make life difficult. You can tell when you're approaching the white-light fringes by looking at the pattern through an interference filter—narrowing the bandwidth of the detection has the same effect on the fringe contrast as narrowing the bandwidth of the source. Thus a 50 nm bandwidth filter will allow you to see fringes several times further away from ZPD, which makes alignment easier. Use a single-point detector and an oscilloscope to measure fringe contrast for interferometer fine alignment, especially in the IR, where it's inconvenient to use the fringe pattern directly. It gives you a single parameter to optimize, which is easier to do when you

get near the end. Mild mechanical vibration (e.g., gently pounding the table with your fist) will often give you moving fringes, which can be detected and displayed on an oscilloscope. Make sure you measure the photocurrent of both beams separately, and calculate the maximum contrast in advance using (12.3)—otherwise you don't know when you're there. The maximum fringe contrast available is

$$C_{\max} = \frac{2\sqrt{I_{\text{sig}} I_{\text{LO}}}}{I_{\text{sig}} + I_{\text{LO}}}, \quad (12.3)$$

which is 100% when $I_{\text{sig}} = I_{\text{LO}}$, but only 2% when $I_{\text{sig}} = 10^{-4} I_{\text{LO}}$.

A contrast of 100% makes the beams blink on and off, and it's a pretty sight to see an interferometer that good. Nonetheless, do remember that beams blinking on and off don't necessarily mean that your beams are in the *right* state of collimation, just the *same* state (much time has been wasted by confusing the two).

12.9.6 Measuring Focal Lengths

This effect can be very useful in precision focal length measurements—throw together a Michelson interferometer using an expanded laser beam, put a lens in one arm, and move the lens axially until the fringes blink on and off with no rings. The lens is now exactly one focal length from the mirror. (You'll want to put the lens in its cell first—otherwise the precision will be lost as soon as you take it out of the setup.)

12.9.7 Aligning Fabry–Perot Interferometers

There are two main ways of aligning plane mirror Fabry–Perots: either by using a laser (preferably one whose wavelength is on the skirts of the HR mirror coating) and make all the beams overlap, or put a large area line source behind it (e.g., a green mercury lamp) and make the ring spacing stay the same as you move your head from side to side. (You have to be looking from the right distance so that the fringes have a reasonable width.) Both of these are pretty sensitive and will get you into the ballpark very rapidly. Fine alignment can be done by maximizing the transmitted intensity of an expanded HeNe beam or some other appropriately isolated laser source.

12.9.8 Aligning Lasers

A laser is a special case of a Fabry–Perot, since no light comes out until it's aligned. After it starts to lase, fine adjustment is a matter of using an analog power meter and fiddling until it reaches a maximum; it's the initial coarse alignment that's the trick.

Since a laser is a Fabry–Perot after all, you can align it the same way, by using an auxiliary laser and aligning the resulting beams. Sometimes you have to do that, for example, when the laser is first manufactured and the resonator isn't focused, but usually it can be avoided in the lab.

There's another technique that is quicker and much more fun: raster-scanning the alignment with a screwdriver. At least one mirror will be sitting on a spring-loaded mount with adjustment screws. Slack off one screw all the way, and jam a pry bar (e.g., a big screwdriver) in next to it. Crank that one quickly back and forth through its range while slowly moving the other screw, starting from one end of its range. You'll get bright

flashes when the slow adjustment is roughly right, so then you pull out the screwdriver and twist the fast knob until you get steady output. This quick method is great for *walking in* a gas laser after changing lines or cleaning the Brewster windows. Even though it's what laser service technicians do, it's a bit shocking to those who haven't seen it done before—almost as if you were using a ball-peen hammer and a tire iron—which makes it even more fun.

12.9.9 Aligning Spatial Filters

One of the reasons spatial filters aren't used more often is that they are fiddly to align if you don't do it properly. There are two phases: finding the beam and finding the sweet spot. Finding the beam is done the same way you align a laser but is a bit easier since the bright light is there all the time, rather than only when you're nearly aligned. If your pinhole is in a shiny surface, the easiest way to do rough alignment is to put in a viewer, even just pellicle or a microscope cover glass at 45° in the incoming beam, shining the reflected light onto a card. Defocus far enough that you can see the image of the pinhole on the card, then zero in by translating the pinhole until it's centered, adjusting focus, centering again, and so on until the pinhole nearly fills the image. By this point, light should be coming through your pinhole. If there is too much interaction between focus and lateral adjustments, your beam is misaligned. If this is not on purpose (remember that slight misalignment is a good idea when using a laser), adjust the lens laterally while racking focus in and out to eliminate the lateral motion of the pinhole on the card.

Fine alignment depends on whether your pinhole is resolved or not. If it isn't, just adjust the pinhole while looking at the detected photocurrent on a meter (analog ones make this a great deal faster). Find the maximum and you're done.

Wider pinholes or slits are a bit subtler. If you don't know which way to go to find focus, look at the far-field pattern of the beam (transmitted or reflected). Passing through focus reverses the direction in which the image of the pinhole moves when the object moves (this is obvious from geometric optics, since all the rays pass through the focus). In the transmitted beam, a pinhole too close to the lens makes an image that moves in the opposite direction. The directions are reversed in the reflected image.

12.9.10 Use Corner Cubes and Pentaprisms

We typically do alignment by overlapping spots in the far field and near field or, alternatively, by looking for interference fringes. Overlapping beams is a lot easier, especially for setup, but requires two measurements per iteration and needs lots of space or auxiliary aids. The most useful alignment aids are those based on accurate corner cubes and pentaprisms, which allow really good angular alignment of counterpropagating beams. You'll need an independent method for aligning them in position. If your cube has sharp enough corners, and your beam is big enough, you can bounce it right off the vertex (grey lines), in which case all is easy; a quad cell will tell you when they're coaxial.

If you can't do that, you'll have to let the returning beam be offset and align the two beams accurately parallel, using a ruler or (for very fine alignment) two quad cells spaced appropriately. You can do the same sort of thing with other constant-deviation prisms (e.g., pentaprisms).

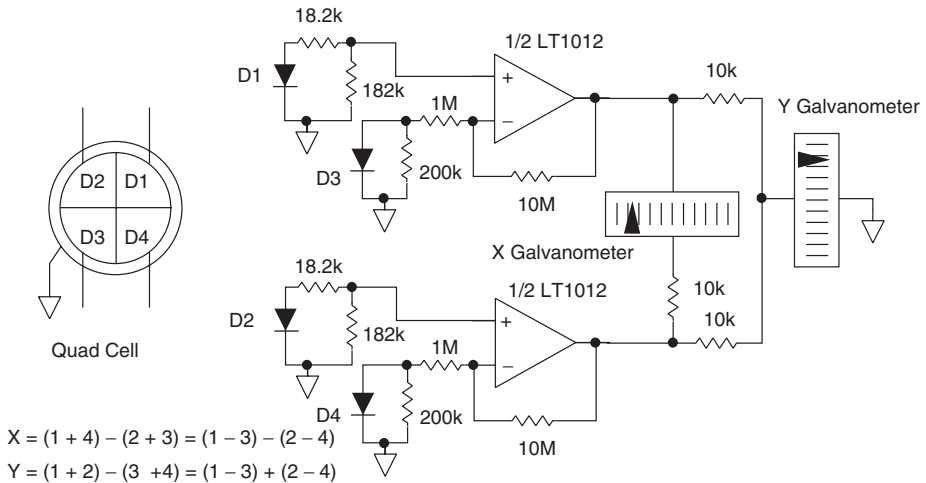


Figure 12.6. Quad cell alignment aid.

12.9.11 Use Quad Cells for XY Alignment

Another approach to aligning copropagating beams is to use a quadrant photodiode and a pair of analog voltmeters, as shown in Figure 12.6. Put a large ($22\text{ M}\Omega$) resistor from each segment to the common terminal, and look at the combinations $(1 + 2) - (3 + 4)$ for the X direction and $(1 + 4) - (2 + 3)$ for Y . You can do all this with a dual op amp and a few resistors, as shown in Figure 12.6 (the trick is to notice that the X and Y signals are given by $(1 - 3) \pm (2 - 4)$). Loading of the open circuit photodiodes by the resistors will start to dominate at low photocurrents, reducing the signal range at the meters. A fancier version, using two high input impedance differential amps (maybe implemented with a quad op amp IC) instead of resistors and a dual op amp, will extend the logarithmic range to photocurrents of $5\ \mu\text{A}$ and below without hurting the strong signal performance.

The reason this works so well is that the open circuit voltage V_{OC} of a photodiode is proportional to the logarithm of the optical power incident on it (see Section 14.6.1). For matched diodes, subtracting the logarithmic signals gives the logarithm of the ratio of the powers on each diode pair, and that ratio is independent of overall power. Thus we have a very simple meter that measures beam position directly in units of the beam diameter (see Section 3.5.7) and can be built from scratch in a couple of hours. If you carefully center a biggish cell (10 mm or so) in an aluminum slug the same diameter as the lenses you use most, you can pop it in anywhere in your setup, any time. It works especially well with optical Erector sets such as Microbench.

Germanium quad cells are a good choice if you work in the visible and NIR at reasonable power levels ($>100\ \mu\text{W}$). At very low power levels, the shunt resistance of the germanium may cause the sensor signal to become power dependent. Silicon is better then. The main objection against Ge is leakage, and who cares about that when we're forward biasing them anyway?

12.9.12 Use Fringes for Angular Alignment

Another way to get accurate angular alignment between laser beams is to force them to interfere perfectly. If the beam blinks on and off across the whole diameter when you

gently press one mirror, it's aligned as well as it's going to get. A viewer helps a lot with this, and to get absolute perfection you can use a big photodiode, a load resistor, and an oscilloscope; adjust for maximum amplitude as you pound gently on the table.

12.10 ALIGNING LASER SYSTEMS

12.10.1 Define an Axis

In most cases, the best thing to do is to define the optical axis with a HeNe laser. Spend a little time getting the laser aligned properly, really horizontal and really square to the table. If you're using commercial mounts, a height of around 150 mm above the table is a good choice; low enough to avoid floppiness, but high enough to get mounts and stages underneath the beam (a common error is to put the beam too close to the table and then have to suspend later elements from above like freeway signs).

Figure 12.7 shows one good way to align the axis: use a jig consisting of a square piece of sheet metal with one side bent down (or a bar attached) to follow the edge of the table, and an optical mount with an iris diaphragm mounted near the opposite side. Slide the jig up and down the table edge, making the beam go through the iris at all points along its path; this guarantees that the beam is parallel to the edge of the table. An even better method is to substitute the quad cell device of Section 12.9.11 for the iris, which improves the accuracy by two orders of magnitude at least, assuming the alignment beam is clean enough.

Surveyors define the horizontal reference with an autocollimation system consisting of a pentaprism and a mercury bath; shining a HeNe down onto the mercury and adjusting it for autocollimation makes a vertical reference, and an accurate pentaprism turns that into a horizontal reference that does not depend on mild wiggles. In the lab we normally care more about being parallel to the table.

12.10.2 Adding Elements

Start from the far end of the system, and add elements one at a time, making sure the front surface reflections from both sides of all elements go back on themselves (this is not subtle). By the time you get back to the laser, you're all aligned.

This procedure may need to be modified if the optical path is extensively folded; clearly you can't use just the sample and laser if the beam path needs to look like a W.




Figure 12.7. Defining an axis.

In that case, set up the folding mirrors and sample, paying close attention to the alignment of each individual section just as before.

This technique works well for beams that are collimated or nearly so; if you have sections where the NA is appreciable (say, 0.05 or greater), build those sections on their own base plates and test them individually.

12.10.3 Marking Lens Elements

You can mark the center of a lens with a small dot of India ink or other easily removed, easily visible material. This helps a good deal with getting the mounts adjusted, or with aligning a previously adjusted subassembly with the system axis. The mark can be positioned by autocollimation testing, or with a lathe or other jig if the lens is sufficiently accurately edged, that is, the axis of the lens coincides with the center of its outline.

12.10.4 Lenses Are Easier than Mirrors, Especially Off-axis Aspheres

Lenses have the nice property that if you shine a laser beam down the axis, the transmitted and reflected rays are at exactly 180° to each other, and the front and back surface reflections overlap completely. Shifting the lens makes the two reflections go opposite ways, and tipping it makes them go the same way. Thus a lens gives you independent cues for position and angle. According to Kingslake,[†] “surface tilt does more damage to an image than any other manufacturing error, and in assembling a lens it is essential to avoid tilting the surfaces at any cost.” He suggests limiting accidental surface tilt to less than 1 arc minute, which is easily done using the front and rear surface reflections.

Mirrors have only one surface, so you get only one cue. That isn't too serious with a spherical mirror, of course, because (apart from a bit of vignetting) all points are equivalent; translations and tilts have the same optical effect, so just shine a laser beam down the desired axis and adjust for exact back-reflection.

Aspheric mirrors are much less forgiving in this regard, especially fast ones. It is usually necessary to mark the center of such a mirror somehow, and this is usually done during fabrication. Failing that, you'll probably have to set up a test bench to look at the interference fringes, align it for best performance, and mark it yourself.

Off-axis aspheres are the worst of the lot; since the axis generally lies outside the element, you can't mark it and can't bounce a laser beam off the axial position either. If you're planning to use one of these, make sure you also plan how you're going to align it—you'll probably wind up using the rest of the optical system plus some auxiliary apparatus.

12.10.5 Use an Oscilloscope

Extremely fine alignment (e.g., for an interferometer system) is best done with an electronic sensor of some sort (e.g., a quad cell or a lateral effect cell) and an analog oscilloscope in X-Y mode.[‡] Stick the sensor in the beam, and adjust the beam pointing until the oscilloscope dot is centered in the screen; then jack up the vertical gains and do it again. Obviously you have to use a sensor where 0.000 volts means the light is dead

[†]Rudolf Kingslake, *Lens Design Fundamentals*. Academic, Orlando, 1978, p. 192.

[‡]Turn the display intensity way down so as not to burn the phosphor.

center, and adjust the offsets on the scope so that 0.000 volt on x and y is in the center of the screen. Doing it this way is much easier on the eyes, and it works very repeatably.

12.11 ADHESIVES

12.11.1 Structural Adhesives

Adhesives are very convenient. Need to hold something in an odd position? A drop of Krazy Glue or Duco Cement or a bit of double-sticky tape and you're on your way, right? Actually, adhesives aren't as useful for optical setups as one might think. Glue is dimensionally unstable, first of all. It shrinks a percent or two on curing and keeps creeping slowly with time. It conducts heat very slowly and may delaminate under wide temperature excursions. It also requires extremely careful surface prep for best results—you have to remove oil, fingerprints, and dust, plus roughening the surface everywhere. Scraping it with a scalpel won't do it; the surface area increase is too small. Sandblasting or bead blasting is good, sanding second best. Another problem is outgassing. There exist low-outgassing RTVs such as GE 566 or GE 142; cyanoacrylate cements such as Permabond 910 or Krazy Glue outgas enough to leave white frost on nearby surfaces. Manufacturers laughingly advertise some types as “low outgassing,” but that's true only by comparison. Two-part epoxy is usually best, with an hour's baking at 100 °C to toughen it up.

12.11.2 Optical Adhesives: UV Epoxy

One way to simplify the mounting problem, and cut etalon fringes to boot, is to cement your optical elements together. For example, you can carefully cement a $\lambda/4$ plate to a polarizing prism to make a transmit/receive beam separator of the kind we saw in Section 6.10.9.

Find a UV epoxy you like with an index of 1.52 (e.g., Norland 65). This greatly reduces reflections from gluing two pieces of BK7 together, which you'll do a lot. If you're using other glasses, pick a glue whose index is midway between them. It is not unknown for the very weak Fresnel reflections from an epoxy joint to cause etalon fringes bad enough to wreck a spectroscopic measurement, so don't use glued joints the way a witch doctor uses a fetish: calculate the worst case reflection over temperature and design around it.

UV epoxy has only a few months' shelf life at room temperature, but most kinds last for some years in the fridge. Get a good quality two-bulb long-wave UV lamp (black light). If you get one of the enclosed ones with windows, you can rest small parts on the upturned lamp, which is very convenient. Really small parts (<15 mm diameter) can be cured in an EPROM eraser. If you're going to be doing it much, it's worth getting one of the nice little hand-held UV guns Norland sells for tacking the surfaces together, because otherwise they will tend to slide around while they're curing.

UV epoxy is easy to use if you're neat with your hands. Use powder-free finger cots or disposable vinyl gloves while cleaning surfaces and applying glue. The glue is best applied by loading it into a disposable syringe with a filter attached. The epoxy flows freely through filters smaller than 0.5 micron, even down to 0.025 micron, so you can ensure that there are no significant bits of crud in your glue (the crud is usually bits of partly cured epoxy from the spout of the bottle). Get the surfaces good and clean. Using

Opticlean works well, since you can dry-fit the surfaces once they're coated with the polymer film, handle them with your fingers, and so on.

Put one element facing upwards in a soft jawed vice, or wax it to a stable surface. Peel the film from the fixed element, and put on a drop of UV epoxy. If the surfaces are flat or nearly so, put it near one edge. Touch the corner of the top element to the outer edge of the glue, and work it all along one rim. Slowly tilt the top element down so that a bead of glue travels outward in the narrowing wedge-shaped gap, leaving a perfect glue layer behind. (The glue should be applied near the center of the lower element if it's steeply curved, and worked outward with a gentle circular rocking motion.)

Using your alignment jig, slide it around until it's in the right place. Press gently to squeeze the glue layer down, and remove any large globs of excess glue. (The glue layer is called the *bond line* in the trade.) Tack the surfaces together before removing the assembly from the jig, using a portable UV lamp or gun. Tacking is not a full cure but turns some of the glue into a gel stiff enough to hold the pieces together. Epoxy can be removed fairly easily at that stage, so double check the alignment before final cure.

How much glue to use is a matter of judgment. The glue layer should be thin, no more than a mil (25 microns), but you really want to avoid bubbles. It may take a couple of tries to get it right—if you get even one bubble, take it apart, add glue, and try again. It's worth practicing this before hand with corn syrup.

If your assembly needs fine alignment, you should make a gluing jig that will accommodate the alignment procedure you plan to use. After the fine alignment (if any), cure the glue as its maker recommends, and perhaps a bit longer just for luck, especially if your glass is thick or doesn't transmit well at 300 nm. UV epoxy isn't quite as strong as two-part epoxy, but this can be fixed by baking it; 15 minutes to an hour at 100 °C does wonders for the strength of the joint.

Epoxy can actually pull hunks of glass out of the sides of your optics if you leave globs there when you cure it. Make sure any significant amount of excess glue is removed before final curing. Temperature swings can cause delamination of larger optics.

12.11.3 Hydroxyl Bonding

Glass and fused quartz can be bonded very simply with a 1:500 solution of KOH in water, applying it like cement and squeezing the bond line down to <10 nm thick. The OH breaks the surface reconstruction of the silicate, leaving dangling silicate bonds that re-form across the bond line as the solvent evaporates. The resulting bond is covalent, highly transparent, and very strong, but the surfaces have to match very closely to begin with.

12.11.4 Temporary Joints: Index Oil and Wax

For temporary use, you may prefer to use index oil or wax, which are available from Cargille with n from about 1.35 to 1.9 in spacings of 0.002, with waxes reaching 2.0. The low index stuff is pretty safe to use, but some of the higher index ingredients are nastier: read the materials safety data sheet (MSDS) and wash your hands a lot.

A drop of the right index oil will reduce front surface reflections by four orders of magnitude, which is not to be sniffed at. Index oil can also be used to bandage damaged surfaces: scatter from a scratch on a lens can be reduced temporarily by rubbing a very

small amount of index oil on it, then polishing it with lens paper,[†] and a pitted window can be helped by putting on a drop of index oil and covering it with a microscope slide.

Oil is applied the same way as UV epoxy, but with still more attention to getting the layer thin, because the surfaces stick together better that way. Even mounted vertically, a good oiled joint with both elements supported from below is stable for months, at least under lab conditions. Wax lasts forever in the lab. You can get much higher viscosity oil for vertical and upside-down use (it's almost as thick as grease).

The temperature coefficients and dispersion of index matching gunk are much higher than those of glass, but that isn't usually a problem; it forms a thin element whose two radii are equal, and the index discontinuity is small, so by the lensmaker's equation it has zero power. Oiled joints also accommodate thermal expansion well.

12.12 CLEANING

We've all seen photographs and slide shows where a large piece of dust or a fiber in the wrong place has made a distracting shadow in the image. The same can happen in instruments, where the consequences are usually far worse than a distraction. Fortunately, most of the surfaces where dust collects are well out of focus, so the shadows are unobtrusive. There is really no way to clean an optic without wear or damage, but some methods are a lot worse than others. Thus it's worth spending some time talking about cleaning.

12.12.1 What Does a Clean Lens Look Like?

Well, if you've ever looked at a camera lens or microscope catalog, you've seen what a lens is supposed to look like. The AR coatings are working beautifully, because there are no fingerprints (which show up bright on an AR-coated surface). We might not have the beautifully positioned lighting of the photographer, but we don't really need it. To examine an optical element, stand with your back turned to an overhead light source (ceiling fluorescents are usually good enough). Rock the element back and forth in your hand, looking at the reflected image of the light source in both the front and rear surfaces (you'll have to refocus your eyes if it's a lens). It should be uniform in brightness and color. Shiny spots are fingerprint grease, which has to go. Next, orient the element so that it reflects some darkish area, for example, the shadow underneath a lab shelf. Rock it back and forth to look for scattered light. Large dust particles are immediately obvious as localized sources of scatter. A broad diffuse scatter is fine dust, condensation (e.g., plasticizers outgassed from plastic mounts), or surface damage such as weathering or etching.

If the lens is to be used for laser work, you have to use a really bright light source to look for even small amounts of scatter. A high powered laser will often burn dust right into the optical surface, destroying it. Scattered light from particles larger than $\lambda/4$ is brighter in the forward direction (small scattering angles), so look on that side. If you're working in the visible at moderate power, use the laser in your system. (Do be careful—most eye damage from lasers happens during setup.) Otherwise a slide projector or microscope condenser works well.

[†]A common photographic darkroom trick to fix scratches on the back of a negative is to smear on a small amount of nose grease with a finger and wipe with lens paper.

12.12.2 When to Clean

Mildly dirty optics are surprisingly innocuous. The areal fraction of the beam that dust occupies is very small, so most of the light is unaffected; the dirt doesn't hurt you until it starts causing so much stray light that your measurement sensitivity is degraded. Dark-field measurements are obviously more vulnerable to dirt than are bright-field ones. The trade-off is between scatter from the dust and scatter from the long thin scratches it will leave when you rub the surface with lens paper or a cotton swab. Scatter tends to depend on the perimeter of the defect more than its area, so this is not necessarily an improvement.

Films are a different matter. A bit of machine oil left in a lens barrel can cause condensation on the surfaces, which really will screw up an optical system. There is a silver lining here, because grease films can be cleaned off without scrubbing (and hence without significant damage), while dust often can't.

Optics don't have time to get dusty when you're actually working on them; most of the dust collects when you're elsewhere. Use dust shrouds, a clean hood, or just a big plastic box with a hinged side that opens when you're working on the system.

12.12.3 Cleaning Lenses

Glass lenses with hard AR coatings on them can be cleaned pretty vigorously without serious damage, as long as you get the dust off first so that it doesn't gouge the surface. Even so, expect the stray light to get slightly worse each time due to accumulated scratches. Rinsing with Kodak lens cleaner or a surfactant solution such as Triton[†] or Kodak PhotoFlo works well unless the lens is heavily fouled with fingerprints, in which case ultrasonic cleaning in a gentle low-suds detergent such as Alconox works well as a first step. Make sure that the optical surfaces don't contact the supports when using ultrasonic cleaning.

If your facility has ultrapure water (18 M Ω deionized or better), then mixing a very small amount of Triton with such good water will make a cleaning solution that will dry without leaving a residue. Mild scrubbing with good quality lens paper, a microfiber cloth, or Q-tips removes most kinds of crud, but will leave fine scratches.

After washing, rinse with very clean water and blow dry with clean dry air or (ideally) clean nitrogen from a LN₂ tank, which is completely oil-free. Don't use machine shop quality compressed air to dry optics.

For everyday cleaning, e.g. getting fingerprints off camera lenses, the microfiber cloths sold for cleaning eyeglasses are quite good. Unlike natural fiber, they don't contain silica and are therefore much gentler. You still have to avoid grinding dust into the lens surface, however.

Some surfaces cannot tolerate such cavalier treatment or cannot be removed for cleaning. Laser tube Brewster windows and soft or water-sensitive IR materials are examples. You can clean these by using a single sheet of lens paper. Put a drop of electronic grade solvent, such as isopropyl alcohol (IPA), acetone, or trichloroethylene (TCE) on the optical surface, lay the lens paper down flat on it (so that the solvent soaks in and wets the whole surface), and slowly drag the lens paper along the surface without applying pressure, as though you were dragging a tarp full of leaves in the garden. Blow dry and repeat, perhaps following a polar solvent such as IPA with a nonpolar one like TCE.

[†]Triton X-100, made by Fisher Scientific, is a very concentrated surfactant that is very clean.

It's important to use electronic grade solvents, since reagent grade ones are not usually pure enough. (Don't use solvents to clean plastic lenses—if you put isopropyl alcohol on acrylic, for instance, it will craze the surface instantly. Stick with aqueous cleaners.)

Some people like to fold the lens paper over several times and hold it in a degreased hemostat, which is easier. If you do this, wear skin gloves while folding the paper, and do it in mid-air. The great thing to avoid in cleaning delicate optics is fouling the lens paper by touching it with your fingers—treat it like a bandage.

Gotcha: Plasticizer Residue. One really insidious source of contamination in a cleaning procedure is leaching of plasticizers from vinyl gloves and plastic parts. These plasticizers are oily substances mixed into plastics to make them soft; the most common one is dioctyl phthalate, also used as vacuum pump oil—its vapor pressure is so low that it stays around forever. As long as they stay put, plasticizers are wonderful, but on your optics, they're a disaster. If you dip your gloved hand into the cleaning solvent, you've just guaranteed yourself dioctyl phthalate all over all of your parts.

12.12.4 Cleaning Gratings

Diffraction gratings must be kept very clean. Contamination will give rise to scatter, and scatter is the enemy of good spectroscopic measurements. The best way to clean a grating is not to get it dirty; cleaning them without damaging them is difficult. A heavily soiled grating may be beyond help even if its surface is otherwise undamaged. Whatever you do, don't touch the grating surface—that's a very expensive mistake indeed.

If the coating is aluminum, the grating should not be soaked long in aqueous solvents because the aluminum may pit slightly, rendering the grating useless. Brief exposure is okay, providing the water is pure enough not to leave residue when it dries—remember, not much residue is needed to produce bright scattered light. If a grating becomes dirty, clean it ultrasonically in a very pure organic solvent such as electronic grade acetone (make sure your ultrasonic cleaner is safe with flammable solvents—some aren't). Gently blow it dry with filtered dry nitrogen or really clean compressed air, chasing the droplets off the ruled surface to avoid residue. Really make sure the air is completely clean and oil-free—one shot of shop-grade compressed air may be enough to kill a grating. Avoid Dust-Off and its ilk, because small amounts of extremely cold liquid may come out with the spray and craze the grating surface.

Some solvents may soften the polymer layer underneath the coating, causing the grating to wrinkle at the edges, so ask the manufacturer first.

Alternatively, you can use a powerful surfactant such as Triton in *very* pure water (18 M Ω deionized preferably, 0.1 M Ω minimum). Avoid using ordinary detergents, or anything but the purest water; they leave spots when they dry, almost no matter what you do. Electronic grade solvents are the cleanest but are not cheap.

Oil contamination wicks down the grating grooves, greatly reducing diffraction efficiency and producing interesting abstract wavefronts. Fortunately, it's not as bad as it looks; rinsing in acetone will completely restore the grating.

12.12.5 Opticlean Polymer

One excellent, although slow, way to get smooth optical surfaces very clean is to use Opticlean Polymer,[†] which is a solution of polyvinyl alcohol plastic in an organic solvent.

[†]Opticlean is made by Bradford Laboratories.

You put it on with a swab (a clean room Q tip is best, if you have them), wait for it to dry in a horizontal position, and then peel it off by attaching tape to the top and peeling the tape. The resulting surface is impressively dust-free, suitable for optical contacting or even frustrated TIR measurements.

The two things you have to remember with Opticlean are to use enough to make the film strong, and to be very careful with the edges of the optic. If you let it spill over onto a ground glass area, such as the sides of a prism or the edge of a lens, it will stick tenaciously, and lead to ripping of the film and fouling of the edges. If you're doing a faceted device such as a prism, don't let the films on adjacent facets merge, or it won't strip.

On the other hand, do use enough at the edges. If the film at the edge is too thin, it won't detach well either. The solution is to very carefully form a raised fluid bead at the edges of the facet. This happens naturally with convex surfaces due to gravity, but has to be done manually with flat ones (concave ones can be dried upside down).

The right quantity to use varies with the formulation you choose, but something like a teaspoonful (5 mL) is about right for a 50 mm square area. It looks like a lot more when it's wet. Certain types of Opticlean exhibit unstable fluid surfaces, so that thin areas spontaneously thin further, limited by the rate of drying and the surface tension. With these types, you have to make the layer slightly more carefully, since the thinned areas are difficult to strip off.

Bradford Labs sells its own swabs and tape strips, but they're very expensive. Use a tape that doesn't shed particles: 3M type 471 white PVC tape is very clean (we use it in Class 1 semiconductor clean rooms) and adequately sticky, although stickier would be better. Ordinary wood-handled Q-Tips work pretty well for spreading the polymer, and if they shed a fiber or two, it gets removed with all the rest when you strip the polymer off.

Resist the temptation to try stripping the film before it's completely dry. You can speed it up a little by mild heating (40 °C or less), for example, by pulling your desk lamp down near it. Don't try to force dry it with a heat gun or oven, because it bubbles, spits, and then refuses to come off.

12.13 ENVIRONMENTAL CONSIDERATIONS

12.13.1 Fungus

Most optical systems will develop fungus growth on their surfaces when used in warm and humid conditions. It was worse back when optical cement was made of tree sap (Canada balsam)—in the Pacific battles of WWII, soldiers' field glasses could become unusable in a few days, as fungus ate the cement. It is still common to see camera lenses with spidery fungus growths inside them. Left unchecked, this will eventually etch the glass, destroying the lens. In general, the same measures you take to prevent mildew in cloth appear to help control fungus—keep the optics dry and either (a) sealed or (b) very well ventilated.

12.13.2 Coating Drift

Coatings, particularly interference coatings, exhibit drift with time. They are rarely fully dense, and since they are deposited at high temperatures and low pressures, dielectric coatings are usually somewhat nonstoichiometric and so only metastable in ambient

conditions. Over time, these coatings hydrate, causing their indices to go up slightly, and sometimes haze develops. Interference filters that are critical for the long-term performance of your instrument should be protected against humidity if possible. (See Section 5.4.6 for more details.)

12.13.3 Lens Staining

Optical glass is susceptible to staining and corrosion due to environmental exposure. Quartz and crown glass don't stain easily, but flints and filter glass are quite vulnerable—a surface layer forms that gradually becomes opaque. Catalog glasses are rated for stain resistance (see Section 4.2.1).

12.13.4 Drift from Temperature and Humidity

In building a complicated optical system needing fine alignment, it's worth calculating what effects temperature and humidity will have on your alignment. For example, the author has a laser built on a cast aluminum optical bench with a coat of paint on it—under the optical mounts. The paint is relatively thick, at least 125 microns; since the plastic binder can easily expand by several tenths of a percent due to humidity, this is a possible source of instability. (It's a drop in the bucket for this laser, as it turns out, but it might not be for yours.)

Signal Processing

One man's noise is another man's data.

—Anonymous

13.1 INTRODUCTION

Most of the time, the end result of a measurement is a digital data set. Preparing the squishy analog stuff coming from the front end amplifier to be digitized is the job of the *analog signal processing subsystem*, which is the subject of this chapter. The objective is to pull the desired signal out of whatever noise and spurious signals it is imbedded in, amplify and filter it appropriately, and digitize it without introducing artifacts.

The basics of designing circuits to do this sort of thing are covered ably and engagingly in Horowitz and Hill, and some of the fine points are covered in Chapters 15 and 18; here we look mainly at the systems engineering and theoretical aspects. Often a high level design is all you have to do, since there may be off-the-shelf components that you can just connect together to make a complete subsystem suitable for demos or for lab use.

The approach advocated here is top-down and iterative; figure out the broad requirements, budget the system specifications for noise, bandwidth, and so on, then trade off the difficulties brought to light. We'll go through some examples along the way to make the process clear. The other common technique is to do it unsystematically; put together a prototype optical system, then string a bunch of electronic modules and lock-ins together with coax cables to do the job. Both of these are valid ways of working, but they get you to different places: the systematic approach gets you a design with optimal performance that can be produced inexpensively, and the unsystematic one gets you data in a hurry from one setup.

Even an experimental apparatus that is going to be built mostly from commercial gear needs a noise budget and a high level system design, because that will help it work well and not have to be taken apart and redone. These are similar to the photon budgets we did in Chapter 1, but there are fewer first principles to work from here, and more idiosyncrasies of components to worry about. This chapter assumes familiarity with basic circuits, including resistance and reactance (and their antipodes, *conductance* and *susceptance*).

Aside: Fourier Transform Sign Conventions. Positive frequencies are easier to handle than negative ones. In the first part of the book, we were doing a lot of spatial transforms, and accordingly took as our plane wave eigenfunction $\exp(i(kz - \omega t))$, whose time dependence is a negative exponential. This convention is standard in physics books and leads to the forward Fourier transform in time having the kernel $\exp(+i2\pi ft)$, and a phase lag (or delay) being $\exp(i\phi)$. (A delay means that the measured phase corresponds to that emitted at a previous time, $t \rightarrow t - \Delta t$.) In electronic signal processing, we do mostly temporal transforms, so the convention is that a positive frequency has a positive exponential time dependence. Thus the forward Fourier transform kernel is $\exp(-j2\pi ft)$ and a delay is $\exp(-j\phi)$. Electrical engineers use j instead of i for $\sqrt{-1}$, to avoid confusion with i for small-signal current, and it's a good reminder here. Since one rarely does signal processing and wave propagation in the same calculation, the opportunity for confusion is slight. Accordingly, we bow to convention and use the $-j$ transform convention in our electronic discussions. This mild form of schizophrenia will enable you to talk to both physicists and engineers without becoming perplexed (by sign conventions, at least).

13.2 ANALOG SIGNAL PROCESSING THEORY

The foundation of signal processing is the theory of linear, time-invariant systems. This theory is so workable and intuitive that nonlinearities or time variations are usually treated as perturbations of linear time-invariant systems. This is similar to the way we treated optical calculations in Chapter 1 (the theorems and definitions are in Section 1.3.8). It's worth having a copy of Bracewell at your elbow while doing this stuff the first few times.

13.2.1 Two-Port Black Box

A sufficiently simple signal processing system can be modeled as a deterministic *two-port network*, that is, a black box with two pairs of wires coming out of it. For any input voltage $g_1(t)$ applied to one port, we get a corresponding output voltage $g_2(t)$ on the other one, which will be uniquely determined by $g_1(t)$ (apart from noise and drift, which we'll treat separately). We assume that the network response is causal, so $g_2(t)$ depends only on current and past values of g_1 , so that for some operator M ,

$$g_2(t) = M(g_1(t')), \quad t' \leq t. \quad (13.1)$$

13.2.2 Linearity and Superposition

A network is said to be *linear* if it obeys two rules for all input functions g :

1. $M(\alpha g(t)) = \alpha M(g(t))$, *scale invariance*, and
2. $M(g_1(t) + g_2(t)) = M(g_1(t)) + M(g_2(t))$, *superposition*.

In algebra, these conditions define a homogeneous linear system, one with no offset term. The inevitable offsets of real networks are put in by hand afterwards, like noise and drift.

The most general causal relationship that satisfies these two conditions is the integral equation

$$g_2(t) = \int_{-\infty}^t m(t, t')g_1(t')U(t-t')dt', \quad (13.2)$$

where $U(t)$ is the Heaviside unit step, $U(t) = 0$ for $t < 0$ and 1 for $t > 0$.

13.2.3 Time Invariance

Most of the time we're dealing with networks whose behavior is not a strong function of time—fixed-tuned filters, amplifiers made with fixed-value components and run at constant bias, that sort of thing. For these networks, we can assume *shift invariance*: if $g_1(t)$ produces $g_2(t)$, then $g_1(t - \tau)$ will produce $g_2(t - \tau)$; that is, there's nothing special about the time origin. This is equivalent to the condition that the kernel function $m(t, t')$ depends only on $t - t'$. Replacing $m(t, t')$ with $h(t - t')$, requiring that $h(\tau) = 0$ for $\tau < 0$, we get the convolution

$$g_2(t) = \int_{-\infty}^{\infty} g_1(t')h(t-t')dt'. \quad (13.3)$$

13.2.4 Fourier Space Representation

As we saw in Section 1.3.8, a convolution in real space corresponds to a multiplication in the frequency domain. This means that the response of a linear time-invariant network can be completely described by a multiplication of the input spectrum $G_1(f)$ by a complex transfer function $H(f) = \mathcal{F}h(t)$. One important consequence of this is that if the input is an exponential function (a δ -function in \mathbf{k} -space), the output is exponential as well, with the same time constant. This applies for both real and complex exponents, although of course complex ones must appear in conjugate pairs, because the voltages g_1 and g_2 are real numbers.[†]

Because g_1 and g_2 are real, $h(t)$ must also be real, since otherwise some choice of g_1 would produce a complex-valued output from a real-valued input. This means that H is *Hermitian*— $\text{Re}\{H\}$ is an even function, and $\text{Im}\{H\}$ is odd. In polar form,

$$\begin{aligned} H(f) &= |H(f)|e^{j\phi(f)}, \\ H(-f) &= |H(f)|e^{-j\phi(f)}. \end{aligned} \quad (13.4)$$

Another way to say this is that in order for $h * g$ to be real, h must apply equal gain and equal and opposite phase shifts to $\exp(j2\pi ft)$ and $\exp(-j2\pi ft)$.

If we put a function, say, $\cos(2\pi f_0 t)$, into the input, then we get

$$\begin{aligned} g_2(t) &= \frac{|H(f_0)|}{2} (e^{j\phi(f_0)} e^{j2\pi f_0 t} + e^{-j\phi(f_0)} e^{-j2\pi f_0 t}) \\ &= |H(f_0)| \cos(2\pi f_0 t + \phi(f_0)). \end{aligned} \quad (13.5)$$

[†]A circuit with a 1 ms time constant can in principle reproduce a 1 μ s real exponential, but the tiny exponential response is masked by the turn-on transient in such cases.

That is, the effect of a linear, time-invariant network on a sinusoidal input is to multiply it by a constant and shift its phase.

Aside: Poles and Zeros. Except in microwave work, we build filters out of lumped elements—individual resistors, capacitors, and inductors, whose impedances go as powers of f . Combining them in series or parallel produces composite impedances and transfer functions that are rational functions of the element impedances, and a rational function of a rational function is another rational function. A rational function is completely specified by a single gain and the roots of its numerator and denominator polynomials, which are the zeros and poles of the function. Thus the behavior of all our circuits and models is determined by the poles and zeros of their transfer functions. So far, so classical. The confusing thing about it for newcomers is that the poles never lie on the real frequency axis, where all our measurements take place.[†] RC poles (such as an ordinary RC decoupling network) lie on the (positive) imaginary frequency axis, and LC poles lie in pairs, distributed symmetrically above the real frequency axis so that there's no phase shift at DC. Analog engineers are so used to this that they say a $10\text{ k}\Omega$ – 15 nF lowpass RC network “has a pole at 1 kHz,” pointing at the -3 dB (half-power) frequency of a limp-looking smooth rolloff, and not at an infinite singularity. It's just mild sloppiness—the -3 dB point is numerically equal to the position of the real pole on the $+j$ imaginary axis, because the transfer function is $H = 1/(1 + j2\pi f RC)$ —the pole is really at $f = j/(2\pi RC)$.

13.2.5 Analytic Signals

The bookkeeping necessary to carry along the explicitly real-valued nature of the signals in the network more than doubles the amount of algebra, which of course increases the number of blunders we make by at least four times. Furthermore, in more advanced network theory we need to use residue calculus a lot, which makes it even more inconvenient, since contours in the complex plane usually have to be chosen differently depending on the sign of the exponent. Accordingly, we make the following formal rules.

1. Set all the components at negative frequency to 0.
2. Double the positive ones to keep the real part the same.
3. Leave the DC alone.
4. Whenever any measurement or nonlinear operation (e.g., computing the power) is to be done, discard the imaginary part beforehand.
 - 4a. If you can be sure that the positive and negative frequency lobes of your signal never overlap significantly or intermodulate with each other, you can use the analytic signal in nonlinear operations (e.g., frequency modulation).

This keeps the real part $g(t)$ the same and adds an imaginary function, $-j\mathcal{H}g(t)$, producing the *analytic signal* $\hat{g}(t)$, as shown in Figure 13.1. The analytic signal is purely a calculating convenience, but as the convenience is considerable, and it is in general use, it's worth understanding and using it.

[†]Zeros can lie on the real axis, but usually don't.

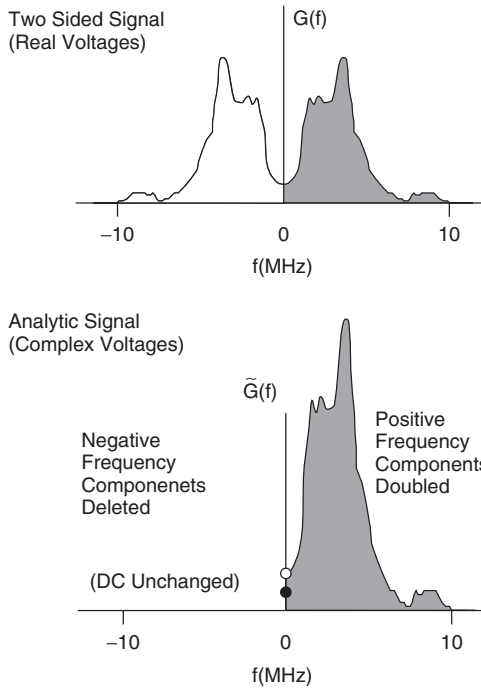


Figure 13.1. Analytic signal definition. Note that DC is a special case.

The operator \mathcal{H} here is the Hilbert transform, which is an application of Cauchy’s integral formula,

$$\mathcal{H}g(t) = g(t) * \frac{-1}{\pi t} = \frac{1}{\pi} \mathbf{P} \int_{-\infty}^{\infty} \frac{g(t')}{t' - t} dt' \tag{13.6}$$

where \mathbf{P} denotes the Cauchy principal value. This is just an organized way of applying rules 1–4 to a given function.

The Hilbert transform phase-shifts every positive frequency component by $\pi/2$ radians, and every negative frequency component by $-\pi/2$. Multiplying by j and subtracting it from the original function cancels the components at $f < 0$ and doubles those at $f > 0$. The principal value makes sure that the component at DC stays the same, which is a good thing since you can’t phase-shift DC in real life. The term *analytic signal* is slightly too restrictive, since the procedure also applies to functions with singularities, such as delta functions. Here are some examples:

True Signal	Hilbert transform	Analytic Signal
1	0	1
$\cos(2\pi ft)$	$-\sin(2\pi ft)$	$e^{j2\pi\mathcal{H}t} \iint f$
$\sin(2\pi ft)$	$\cos(2\pi ft)$	$-je^{j2\pi\mathcal{H}t} \iint f$
$\delta(t)$	$\mathbf{P}\{-1/(\pi t)\}^\dagger$	$j/(\pi t)$
$g(t)$	$\mathcal{H}g(t)$	$g(t) - j\mathcal{H}g(t)$

[†]Delta functions have meaning only inside integrals, and the same is true of their Hilbert transforms. The \mathbf{P} indicates that this operator applies the Cauchy principal value to any integral containing it.

Like the real-space propagators in Sections 1.3.2 and 9.3.2, you rarely use the Hilbert transform in real space; it is much more common to take the transform, chop off the negative frequencies, double the positive ones, and transform back.

The Hilbert transform has practical applications in image reject mixers, which we'll see in Section 13.8.7, and theoretical ones in the synthesis of stable networks and optimal filtering, where the condition that $h(t) = 0$ for $t < 0$ forces a similar analyticity constraint on the frequency domain (see Section 13.8.4).

Analytic signals crop up all the time in signal processing literature and lead to confusing factors of 2 in bandwidths and power spectral densities; for example, the Fourier transform of a 1 second averaging window is 1 Hz wide (-0.5 Hz to $+0.5$ Hz) if negative frequencies are allowed, but only 0.5 Hz on an analytic signal basis. It's obvious that the Hilbert transform has the same squared modulus as the AC part of g , so the AC power of the analytic signal is twice that of g , which is a common source of blunders if rules 1–4a are not kept in mind.

Although amplitude and frequency ought to be local (i.e., instantaneous) quantities, neither is easy to define for a real-valued modulated wave; normally we have to take some average over one or more cycles. The analytic signal allows us to define them uniquely. If the analytic signal is $\hat{g} = A(t) \exp(j\phi(t))$, the instantaneous amplitude is $A(t)$, and the instantaneous frequency is

$$f(t) = \frac{1}{2\pi} \frac{d\phi}{dt}. \quad (13.7)$$

13.3 MODULATION AND DEMODULATION

A sinusoidal signal whose amplitude and phase are not intentionally made to vary with time is referred to as *continuous wave* (CW). A CW signal contains no useful information, except perhaps as a frequency or amplitude standard. Its spectrum is a single narrow spike.

The impressing of information on a CW *carrier* by changing its amplitude, phase, or both is called *modulation*. Modulation can be artificially imposed, as in voice or data communications, but in optical instruments, it is more commonly the result of the physical process we're measuring. Modulation modifies the spectrum of the carrier by reducing the strength of the central spike and putting *sidebands* on it, as shown in Figure 13.2. Sidebands appearing on the high and low frequency sides of the carrier frequency f_C are called upper and lower sidebands, respectively.[†] Some of the more common types of modulation appear in Figure 13.3 and their spectra in Figure 13.4.

Aside: Continuous Wave. Like many other analog signal processing terms, this one has its origins in early radio, where the first transmitters used an energetic spark in a weakly selective network to generate their output directly from a Morse key. Since a spark is a good engineering approximation to a delta function, this was an inefficient use of spectrum, to put it mildly. The first CW transmitters, where a CW oscillator was turned on and off by the key, were a revolutionary improvement, allowing lower power sets to reach much further, and more stations to fit in a smaller geographical area without interfering with each other.

[†]This usage is not commonly applied to FM type sidebands, which are not separately intelligible.




Figure 13.2. Modulating a carrier signal puts modulation sidebands on it. The structure of the sidebands, their phase relationships with each other, and how much of the carrier survives depend on the modulation scheme.




Figure 13.3. Time-domain traces of different modulation types. The modulation waveform in each case is a sine wave at $0.15 f_c$. Note the phase reversals in the DSB modulated signal.

13.3.1 Terms

The discussion of modulation requires a few terms, which you must clearly understand, or the whole subject (which is a pretty one, full of good physical insight) will become opaque very rapidly (Figure 13.4). We’ve already encountered analytic signals, carriers, and envelopes. Here are a few more.

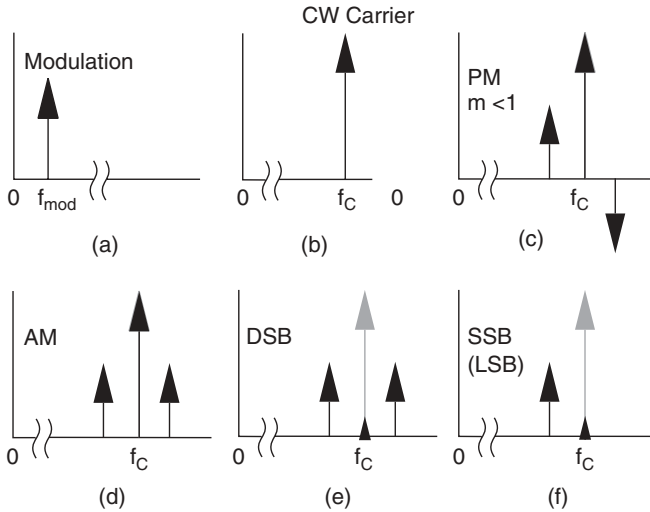


Figure 13.4. Common modulation schemes, in phasor view: (a) baseband modulating signal, (b) CW carrier, (c) narrowband PM, (d) AM, (e) double sideband (DSB), and (f) single sideband (SSB).

Baseband: Signals such as speech or video, which occupy a frequency band extending from near DC to some much higher frequency, are known as *baseband* signals. Typically what emerges from a demodulator is referred to as baseband, and nearly always the signal to a digitizer is baseband (stroboscopic digitizers such as digital receivers are exceptions). It’s a fuzzy but useful concept.

Bandwidth: The output spectrum of our nice linear, time-invariant system is the product of the input spectrum with the system’s transfer function. Because the signal we desire is always corrupted by additive noise before we get our hands on it, we want to choose the transfer function so as to preserve the signal while rejecting as much of the noise as we can. The wider it is, the faster it will respond but the more noise it will let through. The range of frequencies that a system will respond to is its *bandwidth*. A bandpass filter whose transfer function is down by a factor of $1/\sqrt{2}$ (0.5 in power, or 3 dB) at 99 and 101 MHz has a *3 dB bandwidth B* of 2 MHz. This is the most commonly quoted bandwidth, but there are special definitions of bandwidth for different situations. Of these, the most important is the *noise bandwidth*, which is the total noise power divided by the peak value of the output noise spectral density; that is, it’s the equivalent width of the noise power spectrum. The idea is most useful when the spectrum is flat in the middle. The noise bandwidth of an RC lowpass is $1/(4RC)$, which is $\pi/2$ times wider than the 3 dB bandwidth.

Spectral Density: Radio people have very concrete minds, which helps in keeping things straight. Unfiltered shot noise and thermal noise have flat frequency spectra. Other kinds of noise, such as $1/f$ noise and filtered thermal noise, have pronounced frequency variations. What we mean by those statements is that if we connect a noise source to a filter with a constant and very narrow bandwidth *B*, and tune it around, the average electrical power coming out of the filter will vary with

tuning or it won't, and that if B is known accurately, we can normalize it out. For a sufficiently narrow filter, doubling its bandwidth will double the power at its output. The *power spectral density* dP/df in watts per hertz is defined by

$$\frac{dP}{df} = \lim_{B \rightarrow 0} \frac{P_{\text{out}}}{B}. \quad (13.8)$$

Because of the concrete mindset alluded to above, this is almost always called the *1 Hz noise power*, because 1 Hz is usually narrow enough and the units come out right. People also talk about *voltage noise spectral density* or *current noise spectral density*, measured in volts or amps per square root hertz. Since it's power that is linear in bandwidth, these aren't really densities, but there's no getting rid of them at this point. When you have to talk about them, for example, for the transimpedance amplifier designs of Chapter 18, you can call them the 1 hertz voltage and current noise. (Just remember to scale them by \sqrt{B} .)

Furthermore, because the size of the quantities is frequently very small, we often quote them in a mongrel unit, dBm/Hz. The thermal noise is kT per hertz, 4.1×10^{-21} W/Hz or -173.8 dBm/Hz at 300 K. If you multiply this by the bandwidth, you'll get some pretty strange numbers, so either convert back to watts first or stay in decibels and add $\log_{10} B$ instead of multiplying.

Envelope: Modulation is expressed mathematically as multiplying a carrier by a time-varying complex envelope $\Phi(t)$, a widely useful move that we used in deriving the group velocity in Section 1.2.2, the paraxial wave equation in Section 1.3.2, and the eikonal in Section 9.2.3 A modulated signal $g(t)$ has an envelope $\Phi(t)$ defined by

$$g(t) = \Phi(t)e^{i2\pi f_C t}, \quad (13.9)$$

where f_C is the carrier frequency. The modulus of the envelope is the amplitude modulation and multiplies the carrier amplitude. The envelope phase adds to the carrier phase, and so is the phase modulation. If $f_C \gg f_{\text{mod}}$, the carrier peaks trace out a curve of $|G(t)|$ (as on a slow oscilloscope trace), but the envelope phase is invisible unless we use a phase-sensitive detector.

Modulation Frequency: Modulation frequency is different from regular frequency. Modulation frequency is the frequency of the modulation, which is related to but not identical with frequency offset from the carrier; for upper sideband (USB) modulation, $f_{\text{mod}} = f - f_C$, but in LSB, $f_{\text{mod}} = -(f - f_C)$, in AM and DSB each f_{mod} corresponds to two frequency offsets ($\pm f_{\text{mod}}$), and in FM to an infinite number, spaced at harmonics of f_{mod} from f_C . Fundamentally, the modulation frequency is the frequency that goes into the modulator or comes out of the demodulator.

Example 13.1: 5 Hz FM Can Be 1 MHz Wide. A frequency-modulated signal centered at 100 MHz, which sweeps sinusoidally back and forth by ± 0.5 MHz at a sweep frequency of 5 Hz, has a spectrum that looks like a solid blob a megahertz wide, but is really a forest of spikes at 5 Hz intervals (assuming the oscillator is quiet enough that its noise doesn't fill in the spaces completely). The modulation frequency here is 5 Hz, not 1 MHz or anything in between; the baseband signal fed to the modulator was a 5 Hz sine wave.

Sideband Folding: The distinction between frequency and modulation frequency carries over into the time domain as well. The modulating signal in Figure 13.2 looks as though it has a bandwidth of about 0.25 MHz, but the amplitude-modulated signal is twice as wide. However, the temporal response hasn't changed—putting your signal on a carrier doesn't make its envelope vary twice as fast. This is one of those cases where in principle we ought to transform back to the real-valued representation first, but it's easy to clear up the mystery. The one-sided modulation spectrum we plot is actually the analytic signal representation of the modulation, where we ignore the negative frequencies. Those negative frequencies are still there from the system's point of view, so modulation puts both the positive and negative frequency lobes on top of the carrier. The analytic signal math still applies—there is a mirror image of the two-lobed signal around $-f_c$, way off to the left of the axis. (This mild wart is worth tolerating—otherwise you'd have four products to worry about instead of two.)

AC Versus DC Measurement: It may be a bit confusing when we say that an AC measurement has twice the noise bandwidth of an equivalent DC measurement. Why does the bandwidth suddenly go from B to $2B$ when we move 1 Hz away from DC? The reason is that if we move from DC at all, we have to move further than B due to sideband folding. Consider an analog TV signal 6 MHz wide coming in on a laser beam. How would you like to demodulate it? IFs between 0 and 6 MHz are out, because the sidebands will fold over and the resulting interference will destroy information—your picture will be completely lost. DC works because the two sidebands become one sideband, and a high IF works because you can filter out the other sideband. The two schemes will differ by 3 dB in SNR, and there isn't any middle ground.

Noise Gain: The *noise gain* of a system is the partial derivative of its output with respect to a small input perturbation. It may be specified as an average over the total noise bandwidth (*q.v.*) or as a function of frequency. For an op amp, the noise gain is equal to the closed-loop noninverting gain (see Section 18.4.2).

I and Q: In Section 13.9, the envelope Φ of a modulated signal is a complex function, which can be represented in polar coordinates as amplitude and phase, or in rectangular coordinates as real and imaginary parts. In signal processing language, these are called *I* and *Q*, for *in phase* and *quadrature*. Quadrature is an old word for taking an integral, which also survives in the term *quadratic* for a second-order polynomial, and integrating a sinusoid shifts its phase by 90° . (It doesn't refer to a quarter of a cycle, although that's not a bad way of remembering it.) The most familiar sort of *I/Q* demodulator in the lab is a two-phase lock-in amplifier.

13.3.2 Phasors

Besides vocabulary, we also need to be able to visualize what's going on in a modulation process. An intuitive connection between the structure of the sidebands and what they do to the carrier is vital, and the best tool for this is the *phasor*. Phasors are vectors in the complex plane, representing the instantaneous value of each component of the analytic signal, and we visualize them spinning around the origin, as shown in Figures 13.5 and 13.6. They are often plotted in a composite 3D coordinate system, in which the

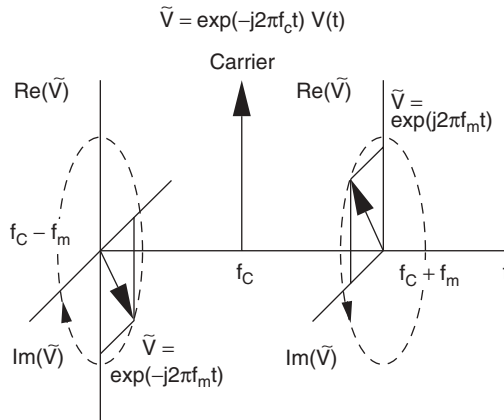


Figure 13.5. Rotating phasor plot of a narrowband PM signal, plotted in coordinates that corotate with the carrier. The oppositely phased sidebands make the amplitude stay constant when the phasors are added vectorially.




Figure 13.6. Two-dimensional phasor construction for narrowband angle modulation (FM and PM) with a modulation index $m = 1$ radian, showing how the sideband phasors add up to a constant radius.

third dimension is frequency, a sort of complex-valued version of a spectrum analyzer display. Since each frequency component spins around at a different speed, visualization is possible only with discrete frequency components (i.e., one or a few sinusoids), but this is enough. A second simplification is to transform to corotating coordinates, so that the carrier phasor is constant and the upper and lower sidebands rotate in opposite directions. Vector addition of the phasors then allows us to see easily that the in-phase sidebands of AM cause the instantaneous amplitude (length) to change but keeps the phase (direction) constant, and the 180° out of phase sidebands of FM always sum to zero length change but cause the carrier phase to wiggle back and forth. (See Figures 13.4 and 13.5)

13.3.3 Frequency Mixing

One of the most useful operations in signal processing is *frequency mixing*. The term *mixing* is familiar from audio, where it describes adding different sound sources together

before feeding them into a recorder or amplifier. In signal processing, mixing refers to multiplying signals together rather than adding them, and the two usages must be kept quite distinct. The typical mixing operation is to multiply a signal at the RF port by a sinusoidal *local oscillator* signal at the LO port:

$$V_{IF} = k V_{RF} \cos(2\pi f_{LO}t), \quad (13.10)$$

so a component of V_{RF} at f produces

$$\begin{aligned} V_{IF} &= k \cos(2\pi f_{LO}t) \cos(2\pi ft + \phi) \\ &= k[\cos(2\pi(f_{LO} + f)t + \phi) - \cos(2\pi(f_{LO} - f)t - \phi)]. \end{aligned} \quad (13.11)$$

The result is that both the RF and LO frequency components disappear, being replaced by new components at $|f_{LO} \pm f|$. There are a couple of things to keep in mind here. One is that the frequencies and phases add and subtract, they don't multiply or divide, and that it's phase and not time delay that matters. If you mix signals at 10 GHz \pm 1 kHz, you'll get outputs at 20 GHz and at 2 kHz. If you delay one of the 10 GHz signals by 10 picoseconds, which is 1/10 cycle, you will phase shift both the 20 GHz and the 2 kHz signals by 1/10 cycle—equivalent to a time delay of 5 ps for the one and 50 μ s for the other. (Note that no causality violation is involved here—we'll talk about what happens with more complicated signals later.) All the derivatives of the phase also add and subtract, of course, so that modulation frequency, modulation index, and frequency offsets stay the same—time-domain multiplication by an exponential is a frequency-domain convolution with a shifted δ -function. The second thing is that in an ideal linear multiplier, the amplitudes of the sum and difference go as the product of the amplitudes of the LO and RF signals, but most practical mixers are run with the LO saturated, so that LO amplitude fluctuations are not impressed on the mixing products.

13.3.4 Amplitude Modulation (AM)

The simplest form of modulation is to change the amplitude of the carrier, by multiplying it by a real, nonnegative envelope function. This is naturally called *amplitude modulation* (AM). Achieving pure AM is not always trivial, since many processes producing it will modulate the phase as well (e.g., changing the gain of an amplifier usually makes it faster or slower). The *Modulation depth* is the ratio of the peak–valley envelope excursion to the peak height (the same definition as fringe visibility). A sinusoidal carrier modulated to a depth d by a sinusoidal signal at frequency f_{mod} is

$$\begin{aligned} v(t) &= \frac{1 + d \cos 2\pi f_{\text{mod}}t}{1 + d} \cos 2\pi f_C t \\ &= \frac{1}{1 + d} (\underbrace{\cos 2\pi f_C t}_{\text{Carrier}} + \underbrace{\frac{d}{2} \cos 2\pi(f_C + f_{\text{mod}})t}_{\text{USB}} + \underbrace{\frac{d}{2} \cos 2\pi(f_C - f_{\text{mod}})t}_{\text{LSB}}). \end{aligned} \quad (13.12)$$

The carrier has grown sidebands of equal strength, spaced above and below the carrier frequency by f_{mod} . A characteristic of AM is that the power in each sideband cannot be more than a quarter of the carrier power, which doesn't help the SNR one bit.

13.3.5 Double Sideband (DSB)

If our modulator can merely turn the amplitude of the sine wave up and down, we're limited to the AM case. If we have a more capable modulator, such as a double balanced mixer (see below), our modulation envelope can go negative as well, becoming

$$\begin{aligned} v(t) &= (\cos 2\pi f_{\text{mod}}t) \cos 2\pi f_C t \\ &= \frac{1}{2}(\cos 2\pi(f_C + f_{\text{mod}})t + \cos 2\pi(f_C - f_{\text{mod}})t). \end{aligned} \quad (13.13)$$

The carrier goes away, and only the AM-type sidebands are left. This sort of modulation is pedantically called *double sideband, suppressed carrier*.

13.3.6 Single Sideband (SSB)

If we use a sharp filter to chop off one of the sidebands in a DSB signal, we get single sideband modulation. SSB is just the shifting of the modulation by the carrier frequency, so that a modulation component at f_{mod} gets moved to $|f_C \pm f_{\text{mod}}|$. Choosing the plus sign leads to *upper sideband* (USB) modulation, the minus sign to *lower sideband* (LSB).[†] The two are not equivalent, since an LSB modulated signal has had its low modulation frequencies flipped around to higher signal frequencies. A common error in analog signal processing is to use an odd number of LSB conversions, so this flipping is not corrected. The symmetry of the sidebands in AM, DSB, and FM make this much less of a concern there.

13.3.7 Phase Modulation (PM)

Phase modulation is just what it says: the modulation signal goes in the argument rather than as a real number multiplier outside. From an analytic signal point of view, sinusoidally phase modulated waveform is multiplied by $\exp[jm \cos(2\pi f_{\text{mod}}t)]$ for some value of m —the envelope magnitude remains 1 but the phase varies. This can be done, for example, with a Pockels cell whose optic axis is aligned with the \mathbf{E} field of our beam (so there are no polarization funnies), or for RF signals with a transmission line built on a ferroelectric ceramic such as barium titanate; it can also be done with variable all-pass networks, carefully linearized (see Section 15.12.2). A PM signal has a more complicated spectrum than an AM one, since it has an infinite number of harmonics:

$$\begin{aligned} v(t) &= e^{j(2\pi f_C t + m \cos 2\pi f_{\text{mod}}t)} \\ &= \sum_{n=-\infty}^{\infty} j^n J_n(m) e^{j2\pi(f_C + n f_{\text{mod}})t}, \end{aligned} \quad (13.14)$$

where J_n is as usual the n th order Bessel function of the first kind, and the constant m is the *modulation index*, which is the peak phase deviation in radians. The Bessel function $J_n(m)$ has significant amplitude only for $n \lesssim m$, falling off exponentially with n for $n > m$. Negative order Bessel functions obey

$$J_{-n}(x) = (-1)^n J_n(x). \quad (13.15)$$

[†]This assumes that $f_{\text{mod}} < f_C$. (Why?)




Figure 13.7. Detail of a wideband PM signal with $m = 5.0$, showing the phase relationships of the sidebands. The lines shown represent phasors, with the imaginary ones shown as dotted lines.

The instantaneous frequency is

$$\begin{aligned}
 f(t) &= \frac{1}{2\pi} \frac{d}{dt} (2\pi f_c t + m \cos 2\pi f_m t) \\
 &= f_c + m f_m \sin 2\pi f_m t,
 \end{aligned}
 \tag{13.16}$$

so (as expected) the frequency swings between the limits $f_c \pm m f_m$, which are roughly the points where the Bessel functions cut off. Figure 13.7 is a phasor construction showing how the FM sidebands add up to produce phase modulation with no amplitude variation. Figure 13.7 shows the complex amplitudes of the carrier and modulation sidebands for a phase modulated signal with $m = 5.0$. Note that the carrier’s phase is inverted,[†] and that the phases of the sidebands are not trivial to predict without using (13.14).

There are three relevant regimes of m : $m \ll 1$, narrowband PM/FM, where only the ± 1 order sidebands are significant; $1 \lesssim m \ll f_c/f_m$, quasistatic wideband PM/FM, where many sidebands contribute; and $m \gtrsim f_c/f_m$, the nonquasistatic case. The first two are the usual situations, where the variation of the envelope phase is slow compared with the carrier period, so that the modulation sidebands die out well before they get to DC, making the analytic signal expression (13.14) valid. The nonquasistatic case, where the sidebands are significant in amplitude down to DC, is a complicated mess—for example, modulating a 100 MHz carrier at 1 GHz f_m , or deviating it ± 6 radians ($m = 6$) at a 20 MHz f_m . For modulation like this, the analytic signal representation has to be discarded, and the phase calculated explicitly as the time integral of the frequency.

[†]The carrier goes to 0 at the roots of J_0 , which is how you calibrate phase and frequency modulators. The first Bessel null is at $m \approx 2.405$.

Messiness is typical of situations where positive and negative frequency lobes start to overlap. Nonquasistatic PM/FM is one such situation; others are aliasing in gratings (Section 7.3.1) and in sampled-data systems (Section 17.4.4).

13.3.8 Frequency Modulation (FM)

Frequency is the time derivative of phase in cycles, so frequency modulation (FM) is a close relative of PM. The two are often called *angle modulation* for this reason. A given modulation voltage moves the carrier by a specific phase offset in PM and a specific frequency shift in FM. In principle, the modulation signal could be integrated and applied to a phase modulator to produce FM. This is not a practical approach, however; a DC term in the modulation waveform causes the average frequency of the signal to change, so that the phase shift grows secularly with time, eventually exceeding the range of any phase shifter. Because of this, FM is in fact generated by putting the modulation on the control voltage pin of a voltage controlled oscillator (VCO) and is demodulated differently as well. The integrating action of the VCO reduces the modulation index of high frequency components by a factor of f_m , so that the total signal bandwidth (approximately $m f_m$) is roughly independent of f_m , a desirable characteristic when f_m varies widely, as in music broadcasting. FM demodulators are relatively insensitive to oscillator phase noise, since the same integrating action boosts signal power at low modulation frequencies by a factor of f_m^{-2} , enough to overwhelm the low frequency noise (see Section 13.6.10).

Aside: FM Preemphasis. In order to maximize the received SNR for a given signal bandwidth, the spectrum of the baseband signal should be white, that is, flat up to its maximum frequency. The spectrum of music is concentrated at low to mid audio frequencies, so FM broadcasting uses *preemphasis*: an RC lead-lag network boosts the high audio frequencies before modulation, and a matching *deemphasis* network after the receiver's demodulator attenuates them again, *along with the high frequency noise*. This is an example of *whitening*, a widely applicable method of improving SNR when the noise is white and the signal isn't—see Section 13.8.10.

13.4 AMPLIFIERS

Amplifiers, as we all know, are devices for making signals bigger. From a signal processing point of view, we want them to do that while remaining nice linear, time-invariant networks, and perhaps having other attributes. Design details are found in Section 14.6.7 and most of Chapter 18. Here we'll stick with their behavior.

13.5 DEPARTURES FROM LINEARITY

Most of the time, the nonlinearities of our networks will be small, so that they can be modeled as perturbations to the behavior of an ideal linear network. The amplitude response of the network can then be written as a perturbation power series in V (valid near $V = 0$),

$$V_{\text{out}} \approx \sum_{n=0}^{\infty} d_n V_{\text{in}}^n, \quad (13.17)$$

which is called the *distortion polynomial*, probably because nobody ever uses more than five terms. It must be underlined that for a real network, the d_n are not fixed in value but depend on the signal frequencies and other things, and that for sufficiently large V , the network will become nonlinear and time varying enough that this series is invalid. The value of the distortion polynomial is that it predicts the amplitude dependence of different distortion products. These products are called *spurs*, a fortunate abbreviation of “spurious” that also connotes small, sharp protuberances, which is what they look like on a spectrum analyzer. Spurs due to distorted modulation (e.g., driving an amplitude modulator too hard) are collectively known as *splatter*, which is what they look like on the analyzer when you shout into the transmitter microphone.

The best thing about RF signal processing is that spurious mixing products land at frequencies away from the desired output, so that they can be removed almost completely by filtering. This luxury does not exist at baseband, for example, in the output of a phase-sensitive detector. (Compression and cross-modulation can still spoil the party—read on.)

13.5.1 Harmonics

The best known consequence of nonlinearity in a network is harmonic generation. If the input is $A \exp(j2\pi ft)$, then the output will be

$$V_{\text{out}} \approx \sum_{n=0}^{\infty} d_n A^n e^{j2n\pi ft}, \quad (13.18)$$

so that the n th harmonic amplitude will go as A^n . Accordingly, a 3 dB change in A will produce a 6 dB change in the second harmonic, 9 dB in the third, and so on. Do remember that this only works in the weak distortion limit. This is because we have broken our rule about going back to the purely real representation before applying nonlinear operations. A trigonometric function raised to the n th power contains components at nf , $(n-2)f$, $(n-4)f$, . . . down to f or 0, depending on whether n is odd or even. Equivalently, the component at nf will depend on A^n , A^{n+2} , A^{n+4} , and so on. Still, since the d_n are not rapidly increasing, the small-signal leading behavior of the n th harmonic will be controlled by the term in A^n .

This is made even more complicated by biasing effects, which make the response asymmetric between positive and negative going peaks. To reflect this, there should notionally be a DC term added to the input signal before all the nonlinear operations, which will mean that the term in A^n will really contain contributions from order n and all higher orders, not just $n+2$, $n+4$, . . . This is part of the reason that the distortion polynomial is mainly a conceptual aid.

How prone a device is to produce harmonics is usually specified by its n th order *intercept points*, which are where the extrapolated power law curve for the fundamental ($n=1$) and n th harmonic cross, as shown in Figure 13.8 (this isn’t the most commonly quoted intercept point—see the next section).

13.5.2 Frequency Multipliers

Sometimes harmonics are useful, for example, to produce a reasonably stable high frequency signal from a very stable low frequency one. Devices that do this on purpose




Figure 13.8. Polynomial behavior of distortion products of a weakly nonlinear network. IP_2 and IP_3 are the second- and third-order output intercept points.

are called *frequency multipliers*. Any modulation on the signal being multiplied will get changed. The multiplier is strongly nonlinear, so it isn't easy to say in general what happens to amplitude modulation, except that it's nothing pretty. The behavior of any PM is much easier to predict: since the frequency is multiplied, the phase is too, and so a frequency n -tupler multiplies the modulation index by n but *leaves the modulation frequency unchanged*. For weak phase modulation (e.g., oscillator phase noise), the sideband power goes as n^2 . Frequency multipliers may exhibit AM–PM conversion (see Section 13.5.6), in which case the phase noise may be worse than this. A phase-locked loop (PLL) with a frequency divider between the VCO and phase detector is another sort of frequency multiplier. In frequency dividers, amplitude information is lost but the modulation index goes down by a factor of N . A divider introduces an N -fold phase ambiguity in its output since we can never be sure which of N states it wakes up in. This causes no problems in straight PLL multipliers (why?) but it does in M/N synthesizers, where both the reference and VCO frequencies are divided before the PD. Thus *direct synthesizers*, using multipliers and mixers, is preferred to PLL synthesizers whenever the absolute phase is important (see Section 13.9.4 for more).

13.5.3 Intermodulation

If V_{in} is not a pure tone but contains several frequency components, the distortion will include mixing products between all combinations of these components, collectively known as *intermodulation distortion* (IMD). Second-order intermodulation is not usually much of a problem, since these products are either near DC or near twice the signal frequency. As with harmonics, unless we're working at baseband, these are easily filtered out. On the other hand, third-order IM products at $2f_1 - f_2$ tend to land in your passband, so they're much more obnoxious.

Intermodulation performance is specified with intercept points, as harmonics are, but since there are more degrees of freedom here, the test conditions have to be carefully

specified. The two-tone test is the most common. Two closely spaced sinusoidal signals of equal amplitude are fed in to the input. The strengths of the n th order mixing products go as V^n , as before. The third-order intermodulation intercept point is where the extrapolated $2f_1 - f_2$ amplitude hits that of one of the desired signals (this is what is most commonly meant by IP_3). When there is a forest of these spurs, the power law dependence is very helpful in identifying the intermodulation order of each one. Exact identification, at least of the lower lying orders, is possible by varying the tone amplitudes separately over a narrow range.

Aside: Optical IMD. You even see this with lasers. If you shine a multiple longitudinal mode HeNe laser on a photodiode and look at the resulting spectrum, you will usually find sharp spurs that sweep aimlessly back and forth from near DC to a megahertz or so. These are fourth-order mixing products, caused by the interference of one mode with a third-order IM product of two others, all four signals being near 632.8 nm (474 THz). The frequency offset and instability result from slight dispersion in the plasma and the mirrors. (Optics people call this *four-wave mixing*.)

13.5.4 Saturation

As the drive level is increased, eventually the ability of the device to drive its output will be exceeded. In high frequency devices, this limit is quite gradual; they don't crash into the supply rail and stop the way op amps do. This behavior is usually summarized by the 1 dB compression point $P_{1\text{ dB}}$, which is the output power at which the fundamental gain G has dropped by 1 dB. This is typically about 10 dB below IP_3 . The compression behavior is governed by the next term in the distortion polynomial that gives rise to a signal at the fundamental, which (remembering the DC bias) is the A^2 term. At a level 10 dB below the 1 dB compression point, the compression is typically 0.1 dB, as we'd expect from a quadratic term. Below $P_{1\text{ dB}}$, where no serious waveform distortion occurs, the compression error δ is

$$\delta \equiv 1 - \frac{P_{\text{out}}}{GP_{\text{in}}} \approx \frac{P_{\text{out}} \ln 10}{10P_{1\text{ dB}}}, \quad (13.19)$$

and the compression factors $(1 - \delta)$ due to cascaded devices multiply. To the same order, the deltas add, so we can define the *compressibility coefficient* CC of a given element in a system as

$$CC = \frac{G_{\text{tot}}}{P_{1\text{ dB}|_{\text{out}}(\text{mW})}}, \quad (13.20)$$

where G_{tot} is the total small-signal power gain (in W/W) from the input of the system up to the output of that element. The overall compression error is then

$$\delta_{\text{tot}} \approx 0.23 P_{\text{in}} \sum_i CC_i \quad (13.21)$$

and the input 1 dB compression point of the system is

$$P_{1\text{ dB}|_{\text{in}}} = \frac{1}{\sum_i CC_i}. \quad (13.22)$$

(For devices such as mixers, where the compression point is specified at the input, use the total gain up to the input instead.) If the input level is turned up high enough, the output level will stop increasing altogether, and the amplifier will become a limiter (see Section 14.7.16). Most don't limit well unless they're specifically designed to.

13.5.5 Cross-Modulation

As with all mixing products, the phases of the spurious signals are additive (since the frequencies add, the phases must too), and the amplitudes are multiplicative. This means that if the input consists of different frequency components, the modulation of the stronger ones will be transferred to the other components.

This most obnoxious behavior is called *cross-modulation*. An amplitude modulated signal mixing with itself produces a baseband product that approximately reproduces the modulation waveform.[†] If this product in turn mixes with a component at f_2 , the result will be an amplitude modulated spurious product with a carrier at f_2 , right on top of the desired component. Thus the desired component will acquire the AM of the interfering signal. This is physically obvious, since if the peaks of one signal are strong enough to reduce the gain of the device, the resulting compression will affect other signals in the input as well.

13.5.6 AM–PM Conversion

One kind of distortion that is not an obvious consequence of the distortion polynomial is AM–PM conversion. Saturating a network changes the time-averaged device parameters (e.g., the phase of the current gain), so that the time delay through the network changes too. Amplitude modulation of the signal will cause the delay to vary with the modulation waveform, thus causing phase modulation. This effect is particularly strong in limiting amplifiers, used for FM and PM IF strips, where it is frequently the largest source of error. Since a limiter is intended to be operated in a strongly nonlinear region, the obvious solution of reducing the signal level to reduce the AM–PM conversion is not applicable. You just have to pick amplifier designs that don't slow down so much when they go into limiting, such as FET differential amplifiers with low impedance drain loads, and reduce the frequency at which the limiting occurs to as low a value as practicable, so that a delay τ produces a smaller phase shift. This is one instance in which a double-conversion system makes good sense.

13.5.7 Distortion in Angle Modulated Systems

Because the information in angle modulated systems is encoded in the phase, it is highly insensitive to amplitude distortion, allowing the use of amplitude limiting in the receiver to suppress AM noise, interference, and signal fading. This doesn't qualify as a free lunch; the information is encoded in phase changes, so phase nonlinearity causes the same sorts of trouble as amplitude nonlinearity in AM systems. (This is obvious from a time-domain point of view.) In the frequency domain, phase nonlinearity causes the carrier and sideband phasors not to sum to a constant length. After limiting, the resulting AM information is lost, so the undistorted phase information cannot be recovered by any

[†]It isn't perfect—signal \times carrier reproduces the modulation, but signal \times signal is a quadratic nonlinearity.

equalizing filter. Even a linear-phase filter can cause distortion by rejecting sidebands. FM sidebands depend nonlinearly on the modulating signal, so each one by itself contains intermodulation products between different baseband frequencies that only cancel out when all sidebands contribute to the detected signal.

Aside: Signal Detection in Noise.

If your experiment needs statistics, you ought to have done a better experiment.

—Ernest, Lord Rutherford

Rutherford was a great enthusiast for clear, cheap, clever experiments. Following this piece of advice would render many sorts of experiments impossible even in principle, but it holds an important point: don't passively accept poor quality data. Spend lots of effort on improving the quality of the data, with the aim of achieving a beautiful, clean result. When that is done, *then* worry about statistical analysis. A properly designed statistical signal processing system can often yield an extra few decibels in sensitivity, which is well worth having, but only after the last decibel has been wrung out of the front end. In any case, don't muck about with parameters whose physical significance you can't explain in less than a paragraph.

13.6 NOISE AND INTERFERENCE

We live in a fallen world, so the signals we process are never free of noise, distortion, and extraneous interfering signals. We must be forearmed against them by knowing what they do to our measurements, and how.

13.6.1 White Noise and $1/f$ Noise

Noise whose power spectral density (PSD) is constant with frequency is said to be *white*, by analogy with light. If you look at the baseband output of some system on a spectrum analyzer when no input signal is present, most of the time you'll see the noise rise to a big peak at DC. Above there, usually, there's a white noise region known as the *flatband*, followed by a rolloff toward high frequency. There may also be a noise peak between the flatband and the rolloff, as we'll see in Section 18.4.2.

The low frequency peak is partly due to the finite bandwidth of the spectrum analyzer smearing out the DC offset, but it's partly real. A system with any DC offset whatever notionally has a component of $\delta(f)$ in its spectrum, and often exhibits a very steep rise in noise near 0. The two major components of this near the low edge of the flatband are drift, which has a characteristic $1/f^2$ dependence in its power spectrum, and a motley collection of pops, clicks, and bangs called $1/f$ noise from its frequency dependence. (Since the real frequency dependence is often far from $1/f$, it's also known as *flicker noise* or, less descriptively, *excess noise*.) Both are statistically nonstationary and generally very badly behaved. Dialing up the time constant on your lock-in amplifier and taking longer to do the measurement will not help these sorts of noise, because for a fixed number of data points, the low frequency cutoff scales with the high frequency one; in that case the variance of $1/f$ noise is independent of the upper frequency limit, and for anything steeper, the variance gets *worse* as the bandwidth goes down. (This is obvious in the case of drift—averaging a linear ramp for a longer time just takes you further from home.) Many a graduate student has come to grief by not knowing this.

Flicker noise comes from sources like contamination migrating around on the surfaces of ICs, conductance and carrier density fluctuations in carbon resistors, electromigration of metal atoms in conductors or interstitial atoms in silicon, and a million more. Because these are conductance variations, $1/f$ noise only shows up when a current is flowing through the noisy part, and the noise power density is proportional to i^2 , just as signal power is, rather than i as with shot noise. Migrating contamination, often due to a cracked IC package or a manufacturing defect, shows up as small jumps in offset voltage, known as *popcorn noise* or *telegraph noise*. Flicker noise is a bear to deal with, since the system will sometimes just sit there being well behaved, only to get noisy again minutes or hours later. It is a strong sign of an unreliable component.[†]

13.6.2 Johnson (Thermal) Noise

A beautiful result of classical thermodynamics is the *fluctuation–dissipation theorem*. It states that any process that can dissipate energy exhibits thermal fluctuations, and tells how to calculate them. What that means to us is that circuit elements that can dissipate electrical power as heat (e.g., resistors and lossy capacitors) are also noise sources. A resistor R at temperature T , connected to a matched load (i.e., another resistor whose value is also R) will transfer noise whose average power in a 1 Hz bandwidth is

$$p_N = kT. \quad (13.23)$$

As usual, we can calculate the total noise power P_N by integrating the noise power per hertz p_N times the squared modulus of the transfer function of the network. For flat-topped frequency responses, this basically amounts to multiplying p_N by the noise bandwidth in hertz.[‡]

Since the resistor is a linear device, the noise can be modeled with a Thévenin or Norton model,[§] that is, as a noise voltage source v_N in series with a noiseless R , or as a noise current source i_N in parallel with a noiseless R , as shown in Figure 13.9. If you open-circuit a Thévenin model, the rms open circuit noise voltage v_N is twice that measured under matched conditions:

$$e_N = \sqrt{4kTR}. \quad (13.24)$$

Similarly, if you short-circuit a Norton model, the short-circuit current is twice the matched value:

$$i_N = \sqrt{\frac{4kT}{R}}. \quad (13.25)$$

[†]If you find a part that produces unusually bad flicker noise, *throw it away*, and seriously consider pounding it into dust with a ball-peen hammer first, so it can never fool you again. (If you find a whole lot of parts like that, save them to pound the manufacturer with.)

[‡]This is in the analytic signal picture: classical equipartition predicts that in the two-sided spectrum, the uncertainty in a 1 second measurement in a 1 Hz bandwidth is $kT/2$, just as in every other classical degree of freedom—the factor of 2 comes from the analytic signal conversion.

[§]Horowitz and Hill have more on Thévenin and Norton models. If you're reading this chapter because you need to design or understand the guts of how low noise circuits work, and you don't own a copy of *The Art of Electronics*, then do not walk but run to get a copy.




Figure 13.9. Noise models of a resistor: (a) Norton and (b) Thévenin. The two are exactly equivalent, so pick whichever one makes the algebra easier.

In a reactive network, the reactances are noiseless, so the open-circuit noise is that of the resistive part of the total impedance,

$$V_N^2 = 4kT \int_0^\infty R df. \quad (13.26)$$

Equivalently, shorting out a parallel network makes the reactive currents zero, so the short-circuit noise current is that of the conductance,

$$I_N^2 = 4kT \int_0^\infty G(f) df. \quad (13.27)$$

Example 13.2: Noise of Resistors in Series and Parallel. Two resistors R_1 and R_2 wired in series behave electrically as a single resistor of value $R_1 + R_2$. We expect their combined noise properties to be the same as well. Choosing the Thévenin equivalent noise circuits, the series resistance is indeed $R_1 + R_2$, and the two noise voltages add in rms fashion (power); thus the total noise is

$$\begin{aligned} v_{N\text{tot}} &= \sqrt{v_{N1}^2 + v_{N2}^2} \\ &= \sqrt{4kT(R_1 + R_2)}. \end{aligned} \quad (13.28)$$

Similarly, using the Norton model for resistances in parallel, the noise currents add in rms fashion, so that the total is

$$\begin{aligned} i_{N\text{tot}} &= \sqrt{i_{N1}^2 + i_{N2}^2} \\ &= \sqrt{4kT(1/R_1 + 1/R_2)}, \end{aligned} \quad (13.29)$$

both of which are just what we expect. A note to those fond of negative impedance converters: negative resistances don't have imaginary noise.

Example 13.3: kTC Noise. This is an alternative derivation of Johnson noise and is helpful in nonlinear situations where R may be changing in unknown ways with time (e.g., due to switching). In a parallel RC network, the electrical connection couples the

capacitor's degree of freedom (i.e., its stored energy) to the temperature reservoir, so we know from classical equipartition that the ensemble average energy on the capacitor is $\frac{1}{2}kT$. Since the average stored energy $W = \frac{1}{2}CV_N^2$, the rms noise voltage V_N must be $(kT/C)^{1/2}$, or equivalently the rms charge noise $\Delta Q_N = (kTC)^{1/2}$. This drops nicely out of (13.24), since the resistance of a parallel RC is $R/[1 + (\omega RC)^2]$:

$$\langle V_N^2 \rangle = 4kT \int_0^\infty \frac{R}{1 + \omega^2 R^2 C^2} \frac{d\omega}{2\pi} = \frac{2kT}{\pi C} \int_0^\infty \frac{du}{1 + u^2} = \frac{kT}{C}. \quad (13.30)$$

Since R and C can be anything, and we can take sums and differences of different bandwidths, this applies in any bandpass; thus the noise must be flat. Since an RC lowpass has a noise bandwidth of $1/4RC$, the spectral density contributed by the resistor is

$$v_N = \sqrt{(kT/c)(4RC)} = \sqrt{4kTR}. \quad (13.31)$$

This uncertainty applies whenever we're resetting the charge on a capacitor with a feedback-free network such as a switch or CMOS gate; since amplifiers need not have noise temperatures near their physical temperature, a feedback network can do better in general. This kTC noise dominates the readout noise of CCDs, because the charge on their readout stages is dumped on each cycle, a problem solved by *correlated double sampling* (see Section 3.9.4).

13.6.3 Shot Noise in Circuits

Most resistors don't exhibit shot noise; metal film resistors especially are known to be very quiet (see Section 3.10.2). Currents in which carriers arrive more or less independently (e.g., via tunneling or photoionization) have full shot noise. These include base, emitter, and collector currents of BJTs, where the conduction is by minority carriers. FETs generally do not have full shot noise, but their other noise sources generally dominate, making FETs noisier than BJTs except in high impedance circuits.

13.6.4 Other Circuit Noise

There are lots of other sources of junk in signals, for example, $1/f$ noise due to conduction fluctuations in carbon and thick-film resistors; microphonics due to high- k ceramic capacitors becoming piezoelectric; jiggling cables with DC bias on them; capacitive and inductive pickup; popcorn noise from electrolytic capacitors; and wideband noise and spikes on power supplies and voltage references. Those ones you just have to foresee and design out. When you're learning, you usually find them during debug and remember them for next time.

13.6.5 Noise Figure, Noise Temperature, and All That

At low frequencies, it is not difficult to measure voltage and current separately. At RF, on the other hand, everything becomes inductive or capacitive. High impedances are difficult to maintain; a good rule is that in a lumped-constant (i.e., not transmission line) circuit, life begins to get difficult at impedance levels of about $100 \text{ k}\Omega/f(\text{MHz})$ (corresponding

to a node capacitance of 1.6 pF). The outputs of things talk to their inputs, and so forth. It's just a more complicated regime.

Reactive mismatches make the noise behavior of high frequency circuits more complicated. The noise of an RF amplifier is usually quoted as a *noise figure*, which is simply how much more noise (in dB) this amplifier puts out versus a noiseless amplifier, if both of their inputs are terminated with a 50 Ω resistor at $T = 300$ K. (It's made more confusing by the fact that although the maximum power is coupled to the amplifier when the source is impedance matched, this does not result in the best SNR because the amplifier's noise usually is not a simple Thévenin model characterized by a single temperature.)

Consider an amplifier with power gain G . If we notionally hook up a noiseless 50 ohm resistor to its input, it will produce a 1 Hz output noise power $p_{N\text{out}}$. This allows us to define its intrinsic 1 Hz input noise power with a matched load, $p_{N\text{in}}$, as

$$p_{N\text{in}} = \frac{p_{N\text{out}}}{G}. \quad (13.32)$$

A resistor at temperature T has Johnson noise power of $p_{N\text{th}} = kT$ in 1 Hz, which allows us to define the noise figure NF:

$$\text{NF(dB)} = 10 \log \left(1 + \frac{p_{N\text{in}}}{kT} \right) \Big|_{T=300\text{K}}. \quad (13.33)$$

We can't actually get a noiseless resistor, but we can do pretty well by doing measurements with the input resistor at different temperatures and using curve fitting to get G and $p_{N\text{in}}$. Other techniques using calibrated noise sources are available, so it isn't always necessary to use liquid nitrogen. See Section 2.5.4 for more on practical noise measurements.

The noise figure of a generic amplifier is around 3–4 dB, and a really good one is below 1 dB.[†] Such low noise figures are usually only needed when the source is in fact not a room temperature resistor. Low temperature sources include cryogenically cooled IR detectors, and antennas pointing away from the Earth and Sun. Photodiode networks are usually badly mismatched (otherwise we'd be wasting half our signal power inside the network), so they don't look like a room temperature resistor either.

For these low noise, oddball applications, the noise figure isn't too meaningful because it's always near 0 dB. Another more useful parameter is the noise temperature,

$$T_N = \frac{p_{N\text{in}}}{k}, \quad (13.34)$$

which is how far you'd have to cool your resistor before its noise equaled the amplifier's. Good low noise amplifiers have noise temperatures of 100 K or below, even though their physical temperature is 300 K. (In case you're worried about the fluctuation–dissipation theorem here, don't be—these devices are that quiet only when they're turned on; the fluctuation–dissipation theorem's assumption is a closed, isothermal system, which breaks down as soon as power is applied.)

Table 13.1 is a handy conversion chart of noise figure to noise temperature.

[†]Even vacuum tube amplifiers got below 2 dB, which is pretty good for something running at 1000 K. Check out the Miteq catalog for amps with guaranteed NFs of 0.35 dB ($T_N \approx 25$ K).

TABLE 13.1. Conversion from Noise Temperature to Noise Figure

T_N (K)	NF(dB)	T_N (K)	NF(dB)	T_N (K)	NF(dB)
650	5.0	100	1.25	25	0.35
500	4.3	78	1.00	15	0.21
400	3.7	75	0.96	10	0.142
300	3.0	50	0.67	5	0.072
200	2.2	37	0.50	1	0.0144

13.6.6 Noise Models of Amplifiers

The origins of current and voltage noise are physically diverse and may have widely different noise temperatures in a single device. These temperatures may or may not be close to junction temperature T_j . For instance, the base current shot noise of a BJT looks like the Johnson noise of a resistance of β/g_M at 150 K, whereas its extrinsic base resistance is a real resistance at T_j , so its voltage and current noise are predicted by the Johnson noise formula. Similarly, the collector current shot noise becomes an input-referred voltage noise contribution of $\sqrt{2eI_C}/(\beta g_m)$ (see Example 18.2). Since base and collector shot noise are uncorrelated, all of these contributions add in power. On the other hand, the base shot noise current times the extrinsic base resistance produces a noise voltage perfectly correlated to that part of the current noise. Most of the time these correlations are small, and when they're not small enough to ignore, a simple physical model beginning with uncorrelated noise contributions is perfectly adequate. Dostal gives the full analysis for difficult cases.

The effective noise temperature of an amplifier in a circuit may be nowhere close to ambient. For instance, an OP-27 op amp has $v_n = 3 \text{ nV/Hz}^{1/2}$, $R_{in} = 6 \text{ M}\Omega$, so $T_N = v_n^2/(4kR) = 22 \text{ mK}$ for voltage noise, and $i_N = 0.4 \text{ pA/Hz}^{1/2}$, so $T_N = i_N^2 R/(4k) = 17,400 \text{ K}$ for current noise. Noise temperature is thus a sensitive function of termination resistance and feedback components, and not a simple Thévenin or Norton model as with combinations of resistors. In addition, noise current and noise voltage behave differently in the circuit; noise current is a real current that will flow out of the amplifier terminals and will affect other circuit elements, as when we parallel amplifiers to get independent noise measurements in Section 17.11.6: the current noise contributions add but the voltage noise remains independent.

Example 13.4: Noise of an Inverting Amplifier. Let's do a concrete example: a CMOS op amp inverting amplifier, as shown in Figure 13.10. The LMC6032 is a jellybean 15 V CMOS dual op amp that typically has a 1.4 MHz gain–bandwidth product and 1.1 V/ μ s slew rate. Its specifications are not exactly ironclad; its data sheet has exactly one (1) guaranteed parameter: the slew rate is guaranteed to be at least 0.8 V/ μ s at 25 °C and at least 0.4 V/ μ s from –40 °C to +85 °C. Not the first choice part for building space probes, but the price is right: under 40 cents each in quantity. At 1 kHz, the LMC6032 typically has 1-Hz voltage and current noise of 22 nV and 0.4 fA.

The noise contributions add as follows. The voltage noise is differential, so we can notionally put it in the noninverting input, where it obviously sees the total noninverting gain of the stage, which is $1 + |A_{V_{cl}}|$. Since feedback is holding V_{in-} at 0, there's ideally no voltage drop across R_I , so all current noise goes out through R_F , contributing $R_F i_N$. (All the current coming into the – input gets summed and multiplied by R_F , which

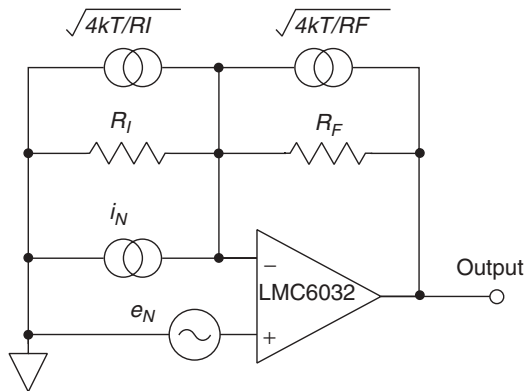


Figure 13.10. Noise model of an LMC6032 inverting amplifier.

is why it's called the *summing junction*.) The resistor noise contributions are calculated using the Norton model, so they contribute $\sqrt{4kT/R}$. The total current noise is multiplied by R_F and summed with the v_N contribution (all sums are RMS of course). Figure 13.11 shows the results, using the typical e_N and i_N at 1 kHz.

As expected, e_N dominates for small R_F and resistor noise for intermediate values. At very large R_F , the amplifier's current noise starts to dominate, but note that it's much less important at higher voltage gain, because e_N gets multiplied by $1 + |A_V|$ and the smaller $R_I (= R_F/|A_V|)$ contributes more noise current. The noise of the inverting ($A_V = -1$) amplifier is the same as that of a noninverting amplifier of $A_V = +2$, because for noise purposes the circuit is exactly the same: R_I and the $+$ input are both grounded. Thus at $|A_V| = 1$, the available dynamic range is 6 dB less in the inverting configuration (see Section 15.11).

13.6.7 Combining Noise Contributions

It is important to remember that these v_N and i_N are RMS values of random fluctuations. Noise sources whose physical origins are diverse produce fluctuations that are completely statistically independent. They *never* cancel, unless derived from the same source at the same moment; when we add or subtract two voltages or currents, the noise contributions from each source always add in RMS fashion.

$$v_N^2 = v_{N1}^2 + v_{N2}^2. \quad (13.35)$$

Circuit designers and textbook writers always toss around noise sources like that, connecting purely notional noise voltage and current sources as though what they were doing had some physical significance, which it may not seem to. Beginners are understandably hesitant about doing this. While this caution is commendable, there are two reasons why the circuit designers' approach is the right one: first, all these noise contributions really do behave that way; and second, it isn't obvious how else to proceed. (Besides, as Figure 13.11 shows, usually one noise source dominates.)

The deeper reasons for the success of this approach are that the circuits we're interested in are operating in a highly linear regime, by design, and that the physical origins of these




Figure 13.11. Noise performance of inverting amplifiers made from a jellybean LMC6032 op amp, calculated from typical specs: (a) $A_V = -10$ and (b) $A_V = -1$.

noise sources are diverse. Besides the naturally good linearity of op amps, transistors, and resistors, the small amplitude of the noise sources and the linearizing action of negative feedback ensure that the effect of all the noise sources combined will be precisely the rms sum of their individual effects. There are cases in which the correlations of the voltage and current noise are important, for example, when using huge resistors with bipolar op amps, so that the bias current noise is converted to a voltage, but the diversity of the fundamental noise sources means that a sufficiently low level circuit description will produce these correlations properly at the output. Since the output is almost always at a very low impedance, either in a 50 ohm system or an op amp output, it is unnecessary to separate voltage and current noise contributions there.

Aside: Noise Correlations. The fully general version of (13.35) includes a correlation term,

$$v_N^2 = v_{N1}^2 + v_{N2}^2 + 2Cv_{N1}v_{N2}, \quad (13.36)$$

where C is defined as

$$C = \langle v_{N1}(t)v_{N2}(t) \rangle / \sqrt{v_{N1}^2 v_{N2}^2} \quad (13.37)$$

and can have any value between -1 and 1 (we write $v(t)$ for instantaneous values and v for ensemble-averaged RMS values). Fully correlated noise ($C = 1$) could come from two cables connected to the same noise source, and nearly fully anticorrelated noise could come from opposite ends of a transformer winding whose center tap is grounded, or from the collectors of a differential BJT pair. Usually if you calculate assuming the noises are all uncorrelated (i.e., $C = 0$), and keep your eyes peeled for situations like these, where simple physics says they have to be correlated, you'll get the right answers.

13.6.8 Noise of Cascaded Stages

Calculating the noise of a cascade (several circuits in a row) is simple. Assuming the circuits are linear, the total noise is the RMS sum of the individual contributions. If there are M stages with noise voltage v_{Ni} , the output noise is

$$v_N^2 = \sum_{i=1}^M v_{Ni}^2 G_{ii}^2, \quad (13.38)$$

where G_{ii} is the total gain from the input of stage i to the output, that is, $G_{ii} = \prod_{j=i}^M G_j$, where G_j is the voltage gain of the j th stage.

Aside: RMS Noise Measurements. Not every instrument does a good job of measuring RMS noise. Many good ones, such as the HP 400EL AC voltmeter and your average lock-in, measure the average value of the rectified signal, and then (perhaps) apply a correction factor to get RMS. The problem is that the correction factor is correct for sine waves but wrong for noise (and everything else): the RMS/average ratio of a rectified sine is $\sqrt{2}$, whereas for Gaussian noise it's $\sqrt{\pi}$, which is 0.98 dB higher. So when doing noise measurements with an average-reading instrument, remember to add 1 dB. If you can get hold of a HP 3400A RMS voltmeter, it's the bee's knees for noise measurement: it has a 10:1 crest factor (i.e., maximum peak/RMS ratio). True RMS meters like the 3400A also get the right answer with more complicated signals, where the fudge factor technique fails.

13.6.9 Interference: What Does a Spur Do to My Measurement, Anyway?

An interfering signal at the same frequency as our signal causes beats,

$$A(\exp(j2\pi ft) + \epsilon \exp[j(2\pi ft + \theta)]) = A \exp(j2\pi ft)(1 + \epsilon \cos \theta + j\epsilon \sin \theta), \quad (13.39)$$

which shift its amplitude and phase,

$$\begin{aligned}\frac{\Delta A}{A} &= \sqrt{(1 + \epsilon \cos \theta)^2 + (\epsilon \sin \theta)^2} - 1 \\ &= \sqrt{1 + \epsilon^2 + 2\epsilon \cos \theta} - 1 \\ &\approx \epsilon \cos \theta\end{aligned}\quad (13.40)$$

and

$$\begin{aligned}\Delta \phi &= \arctan \left(\frac{\epsilon \sin \theta}{1 + \epsilon \cos \theta} \right) \\ &\approx \epsilon \sin \theta.\end{aligned}\quad (13.41)$$

Thus a weak sinusoidal interference term, ϵ times as large as the carrier, causes a maximum amplitude change of ϵ times, and a maximum phase perturbation of ϵ radians. If the two signals are not at the same frequency, the phase θ is not constant, but the time-dependent amplitude and phase perturbations are given by the same formulas. These estimates allow spurious responses to be budgeted based on required system performance.

13.6.10 AM Noise and PM Noise

If the unwanted signal is not a pure sinusoid but a noise waveform, we cannot predict its exact time dependence without knowing that of the noise. We have to settle for predicting the statistical uncertainty of the amplitude and phase caused by the additive noise. This is greatly simplified by the linear dependence of ΔA and $\Delta \phi$ on the (small) relative amplitude ϵ of the interfering signal, which allows us to ignore the effects of different modulation components intermodulating with each other.

This means that for a given type of modulation, the performance of different demodulation schemes tends to be very similar when the SNR is large, and that the performance depends mainly on the strength and character of the noise and the bandwidth of the narrowest filter.

When the SNR becomes small, this convenient approximation is invalid; the phase and amplitude shifts are nonlinear and coupled. Furthermore, the behavior of different demodulation schemes can be very different then, so that extra thought and simulation are needed. For example, a phase-locked loop FM demodulator will lose lock, a limiter will start to suppress the signal in favor of the noise, and a diode AM detector will detect mainly the noise peaks. This is unfortunate, as the ultimate sensitivity of a measurement is defined by how well it handles the low SNR case.[†]

If the noise is filtered additive Gaussian noise, such as thermal or shot noise, then each frequency component is independent of all others, and also of the signal. The phase θ is therefore a uniformly distributed random variable, and the effects due to noise at different frequencies add in power. Thus we can get the statistics of amplitude and phase by taking the rms magnitudes of ΔA and $\Delta \phi$, using (13.40) and (13.41). Doing so, we get

$$\frac{\langle \Delta A \rangle}{A} = \frac{1}{\sqrt{2}} \sqrt{\frac{P_N}{P_C}} \quad (13.42)$$

[†]Hint: For AM and angle modulation of low m , use a synchronous detector; for wideband FM, use a PLL with wideband AGC and no limiter.

and

$$\langle \Delta\phi \rangle = \frac{1}{\sqrt{2}} \sqrt{\frac{P_N}{P_C}}. \quad (13.43)$$

We see that the quality of our measurement is determined by the quantity P_C/P_N , the power *carrier to noise ratio* (CNR). This is slightly different from the SNR, because what we consider signal may be the sidebands rather than the carrier. Because of the importance of this ratio, the noise and interference are usually specified in decibels with respect to the carrier (dBc). Noise and spurs at -40 dBc total will produce rms phase errors of 7 milliradians, and an rms amplitude uncertainty of $\pm 0.7\%$.

Because these contributions enter linearly, the amplitude statistics of small AM and PM fluctuations are the same as those of the noise itself, up to a scaling. Thus weak additive Gaussian white noise generates amplitude and phase fluctuations whose statistics are also Gaussian. This is helpful in predicting the response of a detector to the noisy signal: if you put it into an AM detector, you see the AM noise; if you put it into a phase detector (running near null), you see the PM noise. PM is also the natural description for oscillators whose frequency is not constant, because weak white noise fed into a phase modulator produces flat-topped noise sidebands, like AM but unlike FM. Thus frequency fluctuation in oscillators is generically known as *phase noise*.

An FM measurement with fixed signal power and constant noise PSD gets quieter as m increases. The total phase shift (and hence the demodulator output voltage) grows as m , as does the required bandwidth, so that the noise power is proportional to m ; this means that the total phase noise power goes as m . Since the signal power goes as m^2 , the SNR at the demodulator output will go as m , at least until we leave the high-SNR limit.[†]

Aside: AC Versus DC Again. The assumption here is an AC measurement—once again, you can't phase-shift DC, so all the noise power winds up in the amplitude in a DC measurement. We already knew that, of course, because the noise power is precisely the RMS uncertainty in a measurement of the instantaneous signal power. This is another example of the effect of sideband folding, or alternatively of quoting noise in SSB or DSB terms. (See Section 13.1 for more on AC versus DC.)

13.6.11 Additive Versus Multiplicative Noise

The most basic distinction between noise sources is whether they depend on the level of the signal. Thermal noise remains the same (in $\text{W/Hz}^{1/2}$) whether there is a signal present or not; the noise waveform is just added to the signal. With constant DC photocurrent (e.g., in a heterodyne interferometer or strongly background limited system), shot noise is additive as well. Here the signal-to-noise ratio goes up linearly with signal power.

Multiplicative noise comes from noise modulating the signal source; examples are laser residual intensity noise (RIN) and the phase noise of an oscillator. In oscillators, the low frequency noise of the active device and the resonator, together with external effects such as power supply ripple, will modulate the amplitude and phase of the signal.[‡]

[†]After all our discussion of minimizing measurement bandwidths, this should strike you as odd. It is. (It's closely related to the SNR improvement from pulsed measurements in Johnson noise.)

[‡]While these are theoretically multiplicative, if their main effect is modulating the pedestal and not the signal, you can think of them as additive.

The effects of multiplicative noise get stronger (in $W/\text{Hz}^{1/2}$) as the signal gets stronger, so that the SNR is constant with signal level when multiplicative noise dominates. Multiplicative noise is frequently very obnoxious, since it is systematic in nature, so that the usual noise reduction methods based on averaging of one sort or another (e.g., narrow filters, lock-ins, and signal averagers) often fail. This is aggravated by the tendency of this sort of noise to be strongest at low modulation frequencies, as in baseband $1/f$ noise, which we've already encountered. Multiplicative noise is a form of unwanted modulation that puts noise sidebands on your signal (see Section 13.3.4), so it follows you wherever you go.

13.6.12 Noise Statistics

Noise whose average properties do not change with time is said to be *statistically stationary*. Since in practice most properties of noise waveforms are measured via time averaging, this statement may not seem to have much content—how could an average over all time change with time? The answer is that the properties of noise, such as its rms amplitude, instantaneous probability distribution function, autocorrelation, and so forth, are usually defined in terms of *ensemble averages*. That is, if we took a very large number of identical, noninteracting noise sources, we could measure the instantaneous noise voltage $V_N(t)$ at some time t from every one. From these measurements, we could construct a histogram of V for that t , and so on. This is not quite as silly as it sounds; there are many processes in which this could be done straightforwardly, such as imaging with a nearly uniformly illuminated CCD, where the signals from each pixel form an ensemble (after correcting for device nonuniformities of course).

Example 13.5: Ensemble Averages in Real Life. Two dozen solar telescopes pointing at the solar corona would have outputs consisting of the instantaneous coronal output at each point in the image, plus their own uncorrelated shot noise and technical noise. The coronal component would be highly correlated between receivers, in fact essentially identical apart from time delays depending on their position in space relative to the source. The technical noise would be statistically independent provided it did not depend on the signal level, variations in the common AC power line, local interfering signals, or those sorts of things. The outputs of the telescopes would be different functions of time, but drawn from the same *random process*. A random process is an ensemble of functions representing all possible outcomes of a given experiment.

The corona is far from time invariant, so that the statistics of the signal would be far from time invariant, but that doesn't bother the ensemble average.

More familiar noise sources, such as shot noise and Johnson noise, have statistics that depend on one or more parameters. Shot noise depends on the average current, and Johnson noise on the temperature and resistance. Time variations of these parameters will cause the noise to be statistically nonstationary. This is important in cases where the signal modulation is strong, for example, an interferometer whose beams are nearly equal in strength.

Noise whose time-averaged statistics are equal to its ensemble-averaged ones is said to be *ergodic*, and this class includes most fundamental noise sources, such as shot noise, Johnson noise, and so forth, providing that any parameters are time invariant. This is good, since it makes a lot of calculations and measurements easier, but bad in that it

leads to confusing distinct but related concepts such as statistical autocorrelation and time autocorrelation.

Aside: Time and Statistical Autocorrelations. The time autocorrelation of any signal is

$$g \star g = \int_{-\infty}^{\infty} g(t')g(t'+t)dt'. \quad (13.44)$$

It is the inverse Fourier transform (in time) of the signal's power spectrum, a fact that can be proved from the convolution theorem in about three lines of algebra. This applies to any waveform whatever, assuming the integrals converge. There are different ways of normalizing it, but we'll ignore that fine point.

The statistical autocorrelation is a somewhat more slippery concept: it's still related to the average value of the function multiplied by a shifted replica of itself, but now it's an ensemble average over many different functions drawn from the same random process,

$$\text{Corr}(g, g, t, \tau) = \langle g(t)g(t + \tau) \rangle. \quad (13.45)$$

Its Fourier transform is the ensemble-averaged power spectrum of the random process, but this is a much deeper fact, requiring the famous Wiener–Khinchin theorem, as well as more attention to normalization than we've given it.[†]

13.6.13 Gaussian Statistics

Noise sources such as shot noise and thermal noise arise from the random motions of very many charge carriers. In cases where their motions are not correlated with each other, the central limit theorem of statistics suggests that the instantaneous noise pulses contributed by each electron will sum up into a Gaussian probability density of voltage or current,

$$\frac{dP(V)}{dV} = \frac{1}{\sqrt{2\pi}V_{\text{RMS}}} \exp \frac{-(V - \langle V \rangle)^2}{2V_{\text{RMS}}^2}, \quad (13.46)$$

where V_{RMS} is the RMS voltage, considered as an ensemble average, and $\langle V \rangle$ is the mean voltage. Integrating $P(V)$ from V_0 to ∞ (10 standard deviations is infinite enough for almost everything) gives the probability of measuring a voltage greater than V_0 .

$$\begin{aligned} P(>V_0) &= \frac{1}{\sqrt{2\pi}V_{\text{RMS}}} \int_{V_0}^{\infty} e^{-(V-\langle V \rangle)^2/2V_{\text{RMS}}^2} dV \\ &= \frac{1 - \text{erf} \left(\frac{V_0 - \langle V \rangle}{\sqrt{2}V_{\text{RMS}}} \right)}{2} \end{aligned} \quad (13.47)$$

For signal detection in noise, we're usually interested in finding all separate events larger than some given threshold level. The quantity $P(>V)$ is the relevant one in

[†]Even some of the big guys confuse the Wiener–Khinchin theorem and the autocorrelation theorem. This is a symptom of really not understanding the difference between deterministic and stochastic processes, so if you find some author doing it, keep your powder dry.

situations where you're sampling at predetermined times, and where your sample interval is large compared with the width of the impulse response of your narrowest filter (otherwise adjacent samples will not be uncorrelated).

If you're sampling more rapidly than that, or you have a continuously operating threshold circuit (e.g., a comparator), what you care about instead is the rate of upward going threshold crossings per second. This is a bit more complicated (see Papoulis for details) and depends on the details of the power spectrum of the noise, being proportional to the square root of the second derivative of the autocorrelation of the noise,

$$\lambda_{\text{TC}}(x) = \frac{1}{2\pi} \sqrt{\frac{-d^2 R/dt^2}{R(0)}} e^{-x^2/2} \quad (13.48)$$

(where $x = V_0/V_{\text{rms}}$), which may not seem too useful since that parameter is not exactly at the tips of our tongues. (We could find it easily enough from the transfer function and the derivative formula for Fourier transforms).

The reason it's not a simple function of bandwidth is that a single noise event will tend to resemble the impulse response of the filter. A narrow bandpass filter has noise events that look like tone bursts, and so have many nearly identical peaks. The threshold crossing rate λ_{TC} is then many times higher than the true arrival rate of events. If we want only the first one of the burst, we're really asking for the upward going threshold crossing rate of the envelope of the analytic signal, which is a somewhat different matter. Since we typically need this number accurate only to the $\pm 20\%$ level, we can approximate it by assuming that our filter has a noise bandwidth B , and that the threshold crossing rate of the envelope is the same as if we were using a lowpass of noise bandwidth B . Using a brick-wall lowpass filter of bandwidth B , our upward going threshold crossing rate becomes[†]

$$\lambda_{\text{TC}}(x) = \frac{B}{\sqrt{3}} \exp(-x^2/2), \quad (13.49)$$

which is a very useful result.[‡] Note that if we use a thresholding approach with tone burst signals, or with bandpass filtered noise, we'll need to put in a dead time of a few times the width of the filter impulse response, or we will get multiple triggers per event.

This function falls off astoundingly rapidly with V , as shown in Table 13.2. The last two rows show a decrease of nearly 1000 times for a 1.3 dB change in the threshold voltage. Because Gaussian noise falls off so rapidly at large multiples of V_{RMS} , at some point other noise sources that are less well behaved are going to dominate. Knowing where your nice Gaussian histogram stops falling off and starts being dominated by outliers is a matter of some concern in thresholding applications.

The combined statistics of signal + noise are *not* Gaussian. This is another case where we assume the problem doesn't exist, and fix it by hand afterwards—the power and generality of the Gaussian noise analysis are not to be lightly put aside.

[†]For bipolar thresholds, where a positive going crossing of V_T and a negative going crossing of $-V_T$ are both accepted as events, this rate has to be doubled.

[‡]S. O. Rice, Mathematical analysis of random noise. *Bell Syst. Tech. J.*, **23**, 3 (July 1944) and **24**, 1 (January 1945)

TABLE 13.2. False Alarm Rate Versus Threshold to RMS Ratio for Pure Lowpass Gaussian Noise

Normalized Excursion $(V - \langle V \rangle) / V_{\text{rms}}$	Normalized Probability Density $V_{\text{rms}} \cdot dP/dV$	Probability of Being Above Threshold $P(>V)$	Upward Threshold Crossing Rate $\lambda_{\text{TC}} (B = 1 \text{ Hz})$
0.0	0.399	0.50	0.577
0.5	0.352	0.309	0.510
1.0	0.242	0.159	0.350
1.5	0.130	0.067	0.187
2.0	0.0540	0.0228	0.0781
2.5	0.0175	0.00621	0.0254
3.0	0.00443	0.00135	0.00641
4.0	1.34×10^{-4}	3.17×10^{-5}	1.94×10^{-4}
5.0	1.49×10^{-6}	2.87×10^{-7}	2.16×10^{-6}
6.0	6.08×10^{-9}	9.87×10^{-10}	8.79×10^{-9}
7.0	9.13×10^{-12}	1.28×10^{-12}	1.32×10^{-11}




Figure 13.12. Shot noise statistics. False alarm rate of the ISICL sensor of Example 1.12, as a function of the bipolar threshold level, compared with theory.

13.6.14 Shot Noise Statistics

At high arrival rates, shot noise really is Gaussian. Figure 13.12 shows some false count rate data taken with the ISICL sensor of 1 in a 1.1 MHz bandwidth, compared with threshold crossing theory (13.49); note the close agreement over 10 decades of vertical scale. When translated into an equivalent amplitude shift by plugging the observed false

alarm rate into (13.48) and solving for x , the imputed error is within 0.1 dB over the full range.

13.6.15 Thresholding

Thresholds cannot be set intelligently without knowledge of the noise statistics. We've already covered the Gaussian case, but there are lots of times when it doesn't apply, for example, if the dominant noise source is a fluctuating background. Unless you have reason for confidence (i.e., lots of measurements, not just statistical theory) to support your claim that your noise statistics are Gaussian, proceed with great caution. Don't jack up your thresholds unnecessarily out of fear, though. You're much better off compensating for effects such as source power variations in your threshold setting circuit. The *RCA Electro-Optics Handbook* has a useful discussion of thresholding with matched filters in noise, and Alexander has more depth.

13.6.16 Photon Counting Detection

If there are not a great many photons per second, the counting statistics begin to deviate seriously from the Gaussian case. A well-designed photon counting system with an average dark count rate λ (standard nomenclature, unfortunately) and a quantum efficiency η has dark pulses that arrive in a Poisson process. Shot noise is also Poissonian, but since the average current consists of very many electrons in practical measurements, the shot noise is almost always Gaussian to extremely high accuracy. Photon counting systems, on the other hand, have noise that is dominated by large dark pulses resulting from single photoelectrons emitted by the photocathode. The noise variance in a time t from such events is $R = \lambda t$ and the probability of observing exactly M events in a time t is

$$P(M) = \frac{(\lambda t)^M e^{-\lambda t}}{M!}. \quad (13.50)$$

The probability of observing at least M events in a time t is

$$\begin{aligned} P(\geq M) &= \sum_{n=M}^{\infty} \frac{e^{-\lambda t} (\lambda t)^n}{n!} \\ &\approx \frac{(\lambda t)^M}{M! e^{\lambda t}} \left(\frac{1}{1 - \frac{\lambda t}{M} e^{1/2M}} \right) \left(\frac{1}{1 + \frac{0.4M}{(M - \lambda t)^2}} \right). \end{aligned} \quad (13.51)$$

The approximation in the second line of (13.51), which is valid for $\lambda t > 0.5$, has a relative error of less than 11% for all M for which $P(\geq M) < 0.1$, which covers all relevant cases for signal detection. For $\lambda t < 0.5$, direct evaluation of the sum becomes easy, since a two-term truncation of the sum is accurate within 4% for any value of M . Thus we can set thresholds with some confidence.

At low to moderate count rates, Poisson statistics have different asymptotic behavior than Gaussian ones; the exponential falloff of (13.50) for the $M = 0$ case is much slower than the Gaussian, so that to achieve a very high probability of detection, we need more

photons per pulse for the Poissonian case than for the Gaussian one. On the other hand, assuming that our noise count rate is sufficiently small, the photon counting system can often detect an event with only a couple of photoelectrons. This is especially interesting since the information content is identical (remember the Rule of One). For detection schemes other than straight thresholding (e.g., coherent detection, QPSK, and so on), see Alexander.

13.7 FREQUENCY CONVERSION

Most high frequency measurement systems employ frequency conversion, which is the transferring of the desired modulation from one carrier frequency to another, perhaps moving an entire block of frequencies at once. It is used because signal detection, filtering, and digitizing are easiest at low frequency. Frequency conversion is exactly the same as SSB modulation, but with a different aim: to shift a (possibly modulated) signal from f_1 to f_2 , by mixing with another signal at $f_3 = |f_1 \pm f_2|$. Often we want f_2 to be constant even though f_1 varies, so that subsequent stages can be better optimized. From the early days of radio, the frequencies involved have been known as the radio frequency (RF), intermediate frequency (IF), and local oscillator (LO), respectively.

Frequency conversion is also known as *heterodyning*. The term *superheterodyne* was early radio ad-speak for a receiver that had a tuned RF amplifier before its mixer stage, but it stuck. A superhet radio uses fixed-tuned IF filters, which are easily made with good skirt selectivity and passband accuracy. To cover a broad range of frequencies, the RF section of the radio is a frequency converter (a tuned RF amplifier followed by a mixer), which subtracts the tunable LO frequency from the RF frequency to get to the fixed IF. The choice of IF is nontrivial. (Figure 13.13 has some symbols for these devices.)

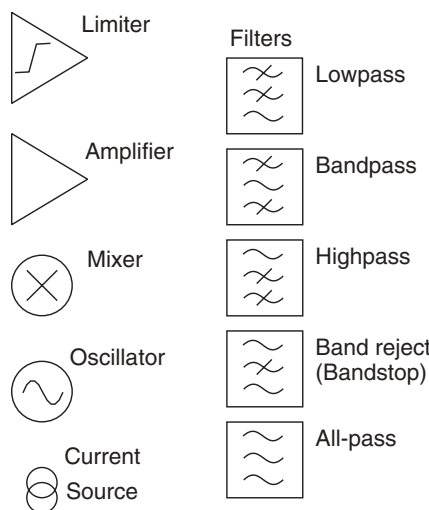


Figure 13.13. Some signal processing block diagram symbols.

13.7.1 Mixers

We encountered mixers in Sections 13.3.3 and 13.3.5 in the context of DSB modulation. Ideally a mixer is a device that algebraically multiplies (i.e., intermodulates) the signal voltage present on the RF port by the local oscillator signal at the LO port, to produce a signal at the IF port.[†] Mixers come in various types, which we'll see in more detail in Section 14.7.6. Any nonlinear circuit will produce mixing products—a mixer just does a better job, particularly a balanced mixer.

13.7.2 Choosing an IF

Frequency mixers never really produce just one output frequency. There are always strong components at $|f_{\text{RF}} \pm f_{\text{LO}}|$, one of which is the IF, and the other one is typically unwanted. A corollary is that a receiver with a given LO and IF can receive signals at more than one frequency. If the RF input contains many frequencies, then we must be concerned with another frequency, $f_{\text{image}} = |f_{\text{RF}} \mp 2f_{\text{IF}}|$, called the *image* frequency, which will also produce a first-order signal at the IF. If the IF is too low, the front end filter will not roll off enough to reject it adequately.

The subtlety of choosing an IF lies in getting good enough image rejection, and spreading out the spurious signals enough that they don't hit your IF passband (which favor a high IF) while maintaining adequate IF selectivity, gain, and phase stability (which favor a low IF). Furthermore, unless you intend to build all your own filters or cost is no object, the commercial availability of good, cheap filters at frequencies such as 455 kHz and 10.7 MHz is an important consideration. It is usually best to start out with a filter catalog in your hand.

There are also the weaker spurious signals in the mixer's output to consider: LO feedthrough, RF feedthrough, and spurious mixing products[‡] at $f_{\text{spur}} = |mf_{\text{RF}} + nf_{\text{LO}}|$, for all integers m, n . Their amplitude falls off toward larger $|m|$ and $|n|$, generally much faster with m . The LO harmonics are worse because mixers work much better if we drive the LO port hard: conversion loss decreases and overload characteristics improve.

Ideally, all these unwanted products will miss your IF passband, over the whole tuning range of the instrument. In practice, unless your measurement is narrowband, this may be difficult to accomplish. In radio applications, the RF input will probably have unwanted signals many times stronger than the one you want, but in electro-optical instruments the signal sources are much better controlled. In any case, make very sure that the strongest spurs (especially the LO and RF frequencies) don't cross your IF passband.

13.7.3 Image Rejection

If your RF signal varies from f_{low} to f_{high} , then in order that the image frequency never be in the passband, the IF must be greater than $(f_{\text{high}} - f_{\text{low}})/2$, plus some extra to allow the RF and IF filters to reject the image signal; this is why FM entertainment radios ($f_{\text{RF}} = 88\text{--}108$ MHz) have a 10.7 MHz IF. This assumes that $f_{\text{LO}} > f_{\text{IF}}$; for the more

[†]This is not the same as an audio recording mixer, which sums instead of multiplying.

[‡]These spurious products are closely analogous to unwanted diffraction orders in grating spectrometers and have to be avoided in analogous ways.

general case, the image rejection condition is

$$\min(f_{LO}, f_{IF}) > \frac{f_{\text{high}} - f_{\text{low}} + B_{IF}}{2}. \quad (13.52)$$

It was once customary to use double conversion, with a high first IF for image and spur rejection, and a lower second IF for high gain and good selectivity. Double conversion really works to reject the image but is often not as successful at rejecting spurs, because the two conversions generate a veritable forest of extra mixing products of their own.

Nowadays, high gain, high frequency amplifiers are not nearly so expensive as they once were, and good high selectivity filters are readily available at high frequency. Thus single conversion with a fairly high IF is usually best. The IF can be several times higher than the RF frequency in some cases.

Aside: Image Frequency Noise. It might seem that image rejection is not important in an instrument that doesn't have to cope with the vagaries of on-the-air reception, but that is not the case. Even if it doesn't contain interfering signals, the image contains shot and Johnson noise, which is uncorrelated with the desired signal's noise. All the noise contributions add in power at the IF, so failing to filter out the image noise before the mixer (or use an image reject [SSB] mixer) will reduce your SNR by 3 dB for no good reason. (This is the same effect as the noise difference between AC and DC measurements, which can also be avoided by SSB mixing—see Section 10.3.8.) If your devices are rolling off with frequency, the image noise may be even stronger than the signal's noise, making the SNR loss correspondingly worse.

Example 13.6: The AM Radio. Reducing noise and interference is not a new requirement. Your AM radio does an excellent job of rejecting out-of-band signals while pulling in the traffic report giving you the bad news about your commute. Similar techniques are very useful in electro-optical instruments.

Figure 13.14 shows how it works. The first stage is a tunable RF amplifier (i.e., it contains a tunable filter). This relatively crude filter rejects the worst out-of-band signals, reducing the opportunity for intermodulation, and also helps reject the image frequency. The gain helps overcome the loss and noise of later stages. The IF is 455 kHz, less than half the width of the AM band (540–1630 kHz). It is thus too low for the IF filter to adequately reject the image, which is one reason stations cluster near midband.

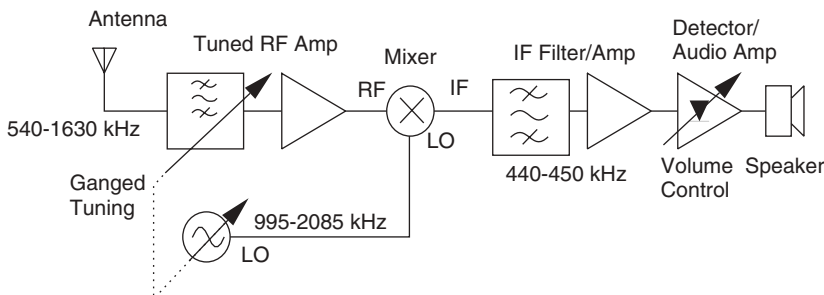


Figure 13.14. Block diagram of a superheterodyne radio, an archetypal noise-and-interference rejecting signal processing system.

A tunable LC local oscillator at 995–2085 kHz tracks the RF stage tuning, either by ganging the tuning capacitors together, or by adjusting the tuning voltage ranges if the stages are electronically tuned.

The mixer stage is usually a single bipolar transistor running in a strongly nonlinear regime. This is a poor mixer, but it's adequate to the purpose and is very cheap. AM radio ICs have better mixers, usually Gilbert cell multipliers.

A fixed IF filter, usually two solid ceramic resonator devices in cascade, provides about a 10 kHz bandwidth, with good skirt selectivity. A high gain IF amplifier with slow automatic gain control (AGC) presents a nearly constant average power level to the detector, which produces low level audio output. The AGC stage ensures that there is enough power delivered by the IF to the detector, without distorting the envelope by overdriving the amplifiers. The RF amplifier is also occasionally AGC controlled. Discrete detectors are usually half-wave rectifiers followed by RC filters, but integrated ones are almost always multipliers, which multiply the signal by a clipped replica of itself. An adjustable gain audio amplifier and speaker complete the signal path.

The ability to use good fixed-tuned filters means that the filter characteristics do not vary with tuning, as they inevitably do with tunable stages. Even more than that, though, the fixed IF frequency relaxes the requirements on the signal and AGC detectors, and allows the IF strip and back end to be tested easily as a separate subsystem, which is very important in system development and production.

13.7.4 High Side Versus Low Side LO

A frequency plan is more than just choosing an IF; there's also the LO to consider. Unless your IF is at DC, there are always two possible choices of LO: $f_{LO} = |f_{IF} \pm f_{RF}|$. In a wideband measurement involving down conversion, the choice is usually obvious. For example, in an AM radio, where $f_{RF} = 540$ to 1630 kHz and the $f_{IF} = 455$ kHz, a high side LO would be 995 to 2085 kHz, whereas a low side LO would be 85 to 1175 kHz.

A low side LO would be disastrous for two reasons. First, spurious responses: the LO and its second through fifth harmonics would cross the IF, leading to huge spurs from LO feedthrough. Near the low band edge, the LO harmonics would lead to a huge number of spurious responses to other stations throughout the band. Second, circuit requirements: it is much more difficult to make a good stable oscillator tunable over a 14:1 range than a 2:1 range. We couldn't use an LC oscillator, because the capacitance would have to change nearly 200:1. Fortunately, in an AM system, the symmetric sidebands allow LSB conversion with no frequency-inversion problems (see Section 13.3.6).

13.7.5 Direct Conversion

Synchronous receivers such as lock-in amplifiers generally work with the LO at the center of the RF passband—that is, the IF is at DC. Mathematically, this is equivalent to multiplying $A \cos(\omega t + \phi)$ by $\cos \omega t$ —the envelope is preserved, but the quadrature component of the modulation is lost due to sideband folding. To preserve it, we can use SSB mixing (see Sections 10.3.8 and 13.8.7), use I/Q demodulation, which is what a two-phase lock-in does[†], or phase-lock to the carrier (Section 13.9.3).

[†]See also Section 13.9.6.

13.7.6 Effects of LO Noise

It is often said that mixers are insensitive to LO noise, but this is somewhat misleading. A diode bridge mixer, for example, is usually operated in a switching mode, with a LO signal large enough to turn the diodes completely on and off each cycle. A small variation in LO power will not significantly change the diodes' resistances, so it causes a much smaller than linear change in the IF output amplitude. LO phase noise is of course passed unattenuated; the instantaneous phase of the IF output is the sum or difference of the instantaneous phases of the RF and LO signals. This will degrade FM/PM and SSB measurements. Phase noise in the LO does not greatly affect AM measurements, because the envelope amplitude is independent of the carrier frequency. The only time you get into trouble with phase noise in AM is if it's so large that it extends to the skirts of the IF filter. In that case the filter will turn the PM into AM, so you have to increase your IF bandwidth to accommodate it. It's the same in optics—a LED has horrible phase noise but excellent amplitude stability.

13.7.7 Gain Distribution

For stability and low spurious generation, it is important to distribute the gain between the RF, IF, and baseband stages wisely. The mixer usually has the lowest dynamic range of any stage, so put just enough RF gain in to override its noise. Putting too much gain in the IF is liable to lead to oscillation, especially if the layout is too tight. Because any pickup from the baseband stages is at a different frequency, it is rejected by the interstage filters and does not cause oscillation. (It's also at lower frequency, so it's much less likely to couple where you don't want it.) There are power issues as well: there's no point in pumping a 1 kW signal through a 100 dB loss—unless you're building a radar, you can get better SNR at much lower cost and lower power by distributing the gain and loss a bit more evenly.

13.8 FILTERING

Just about any transfer function describes a *filter*, loosely defined as something that treats different frequencies differently. Filters are very important, because they allow us to choose the frequency components we want, and to modify their amplitude and phase as required by the signal processing scheme. The details of the different kinds of filters, their characteristics, and some design information are covered in Section 13.8.9. Here we stick to the systems design and theoretical issues.

Most of the time, what we want is a filter that passes what we care about and rejects everything else. Typical filters of this type are lowpass, which reject everything above the cutoff frequency f_c , and highpass, which reject everything below f_c . Bandpass and bandstop (band reject) filters similarly pass or reject everything between f_1 and f_2 . We'd like to reject the bad stuff (the stopband) completely, while keeping all of the good stuff (the *passband*), and have a vertical cliff separating the two regions, but are willing to allow a *transition band*, as shown in Figure 13.15. The ratio of output power to input power is called the *insertion gain* (its reciprocal is the more commonly heard *insertion loss*), and the phase shift between input and output is the *insertion phase*. The flatness of the passband is the ratio of the peak-to-average insertion gain ratio, or alternatively the square root of the peak-valley ratio, and is normally specified in decibels.




Figure 13.15. Filter specification, showing passband (left), stopband (right), and skirt selectivity requirements, together with a Chebyshev filter design that just meets the specification.

In signal processing, we often want at least one edge of the filter to fall off very steeply. The trailing-off edges of the filter's passband are called the *skirts*, and the skirt selectivity is expressed by the *shape factor*, the ratio of the filter's bandwidths measured at 60 dB down and 3 dB down (f_3/f_2 in Figure 13.15). A shape factor of less than 3 is pretty sharp, but fancy filters can sometimes achieve 1.1. (The time-domain response of such a filter is not pretty to look at.) A one-pole RC has a shape factor of 1000.

13.8.1 Cascading Filters

Most signal processing systems involve at least a few stages of filtering. Remember that the transfer functions of cascaded filters multiply each other. In the process, the passband gets narrower, and any ripple usually gets worse. Consider two identical bandpass filters, whose 3 dB bandwidth is 10 to 13 kHz. When the two are cascaded, these frequencies are down by 6 dB instead of 3. If the top of the filter is roughly Lorentzian in shape, like a single *LC* section, the 3 dB bandwidth will have decreased by a factor of $1/\sqrt{2}$. Filters with more nearly rectangular transfer functions will narrow less; pointy ones will narrow more.[†]

If the filters are passive *RLC* units, whose transfer functions depend on the source and load impedance, then cascading two or more of them will change their shapes, because a passive filter usually looks like a nearly pure reactance in its stopband, which effectively changes the *LC* values of the end sections of the filter. Putting a pad (attenuator) or a well-matched amplifier in between will help.

Deep in the stopbands, the rejection may not be as good as the product of the two filters. Eventually you're going to be dominated by other signal paths: capacitive and

[†]Examples are skin-depth attenuation in coaxial cables, which goes as $f^{-1/2}$, so that doubling the length of a cable will reduce its bandwidth by *four times*.

inductive pickup, power supply coupling, slightly wiggly grounds, that sort of thing. Don't be surprised if cascading two filters with 70 dB rejection results in 92 dB rejection rather than 140 dB, at least at first. A useful rule of thumb is that in a filter you can get 60 dB of attenuation per inch of separation on the board, but that anything more is hard, requiring shields and extra care.

13.8.2 Impulse Response

We specify filters most often in the frequency domain. Since the transfer function multiplies the input spectrum, the frequency response H is the transform of the output signal produced by an input whose transform G is 1 for all f . Equivalently, h is the output signal in the time domain produced by an input g whose transform is 1—that is, an impulse function, $\delta(t)$. For this reason, $h(t)$ is called the *impulse response* of the filter.[†]

13.8.3 Step Response

In many applications, the response of the system to a step change in the input is the most important parameter. After a sudden change, the filters and amplifiers must slew to the new value and settle there to high accuracy before a measurement can be made with confidence. The figures of merit are the rise and fall times, usually quoted between 10% and 90% of the total change, and the settling time, which must elapse following the input step before a good measurement can be made. Rise and fall times are set almost entirely by the filter bandwidth, and the settling time by a combination of the bandwidth and group delay characteristics. (This assumes that the network is highly linear; nonlinear settling is considered in Section 14.7.15.) Lowpass filters settle roughly twice as fast as bandpass filters with the same bandwidth. (As usual, this assumes that the low frequency cutoff of the bandpass filter isn't too close to 0—the bandpass is even slower if it is.) Normally we don't worry a great deal about the settling performance of a highpass filter, because we want the highpass response to be constant over our measurement time (i.e., slower is usually better).

13.8.4 Causality

The causality condition says that effects do not precede their causes, so $h(t) \equiv 0$ for $t < 0$. Enforcing this condition in the Fourier domain is exactly analogous to what we did in the time domain to form the analytic signal—chopping off all negative frequencies—except with the opposite sign of j in the imaginary part due to flipping the t axis. (See Section 13.2.5.) Thus the real and imaginary parts of H must be a Hilbert transform pair:

$$\Im\{H(f)\} = \mathcal{H}\Re\{H(f)\}. \quad (13.53)$$

This constraint is important in the design of filters, since it enables us to take a frequency-domain specification for $|H|^2$ and turn it into a possible filter design. A less palatable consequence is that it specifies a minimum *delay* for a given filter response.

[†]For the mathematically inclined, it is the Green's function for the differential equation expressed by the transfer function of the filter, with the initial condition on the derivatives that $h^{(n)}(0^-) = 0$ for all n .

It is reasonable that a filter which can distinguish sharply between frequencies f and $f + \Delta f$ should require a time delay on the order of $1/\Delta f$ seconds to do it. More subtly, a filter with a cutoff frequency of 1 kHz should be able to distinguish between DC and 2 kHz faster than it can between 900 and 1100 Hz. Thus the time delay through a filter conceptually need not be exactly the same for all frequencies, and in fact generally is not.

Aside: Stable and Unstable Transfer Functions. One of the most basic constraints on the impulse response of a filter is that it should not grow without bound as $t \rightarrow \infty$ —otherwise we’ve inadvertently built an oscillator. This happens all the time with active circuits, but would be strange to say the least in a pure *RLC* filter. The Fourier transform of a rational function can be found by residue calculus (don’t worry, no branch cuts here—just isolated poles). It’s worth doing this, because it shows the deep connection between stability and causality.

Consider a rational transfer function $H(f)$, which has been decomposed into partial fractions:

$$H(f) = \sum_{k=1}^N \frac{r_k}{f - f_k}, \quad (13.54)$$

where r_k is a complex constant. (This form assumes that the order of the numerator is not greater than that of the denominator, which must always be true for passive networks. Why?) The inverse Fourier transform of this is

$$h(t) = \sum_{k=1}^N \int_{-\infty}^{\infty} \frac{r_k}{f - f_k} e^{j2\pi f t} df, \quad (13.55)$$

which we propose to evaluate by residues, using a contour consisting of the real axis plus a semicircle of large radius R , which we will let go to infinity. (Note that if $h(t)$ does not die away at large times, the integral along the real axis will not exist, so we can’t use this particular contour with unstable transfer functions.) As always, poles inside the contour of integration contribute to the integral, whereas those outside do not. Poles above the real f axis will contribute residues with exponentially decaying time dependence, so their contributions are stable as $t \rightarrow +\infty$. Poles below the axis will contribute an exponentially growing time dependence, so stable transfer functions must not have poles in the lower half plane.

When we go to close the contour, using Jordan’s lemma to let us neglect the contribution from the large semicircle, we find that the upper half-circle contributes 0 for positive times, but that the integrand is unbounded there for negative times. We therefore close the contour above for positive times, and below for negative times.

Requiring causal behavior means that there can be no poles inside the contour at negative times, so causality forbids any poles whatsoever to be in the lower half-plane. Thus our requirements on stability and causality lead to the same condition on the poles of H , a remarkable result first proved by Titchmarsh.[†]

[†]E. C. Titchmarsh, *Introduction to the Theory of Fourier Integrals*, 2nd ed. Oxford University Press, Oxford, UK, 1948.

13.8.5 Filter Design

If you don't care much about the time-domain response of your filter, or if it is much wider than the bandwidth-setting filter (usually the second IF filter), you can just pick the power spectrum $HH^*(f)$, factor the denominator into complex-conjugate poles, discard the ones in the lower half-plane, and you have a causal filter function. Then you have to synthesize a network with that transfer function, which is the hard part. We'll talk more about the details in Section 15.8, but first we need to understand the problem.

13.8.6 Group Delay

In Section 1.2.2, we encountered dispersion, where different colors of light traveled at different speeds, and group velocity, where the modulation traveled at a different speed than the phase (always more slowly). All causal, frequency-selective systems exhibit the same sort of behavior. In optics, the causality condition is called the Kramers–Kronig relation, but it's just the same idea.

If the phase shift of a sinusoidal signal passing through the filter is $\phi(f)$, then the phase delay is just $\phi/(-2\pi f)$, but the group delay, which is what the modulation sees, is $(-1/2\pi)\partial\phi/\partial f$.

A pure time delay of τ phase shifts each component by $-2\pi f\tau$ radians; its graph is a straight line with no offset. This delays the modulation but does not distort its shape. Networks whose delay has this characteristic, at least approximately, are called *linear phase networks*. Coaxial cable and Bessel filters are two examples. Variations in the group delay across the passband distort the relative phases of the Fourier components of the modulation envelope. Filters whose group delay varies all over the place in the passband have weird time-domain behavior, typically surfacing as severe overshoot and ringing on their step responses. Group delay in the stopband is not of much consequence, since signals out there don't appear in the output, but the transition band matters nearly as much as the passband.

As we might expect, there is a trade-off between sharp cutoff and good group delay performance. For a given filter order (number of poles + number of zeros), sharper filters have larger group delay in general, with delay being greater in the passband than in the stopband. Sharp filters also exhibit tall, narrow group delay peaks near the band edges, which is reasonable, as we mentioned earlier. Elliptic filters are the worst, followed by Chebyshev, Butterworth, Bessel, Gaussian, and the champion, the equiripple group delay filter (see Zverev).

Linear phase filters whose group delay is as small as causality permits (the so-called minimum phase condition) all have a rather gentle slope in their transition band and significant bowing in the passband, which makes them far from ideal for rejecting spurious signals and especially for cascading. This property is not improved by adding more poles to the filter, because when you do this the filter approaches a Gaussian response, which is no flatter.

All is not lost, however; it is possible to synthesize networks that have varying group delay but whose transfer function has unit magnitude. These are generically called *all-pass filters*, in this case specifically *group delay equalizers*. Putting in an equalizer will at least double the number of elements in the filter but will yield sharp filters with flat group delay (albeit several times longer than the peak delay of the original

minimum phase filter).[†] The required order of the equalizer increases more rapidly than the filter order, and especially with Cauer filters, it becomes prohibitive quickly as the filter becomes sharper. Have a look at Zverev or *Reference Data for Engineers* for lots more design information. You can buy filters like this, so you don't need to become a filter design whiz to make your measurement work. If group delay is really killing you, consider using a FIR digital filter, which can have exactly constant delay (see Section 17.6).

Aside: Ladder Networks. Most filters built from discrete *RLC* sections are *ladder filters*—they are connected with series and shunt sections, and their schematic looks like a ladder. All passive ladder networks are automatically minimum phase,[‡] so group delay equalizers and other non-minimum-phase filters must have more complicated connections (e.g., the lattice filter of Section 15.20).

13.8.7 Hilbert Transform Filters

It is frequently very useful to construct the envelope of a function, e.g. a tone burst of several cycles; this makes an especially nice AM detector and is good for detecting the amplitude of pulses. Other times, we want to build image reject mixers (single sideband mixers), where we use the phase relationships between upper and lower sidebands to reject one and pass the other. SSB mixers are useful in rejecting image frequency noise, which is sometimes good for a 3 dB improvement in measurement sensitivity. Supporting both of these worthy aims is the *Hilbert transform filter*. An ideal Hilbert transformer, remember, would apply a 90° phase shift to all positive frequencies and reject all negative ones. Thus we want to apply a phase shift of exactly 90° to all signals we care about, without changing their amplitudes. Unfortunately, it's impossible to do this. What we actually do is make two all-pass filters, whose phase shift differs by 90°, within some accuracy, over some given bandwidth. This is much easier, especially for relatively narrow bandwidths (an octave or less). The two outputs, notionally 0° and 90°, are known as *I* and *Q*, respectively, for in-phase and quadrature. The basic scheme is called *vector modulation*.

Provided the modulation sidebands don't go down too close to DC, the analytic signal is $I + jQ$, and its magnitude is

$$|\Phi|^2 = I^2 + Q^2. \quad (13.56)$$

This can be constructed by several means. Section 15.5.4 has a cute trick used in radar to make a good approximation of $|\Phi|$, very fast and without massively increasing the required dynamic range by squaring the signal.

Example 13.7: Building an SSB Mixer. Single sideband mixers can be made from two power splitters, two mixers, and a power combiner (which is a splitter run in reverse). The LO and RF splitter/combiners are quadrature devices,[§] which contain the *I-Q* filter

[†]Since we can't go backwards in time, equalizers have to work by adding group delay to the quick parts, not taking it away from the slow ones. Thus the total delay through the equalized filter will be greater than the maximum delay of the original one.

[‡]H. W. Bode, *Network Analysis and Feedback Amplifier Design*. Van Nostrand, New York, 1945, p. 243.

[§]The 90° power splitters are also called quadrature hybrids.

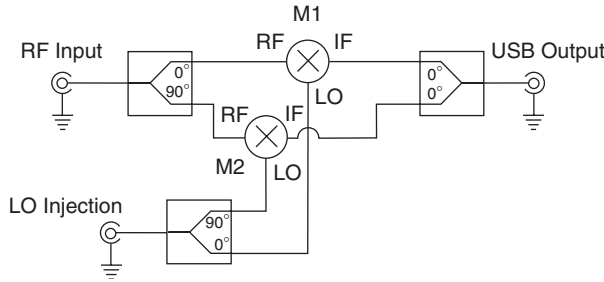


Figure 13.16. Single sideband (SSB) mixer: due to the opposite phase dependence of upper and lower sidebands, making two DSB signals with out of phase sidebands allows the unwanted sideband to be canceled, producing an SSB output.

pair inside them. The block diagram is shown in Figure 13.16. Assuming the RF input is $\sqrt{2} \sin \omega t$, the LO injection is $\sqrt{2} \sin \Omega t$, and ignoring for the moment the common-mode phase shift in the quadrature splitters, the outputs of M1 and M2 and the resulting summed IF voltage are

$$\begin{aligned}
 V_{M1} &= \cos \omega t \cos \Omega t \\
 &= \frac{1}{2} \cos(\omega + \Omega)t + \frac{1}{2} \cos(\omega - \Omega)t \\
 V_{M2} &= \sin \omega t \sin \Omega t \\
 &= \frac{1}{2} \cos(\omega + \Omega)t - \frac{1}{2} \cos(\omega - \Omega)t
 \end{aligned}
 \tag{13.57}$$

$$V_{IF} = \cos(\omega + \Omega)t$$

which is a pure USB output. Switching the outputs of one of the quadrature splitters will produce an LSB output instead. Equivalently, you can subtract the two IF signals instead of adding them, and even have both an adder and a subtracter, producing both sidebands separately. The unwanted sideband rejection obviously depends on excellent phase and amplitude matching between the two signal paths. The degree of matching required can be determined by the kind of simple analysis we used in Section 13.6.9: 40 dB sideband rejection requires an amplitude error of $<1\%$, a phase error of <0.01 radian, or some combination of smaller errors in each. Such close matching of signal processing components will require one or two trims in general.[†]

The 0° and 90° phase LO signals can also be made by dividing a signal at $4f_{LO}$ by 4 in a synchronous $\div 4$ walking ring counter (Johnson counter), or at UHF by putting a quarter-wave piece of transmission line in one arm of a 0° splitter.

13.8.8 Linear Phase Bandpass and Highpass Filters

As we'll see in Chapter 15, the usual procedure for designing a bandpass or bandstop filter with a 3 dB bandwidth from f_1 to f_2 involves taking a prototype (i.e., canned)

[†]The difficulty of this is one reason for the current trend toward DSP radios, in which the multiplication is done numerically after digitizing. That takes a lot of engineering and isn't always possible.

lowpass filter design whose bandwidth is about $\sqrt{f_1 f_2} - f_1$, and resonating all its elements at $f_0 = \sqrt{f_1 f_2}$.[†] The resulting frequency transformation maps the DC response of the lowpass to f_0 and makes the skirts symmetrical when plotted against $\log(f)$. Similarly for a highpass, you swap inductors and capacitors, keeping $|X(f_c)|$ the same, which maps f into f_c^2/f . Unfortunately, these nonlinear frequency stretches mess up the linear phase characteristic completely, so you can't build linear phase highpass and bandpass filters this way. Packages exist for designing filters of this sort, but we often wind up writing programs to do it.

Another subtlety is that in order for a bandpass or highpass filter to have good pulse fidelity, the trend line of the linear phase has to pass through the origin, or at least be offset by $2n\pi$ radians, and this need not be the case. Roughly speaking, if it doesn't, the carrier will be phase shifted with respect to the envelope, just as in group velocity dispersion in optics. This isn't too serious for narrowband filters, but for wider ones (say, >15% fractional bandwidth) you should check for it, because the position of the carrier peak within the envelope matters. Remember that if the trend line goes through (0 Hz, $\pi/2$ rad), you've built a bandpass Hilbert transformer, which may not be what you want.

13.8.9 How to Choose a Filter

Choosing a filter for your application starts with knowing what you need it for. A system that needs good pulse response while minimizing noise bandwidth will require a filter with good group delay characteristics. If the filter is inside a feedback loop (e.g., a fast AGC loop), the total delay is important as well as its flatness. In cases like that, a Bessel, Gaussian, or equiripple group delay filter is indicated. Most of the time, though, a little extra delay is not an excessive price for better skirt selectivity and narrower noise bandwidth for a given passband flatness, so that a group delay equalized Chebyshev or Causer filter will do a better job.

The group delay characteristic is normally of interest only in the narrowest filter in the signal path. Elsewhere, we are concerned with rejecting images and spurious signals, and minimizing the noise bandwidth, while using the cheapest filter that will do the job. Cheapness has different meanings depending on what the intended use of our system is, and on how many are to be made. For one-offs or experimental systems, the optimal setup is a drawer full of Mini Circuits filters for chopping off harmonics, plus a few hand-wound coils and variable capacitors for the slightly more difficult cases, and perhaps a custom filter with a very flat passband and good skirts for the bandwidth setting filter.

Bandpass or bandstop filters with very wide bands (an octave or more) are difficult to build, so consider cascading a lowpass and a highpass instead. Narrow bandpass and bandstop filters have much steeper skirts than highpass or lowpass filters with the same complexity and cutoff frequency—the characteristic scale of the decay is δf rather than f_0 . This is because a bandpass filter at $f_0 \pm \delta f$ is really a lowpass filter of bandwidth $2\delta f$ (see Section 15.8.6) whose elements have each been resonated at f_0 by an appropriate series or parallel element paired with each one; thus the skirts tend to roll off as inverse powers of $(f - f_0)/(\delta f)$ until the finite Q of the inductors limits them.

[†]Since the passband is mirrored about f_0 , why isn't the equivalent lowpass bandwidth $(f_2 - f_1)/2$?

If you need to make your own *LC* filters, have a look at Chapters 15 and 19 for how to design and adjust them. There's a whole lot about practical filter prototyping in the ARRL handbook.

13.8.10 Matched Filtering and Pulses

Consider a narrowband signal in constant-amplitude additive white noise (e.g., Johnson noise but not shot noise), running into a linear phase filter of adjustable bandwidth. For discussion's sake, the signal is a train of narrow tone bursts 10 carrier periods long, spaced by 1000 carrier periods. A very narrow bandwidth will essentially sample the signal PSD at the carrier frequency, and the SNR will be just (signal PSD)/(noise PSD), evaluated at the carrier frequency. Seemingly, this is the best case, because the signal PSD drops off from there, and the noise PSD is constant. It would seem that the only reason to open up the filter would be to trade off SNR versus measurement speed, or possibly to overcome additive noise in later stages.

This analysis is exactly correct for a linear, time-invariant system with no additive noise after the filter, and so is relevant to unmodulated carriers. With a pulsed waveform, though, there is no information between the pulses, so we can turn off the receiver then and eliminate most of the noise. We'll put in this time-varying behavior by hand as usual, by paying attention to the instantaneous SNR within the pulse, instead of the average SNR.

Figure 13.17a shows a series of pulses in additive noise, with a SNR near 0 dB. Applying a narrow filter as in Figure 13.17b gets rid of most of the noise but also smears the pulses out completely, leaving only the DC, and that at a very low level because of the small duty cycle. As the filter gets wider as in Figure 13.17c, the total SNR starts out constant and gradually degrades, because at some point each additional hertz lets less and less signal get in, but the same amount of noise. The key observation, though, is that not only does the average power increase linearly with bandwidth (because of the nearly constant PSD close to the peak), but the pulses get shorter by a factor of $1/B$ too. Thus the instantaneous signal power in the pulses goes as B^2 for small bandwidths, and the instantaneous SNR within the pulse goes as B .

We keep winning by increasing B until the filter no longer limits the pulse width, and then even the peak-power/noise ratio starts to decline again, as we'd expect. Thus from the instantaneous SNR point of view, the best filter for this application is a linear phase bandpass, centered on the carrier, adjusted so as to stretch the pulses by a bit but not too much—say $\sqrt{2}$ times. The overall SNR improvement is maintained by thresholding (i.e., ignoring signals weaker than a threshold), by turning off the receiver between pulses, or ideally (as in Figure 13.17d), multiplying by the input pulse shape.[†]

Adding group delay dispersion to the filter will change it. This can be a good thing in special cases, for example, grating pairs used as pulse compressors in femtosecond lasers, but usually it hurts you by smearing out the pulses and adding artifacts without improving the noise.

As we can see, the choice of the narrowest filter in the signal path is the thorniest and really requires you to know what you're trying to optimize. If you like calculus of variations, you can easily prove that the maximum instantaneous SNR for a statistically stationary, deterministic signal in the presence of additive white noise is achieved with a

[†]These are both nonlinear operations, which is how we can get better total SNR than a linear analysis suggests.

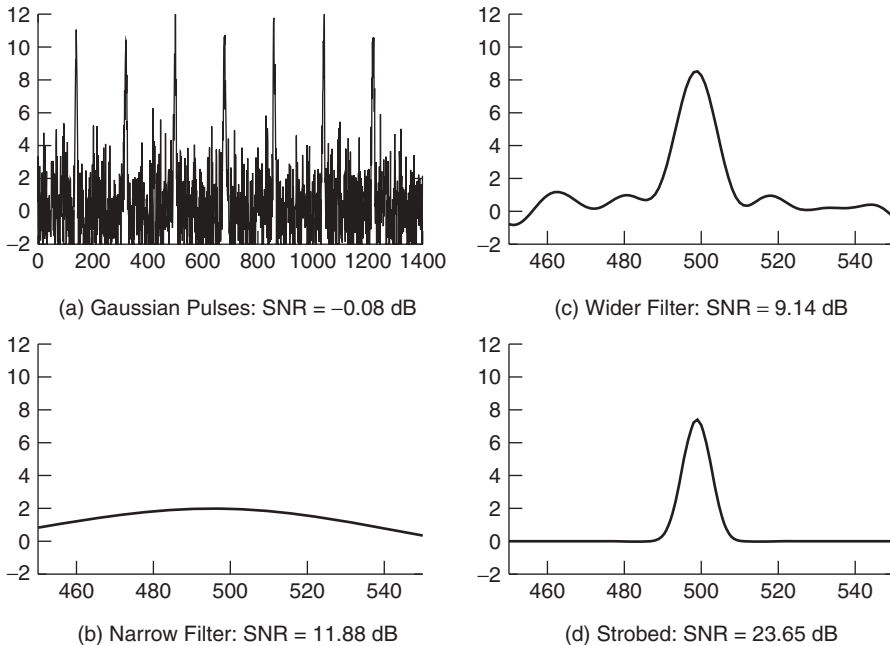


Figure 13.17. Pulse detection in the time domain. This is a fairly typical case, with a duty cycle on the big side (3%). The improvement is more dramatic with narrower pulses. Strobing is done by multiplying by the input pulse shape (Gaussian).

matched filter, that is, one whose transfer function is the complex conjugate of the signal spectrum (the complex conjugation undoes any group delay smear, see Example 1.12). If the noise isn't white, the optimal strategy is to whiten it (i.e., apply a filter that makes the noise PSD constant), apply the matched filter, and then unwhiten again to recover the narrow pulse shape.

If you are a bit unsure about what the best possible filter is for your application (and who isn't, at least to start with), choose a filter whose passband is fairly flat, and whose 3 dB bandwidth coincides with the 6 dB bandwidth of your signal waveform. Get one that has good group delay performance, no more than $\pm 2^\circ$ deviation from linear phase over the passband and transition band. This filter will be reasonably close almost always, and it errs in the right direction, namely, by being slightly too wide. Once you have your system working with this filter, you can fiddle with optimizing it to wring out the last decibel or so here, and find a cheaper filter in the process. The optimum is pretty flat, so it is unnecessary to bleed over errors of 10–15%. Note that we have implicitly assumed here that the desired signal is slowly varying, so that we don't need to consider sampled-data effects. Depending on the situation, these become important for signal bandwidths above 1–15% of the rep rate f_{rr} ; see Chapter 17 for how to handle this case.

Aside: Matched Filters and Signal Averaging. In case you're worrying about signal averaging in this situation, a repetitive signal can indeed be improved by signal averaging after detection in a wide bandwidth, but in that case the signal PSD is zero except near

harmonics of $1/P$, which are the only frequencies that survive averaging, so the signal averager is functioning as a matched filter.

Example 13.8: SNR Improvement with Pulsed Measurements. You don't do a pulsed measurement by shuttering a CW beam, but by crowding the available photons together in a high peak power pulse. If the average optical power is being kept constant as the duty cycle is reduced, the peak optical power and the peak photocurrent both go as d^{-1} , and peak electrical power as d^{-2} . If the measurement bandwidth needs to be at least f_{\max} , the pulse repetition rate f_p must be greater than $2f_{\max}$, so that the signal and noise are multiplied by the gating function $s(t)$,

$$s(t) = \text{rect}\left[\frac{f_p t}{d}\right] * \text{III}(f_p t), 0 < d \leq 1, \quad (13.58)$$

where III is the sampling function (Cyrillic capital *sha*), a train of unit strength delta functions spaced at unit intervals. Note that this function becomes identically 1 as d goes to 1, so that we can connect the results for small d to the continuous case easily.[†]

The increased electrical power will be spread out over a wider frequency band, however. The signal has been multiplied by a pulse train, which smears its spectrum out through convolution with the transform of the pulse train,

$$S(f) = \frac{d}{f_p^2} \text{sinc}\left[\frac{fd}{f_p}\right] \text{III}\left[\frac{f}{f_p}\right], \quad (13.59)$$

so that the spectrum of the sampled signal is turned into a lot of little copies of itself, centered at all harmonics of f_p (including 0), whose amplitudes are governed by the sinc function envelope. To reconstruct our signal, we could just lowpass filter the pulse train, to reject all but the lobe at 0, which is identical with our original signal, no more and no less; all the extra electrical power has gone into the other lobes, and the noise bandwidth is just what it was before, so we will have gained exactly nothing by the exercise.

On the other hand, if we widen the filter bandwidth to about f_{\max}/d , we'll gain the d^{-2} advantage in signal power, at the price of a d^{-1} increase in the noise, even if we do nothing else. We can do even better, by time gating the noise as well, by using the same pulse waveform to gate the received signal. This of course does nothing to the signal but reduces the noise power by a factor[‡] of d , so that our electrical SNR improvement is of order d^{-2} . Note that all this depends on our being able to pack the same number of photons into a shorter time. If the peak optical power doesn't grow as we shrink the duty cycle, we gain nothing—we've made a perfectly ordinary sampled-data system.

13.8.11 Pulsed Measurements and Shot Noise

A pulsed measurement concentrates all the signal power in a small time slice, where it can dominate the Johnson noise in the difficult photocurrent range from 1 pA to 1 μ A.

[†]What happens to the spectrum as $d \rightarrow 1$?

[‡]Assuming the noise is white.

That's the only noise benefit, though, and you trade a lot away to get it.[†] Pulsed lasers are dramatically noisier than CW ones, have poorer spatial behavior, usually exhibit serious timing jitter, and are generally hard to use. Photodetector nonlinearity depends strongly on the peak photocurrent, so it gets far worse at short duty cycles. There is lots of spurious crud flying around at f_{rr} and its harmonics, from power supplies, Q switch drivers, and pulse gating electronics. Shrinking the duty cycle makes all these things worse rather than better.

Furthermore, shrinking the duty cycle doesn't help shot noise at all. There's no shot noise between pulses, of course, but if you're going to collect 10^{14} photoelectrons in one second, you're going to have 10^7 electrons rms noise in that time, whether the duty cycle is 1 or 10^{-6} . Thus as we improve the rest of the measurement, the need for pulses tends to go away; unless we're really stuck for photons, or there is some absolutely unavoidable noise source, for example, the shot noise of the self-luminosity of the sample,[‡] or you need the tall narrow pulses for another reason, for example, nonlinear optics or stroboscopy, you're a lot better off concentrating on getting to the shot noise than on fighting the annoyances of pulsed measurements.

13.9 SIGNAL DETECTION

13.9.1 Phase-Sensitive Detectors

From a signal processing theory point of view, a phase-sensitive detector is a mixer stage whose IF is at zero hertz. There are four common kinds of phase detectors that accept analog inputs: switching ones built from CMOS analog gates, track/hold amps, Gilbert cell multipliers, and diode bridge mixers. With a sinusoidal input, the amplitude and phase changes in the RF signal are converted to a baseband signal,

$$V_{IF} \approx A_{RF} \cos \theta. \quad (13.60)$$

Because phase detectors preserve phase information, filtering after the phase detector has nearly the same effect as doing it before; the only difference is the sideband folding, where USB and LSB signals produce the same output, so you can't separate them with a single mixer. This is in decided contrast to an ordinary rectifying detector, which will merrily intermodulate all noise components together, so that even distant noise will wind up landing on top of your desired signal. Even the sideband folding problem can be eliminated by using an SSB, making baseband filtering precisely equivalent to IF filtering.[§] This is a very useful fact, because you can make a 10 Hz wide filter at baseband with a 150 k Ω resistor and a 0.1 μ F capacitor, which is a lot easier than doing the same trick at 10.7 MHz.

[†]It isn't the same as an ordinary AC measurement, because the signal power gets spewed out over a huge number of harmonics of the rep rate. We don't even get the benefit of getting out of the $1/f$ noise, because only the DC term is immune to timing jitter and other nonideal behavior.

[‡]Even if the sample is strongly self-luminous, a correctly designed CW measurement will usually win; the ISICL sensor of Chapter 1 needs only 50 photons to detect a particle in a plasma that's too bright to look at.

[§]Of course, if your sideband folding happens at optical frequencies (i.e., at the optical detector), you're dead. This is the origin of the 3 dB noise advantage of homodyne (IF at DC) measurements over heterodyne ones in optics. It's only the low frequency noise that tends to make homodyne measurements worse.

13.9.2 AM Detectors

There are two basic kinds of AM detectors: rectifying and synchronous. The classical AM detector is a 1N21 UHF diode wired up as a half-wave rectifier with an RC filter after it (nowadays you'd use a Schottky or back diode)—what could be simpler?. Unfortunately, this leads to horrible nonlinearity due to the soft forward turn-on behavior of the diode. The linear operating range of a diode detector is thus poor, about 20 dB, although there are circuit hacks that can be applied, as we'll see in Chapter 15.

Synchronous AM detectors are just phase detectors operated at $\theta = 0$ or π . Because of their excellent linearity, these make the best AM detectors, but of course there's the problem of where to get the LO signal. If you know its frequency and phase a priori (because you're generating it elsewhere), everything is easy—use the transmit signal as the LO, or put a very narrow filter on the input signal and use that. If you're not that lucky, you can fake it by putting the raw input signal through a comparator or limiter to make the LO—a pseudosynchronous detector. You have to worry about phase shifts and AM–PM conversion. A poor man's version of this that works only at low frequency is the (badly misnamed) perfect rectifier, which we'll encounter in Chapter 14. A common malady of these pseudosynchronous detectors is that they cannot distinguish signal from noise; the additive noise makes the LO phase noisier and noisier as the signal level drops, so that the low-SNR behavior of pseudosynchronous detectors is similar to that of rectifying ones.

A more complicated approach that yields excellent performance is to use an I/Q mixer to construct the envelope of the analytic signal directly, using the trick from Chapter 15. AM detectors are tougher to get right than you'd think. Design yours carefully.

13.9.3 PLL Detectors

Another approach, which will separate the signal from the noise a great deal better, is to phase lock to the signal (see below), and use the PLL output as the LO. This approach is probably the best if you anticipate working at low SNR a lot and want to be able to narrow the noise bandwidth after detection, as in a lock-in amplifier. This is not something for nothing; the PLL can be narrowband, so the SNR will be much better in the PLL's bandwidth than it is in the IF bandwidth. A corollary is that you can't usefully make the measurement bandwidth narrower than the PLL bandwidth. A minor drawback is that you need two phase detectors driven in quadrature, since one will be held at null by the operation of the PLL, and you need an in-phase LO for the AM detector. The walking ring counter approach is a good match here, since your LO is usually low enough that garden variety logic ICs can work at $4f_{LO}$. In a real instrument you'll need the second PD for lock detection anyway (see Section 15.6.4).

13.9.4 FM/PM Detectors

FM and PM are indistinguishable if f_m is constant, but they're very different for wide modulation frequency bandwidths. They are generated differently, and you detect them differently too. For small modulation indexes $m \lesssim 0.1$ rad, you can just shove the signal into a phase detector operating near null, and use the IF output as a direct measure of the phase deviation (of course, you have to know the average phase in advance). Although the null is usually pretty stable with signal amplitude, the scale factor will change linearly,

so that there is a danger of severe AM–PM conversion away from null. This problem is typically fixed by putting a limiter ahead of the phase detector.

Where to get the LO signal for an FM detector is a problem just as in the AM case. The usual way is to use an LC phase shift network whose phase varies with frequency (see Chapter 14). For wideband applications, the delay-line discriminator of Section 15.5.12 is useful. This is nothing but a mixer, whose LO is just the (limited) RF signal passed through a delay line made of a piece of coax cable cut to length. Its output voltage is theoretically proportional to $\cos(2\pi f\tau)$, where τ is the cable delay, but because of the nonsinusoidal characteristic of the mixer, it will actually be flattened a little on top and have a slight DC offset.

Phase demodulators are always limited to $|\Delta\phi| < \pi$, and those that operate on sine waves are limited by the nonlinearity of their output characteristic, $V(\phi) \approx V_0 \cos(\phi)$. The linear range can be extended to $\pm\pi/2$ by using square waves fed into a CMOS exclusive-OR gate, such as a 74HC86.[†] After averaging the output with an RC filter, the result is a triangular function of the relative phase, so almost the full π range is available (the finite rise time of the edges of the output waveform make it very nonlinear if you go too close to 0 or π). Edge-sensitive phase detectors can accommodate a range of nearly 2π , also limited by rise time effects. Linearity requires flat-topped signal waveforms. They need not be square, though the steeper the edge, the more of the phase range we can use without running into edge effects. Thus we can get the same linear slope with sine waves by overdriving the inputs of a Gilbert cell multiplier, because overdriving a differential pair produces a roughly trapezoidal signal waveform with a very flat top (assuming that the collector load resistors are small enough that the transistors cut off before they saturate).

For ranges beyond 2π , it is possible to use frequency dividers in both the signal and comparison arms, so that the phase range can be as large as desired, at the expense of time delay[‡] and reduced sensitivity. A more subtle problem with dividers is the phase ambiguity, as we saw in Section 13.5.2. When the power is turned on, a divide-by- N counter can wake up in any one of N states, so the phase detector output can be at any of several places in its range, independent of the real phase relationship between the input signals. Resetting the counters externally can bring the phase detector output to a known state, but this will reflect the true phase shift modulo 2π , which does not fix the phase ambiguity. This problem remains difficult unless it is feasible to bring the two true phases into a known relationship before resetting the counters.

Besides limiters, another approach to eliminating the AM–PM conversion of a mixer is to force it to run at null, by putting it in a feedback loop with a phase shifter, and using the phase shifter control voltage as the output. You can do this in analog, with an op amp doing the feedback, or digitally by successive approximation.[§] (This is a lot quieter than fringe surfing. Why?)

[†]An exclusive-OR gate performs a cross-correlation between its inputs, if the two logic states are assigned numerical values of ± 1 and the result is time averaged.

[‡]Time delay in a phase detector is somewhat subtle and may be of crucial importance in the frequency compensation of phase-locked loops. A digital phase detector's output cannot change until the next clock edge, because there's no phase information in a featureless logic level. On the other hand, sine waves exhibit phase information all the time, so an analog phase detector's output can begin changing immediately. The trade-off is a bit subtler than this because of the delays in any filters required to make the nice sine wave.

[§]P. C. D. Hobbs, High performance amplitude and phase digitizers at 60 MHz. *Rev. Sci. Instrum.* **58**(8), 1518–1522 (August 1987).

13.9.5 Phase-Locked Loops

A more common type of loop involving phase detectors is the phase-locked loop (PLL). PLLs are discussed in more detail in Section 15.6, but it's worth mentioning their use in FM and PM detection. A PLL has a voltage-controlled oscillator, a phase detector, an op amp, and some sort of filter connected in a closed loop. The feedback loop forces the VCO to follow the exact phase (and *a fortiori* frequency) of the input, for modulation frequencies well within its bandwidth. Near the unity gain bandwidth of the loop, the tracking degrades, and well outside the bandwidth, it hardly tracks at all.

You build an FM demodulator from a PLL by making the loop fast and using the VCO control voltage as the output. Since the VCO voltage has a 1:1 mapping to the frequency, this is a complete solution apart from perhaps needing linearizing. Even the linearization problem is not too severe, since it's easy to measure the VCO characteristics with a counter and voltmeter—it's a 1D calibration problem. PLL FM detection is a lowpass filtering operation and is slightly nonlinear if the signal has strong phase modulation near the edges of the loop bandwidth. An important side benefit is that you usually don't need a limiter with a PLL detector, so that the noise degradation at low SNR is avoided. (See Section 14.7.16.)

A PLL PM demodulator is made by making the loop as narrow as reasonably practicable, to reduce the noise impressed onto the oscillator by the loop, and using the phase detector output. This is inherently a highpass filtering operation, although you can sometimes use time gating of the control signal to approach DC-coupled operation. The VCO itself should be intrinsically quiet, so a voltage-controlled crystal oscillator (VCXO) is often used here.

13.9.6 *I* and *Q* Detection

In an interferometer with significantly unequal arms, (13.14) shows that we can get *I* and *Q* by a sleazy trick with the unintuitive name of *modulation-generated carrier*. If we pick *m* to be about 2.6, we get equal amounts of fundamental and second harmonic, and together they contain about 85% of the total AC power. Using a couple of plain AM detectors, we get *I* and *Q* almost free, because of the factors of *j* in the terms of (13.14). This trick is very convenient with a diode laser source, which can be frequency modulated by changing its bias current. It is especially popular in fiber optic sensors.

13.9.7 Pulse Detection

Many sorts of measurements require accurate determination of the height of a pulse. Examples are sizing of aerosol particles, pulsed laser spectroscopy, and time-domain reflectometry. It is tempting to use an analog peak detector circuit, but this is almost always the wrong thing to do, unless the pulses are very long (at least several microseconds), because peak detectors are inaccurate due to overshoot and sensitivity to the exact pulse shape. Unless you have a special situation where your pulses have all different shapes and widths, and what you need is really the maximum height, use a gated integrator instead. There's more in Chapter 14.

The other approach is to sample digitally, and do the peak detection in software. Provided the data rate is not too high, this is a reasonable approach, but sampling a waveform at 14 bits, say, and a 10 MHz sampling rate is difficult and expensive compared with a gated integrator of 30 ns time constant, with a fast comparator and a few logic

gates to do the gating and a slow digitizer to follow. This is especially true if the event you're waiting for is infrequent, so that the 1000 MIPS DSP wastes most of its time idle.

13.10 REDUCING INTERFERENCE AND NOISE

Once your measurement has been moved away from 0 Hz, and the optical and environmental background, drift, and interference have been reduced adequately, the principal remaining way of rejecting noise and interference is by filtering. Sometimes the measurement is at a fixed frequency, in which case everything is easy: just pick a bandpass filter with the right bandwidth and adequate skirt selectivity, and you're almost done. This can be either an actual *LCR* or active filter, or else a bandwidth narrowing instrument such as a lock-in amplifier or signal averager.

Measurements where the carrier frequency uncertainty is large compared to the measurement bandwidth are a bit more difficult. For that, you'll be forced to choose between picking one frequency at a time and throwing the rest away, as in an AM radio, or using more hardware and detecting many frequency bins simultaneously, with a filter bank (analog or digital). It is also possible to increase the signal level while leaving the noise alone, without needing more photons, via pulse techniques such as time gating.

13.10.1 Lock-In Amplifiers

Lock-in amplifiers are described in Section 14.7.17. You probably already know that they're combination narrowband tunable filters and phase-sensitive detectors, designed to have a reasonably large dynamic range. Like all narrowband filters, lock-ins are slow, slow, slow. The speed problem inevitably accompanies the SNR improvement. A system with bandwidth B takes $10/B$ seconds to settle to 12 bits, even in the best case; if the number of data points is proportional to $1/B$ as well, as is the case in continuous-scan measurements, for example, the measurement time increases as $1/B^2$. A measurement that's 20 dB too noisy can take a graduate student's lifetime to complete if the approach to SNR improvement is just to dial down the lock-in bandwidth (this is even more of a problem if you don't have any graduate students). Such a slow measurement is highly vulnerable to drifts and external disturbances, e.g. bumping the table by accident and spoiling an hour's work, or someone opening the lab door and causing a temperature transient. Remember that that 20 dB SNR improvement will cost you a factor of 100 in speed for a fixed number of data points, and a factor of 10,000 for a continuous sweep. See also Section 14.7.4.

13.10.2 Filter Banks

Lock-ins and superhets work very well provided you know your signal frequency in advance. Sometimes this is not the case, as in Doppler shift measurements. What then? One approach is the use of filter banks. The RF signal is split into N separate transmission lines, each with its own sharp bandpass filter, chosen to be a good match for the expected signal bandwidth. The filters are overlapped somewhat, so that signals don't get missed if they happen to land in the hole between two adjacent filters. This gets expensive if you need more than 4 or 5 bands.

At sufficiently low frequency, we can use a fast Fourier transform to turn N samples into N frequency bins; this is an inexpensive way to get a whole lot of channels. The

overlap between bins is controlled by windowing the data (multiplying it by a function that drops off toward the ends of the data run). The filter shapes are the DFT of the window function, so they're all identical, which is very useful. There's more on this in Section 17.5.

13.10.3 Time-Gated Detection

We've talked about improving the SNR by reducing the duty cycle d while keeping the average optical power constant. In that case, we gained SNR as the reciprocal of the duty cycle by simply widening the filter bandwidth, and by d^{-2} by time-gating the receiver as well. If we are using statistics-based methods such as thresholding, the situation is a bit more subtle.

A measurement with a 1% receiver duty cycle will have a false alarm rate 100 times smaller than a continuous measurement under identical conditions. This can really make a difference in photon counting experiments, which are limited by the dark count rate. Where the noise probability density is really exponentially decreasing,[†] this may make only a decibel or two of difference in the threshold, but if there are lots of outliers in the amplitude statistics, the gain can be much greater. If the peak optical power goes up by d^{-1} , the threshold can be raised considerably too.

Often the undesired pulses do not arrive in a Poisson process. In an optical time-domain reflectometer (OTDR), which sends pulses down a fiber and looks for reflections from breaks and discontinuities, there are big pulses at the beginning and the end of the fiber, which we'd like to reject well enough that they don't peg some amplifier and so mask real signals nearby. Time gating is good for that too.

The time gating need not be on-off. Some measurements have a signal strength that varies after a transmit pulse in a reproducible way. Time of flight ranging and optical tomography are good examples. It would be useful to change the system gain with time, so that huge signals from nearby objects can be accommodated along with faint ones from distant objects that arrive later. Variable PMT, MCP, or APD bias voltages can do this.

13.10.4 Signal Averaging

Pulsed measurements have another advantage, which is not always realized: most of the signal energy gets moved away from DC, up to harmonics of the sampling rate. Since most systems have $1/f$ noise and DC drifts at some level, this is an unequivocal benefit unless the $1/f$ noise of the pulses themselves dominates, as it frequently will.

Even CW measurements can benefit from this, through fast scanning plus signal averaging. Say you're making a scanned measurement, such as current-tunable diode laser spectroscopy. Instead of scanning slowly and accumulating a spectrum in one scan of a few seconds' duration, scan at 1 kHz and take N (say, 2000) spectra. Because the signal will be periodic, it will now occur in 0.5 Hz wide bands around each of the harmonics of 1 kHz, and, crucially, stuff happening much below 1 kHz can only produce a baseline shift. If your spectrum needs a frequency resolution of 1000 points, use a filter of 500 kHz to 1 MHz bandwidth, because you'll need at least the first 500 harmonics.

[†]For example, Gaussian or Poissonian.

When these traces are added, the signal will add coherently, and the noise only in power, leading to an SNR improvement of N times. This is equivalent to the slow scan approach for white noise, but with extra low frequency noise, it's a lot more stable and accurate. Do watch out for the temporal response of your filter—try scanning more slowly at first, and look for changes in the shape of the sharp features as you increase the scan rate, or use a triangular scan rather than a sawtooth, and force the retrace to lie on top of the forward scan on the scope. There's a more sophisticated look at this in Section 17.11.5.

13.10.5 Frequency Tracking

There are some signals, such as space-probe telemetry, where the carrier frequency of a narrowband signal varies slowly over many times the signal bandwidth. The usual method is a phase-locked loop with some lock-acquisition aid such as a slow frequency sweep, as in Section 15.6.5.

13.10.6 Modulation-Mixing Measurements

Section 10.9.4 gave some examples of modulation mixing. This trick is so useful that it's worth looking at it in some detail. The goal is to put the desired signal all by itself at some frequency that is easily filtered out from the various drive signals in the system—sort of an electronic version of a dark-field or fringe-based measurement (see Section 10.7.2). The usual method is to modulate the sample at one frequency and the beam at another frequency, then detect the difference frequency, using (13.11).

For instance, you can improve the sensitivity of a bolometer by biasing it with AC instead of DC and chopping the incoming light; signals at $f_{\text{chop}} \pm f_{\text{bias}}$ have to come from the mixing of the two, and you can put them at some distance from the effects of air currents and $1/f$ noise. The desired mixing is detected using a lock-in amplifier or tuned IF filter and detector (as in a superheterodyne radio, see Example 13.6). All the discussion of Section 13.7.2 on how to choose the frequencies to avoid spurious signals applies here as well.

Other examples are the Doppler-free spectrometer of Example 1.1, where the signal was at the sum and difference of the chopping frequencies, and many kinds of pump/probe experiments such as photothermal spectroscopy, where the pump and probe beams are modulated at different frequencies.

13.11 DATA ACQUISITION AND CONTROL

13.11.1 Quantization

A perfect N -bit digitizer will produce a numerical output indicating which of 2^N equal width bins the instantaneous signal voltage fell into at a time exactly coincident with the edge of its clock signal, and ignore what it does the rest of the time. This is obviously a nonlinear operation, and time variant too. How does it fit into our linear, time-invariant network picture? Well, by now you should foresee the answer: it doesn't, but it does nearly, and the errors can be fixed up by hand. The main things to keep straight are Nyquist's theorem and the noise due to quantization. Nyquist's theorem is covered in more detail in Chapter 17, but what it says is that if our signal contains no frequency

components above f , then provided we sample periodically in time at a rate of more than $2f$, we can reconstruct the input signal perfectly from the samples at a later time, by interpolating in a certain way. This gives us confidence that sampling the data does not itself degrade it.

Quantization effects are usually dealt with by treating them as additive noise. Provided our signal swing is at least several ADUs,[†] the main effect of quantization is to contribute additive white noise of RMS voltage $1/\sqrt{12}$ ADU.[‡] This is of course not necessarily true if the signal has large components of the clock signal in it, or if your ADC autoranges. Together these give us a basis for specifying the digitizer subsystem. Real A/D converters are not of course perfect.[§] Unless you plan to do extensive numerical simulations of how different A/D errors contribute to your measurement inaccuracy, mentally reduce the number of bits of a converter by looking at the differential nonlinearity specification. The DNL of a perfect converter is 0, and of a good one is 0.5 ADU or less. If the DNL is 2 ADUs for a 12 bit converter, then what you've got is a 10 bit converter with some nice extra bits that may sometimes mean something, and won't make life worse. Integral nonlinearity and offsets are usually less serious, as something else in your system will probably be nonlinear as well, so calibration is likely to be needed anyway.

It's easy to see that the quantization error can't really be independent of the input signal, but (as the Widrow reference shows) assuming the signal is at least several ADUs in size, quantization noise is spread out pretty evenly through the fundamental interval of $[0, v/2)$. Since the total noise power is fixed, the noise PSD goes down as the bandwidth goes up, making quantization noise relatively less important in wideband systems for a fixed resolution.

Aside: When Is the RMS Error What Matters? So far, we've blithely assumed that the RMS error is what matters. That's generally true when we're taking many measurements where the quantization error isn't likely to be correlated from sample to sample. On the other hand, if a maker of $3\frac{1}{2}$ digit DVMs claimed that they were accurate to 1 part in $2000\sqrt{12}$, that would clearly be specsmanship of the deepest dye; we're not making repetitive measurements of uncorrelated quantities there, and the relevant number is 0.5 ADU, or maybe 1 ADU, even if the digitizer is perfect. Thought is required.

13.11.2 Choosing a Sampling Strategy

Choosing how you'll filter and sample your data is a great deal like choosing an IF, so draw a frequency plan just as we did in Section 13.7.2. If at all possible, pick your sampling clock rate so that strong spurious signals (e.g., harmonics of 50 or 60 Hz from AC-powered equipment, or harmonics of your PWM motor controller clock) don't alias down into your frequency band of interest, and filter as close to the ADC as possible. If you can't avoid these spurs, try to put them where you can filter them out digitally, and pay very close attention to your grounding and shielding (see Chapter 16).

[†]B. Widrow, A study of rough amplitude quantization by means of Nyquist sampling theory. *IRE Trans. Circuit Theory* **3**, 266—276 (1956). A. B. Sripad and D. L. Snyder, A necessary and sufficient condition for quantization errors to be uniform and white. *IEEE Trans. Acoust. Speech Signal Process* **ASSP-25**, 442—448 (1977).

[‡]This noise has a peak-to-peak value of 1 ADU and is therefore not Gaussian.

[§]ADCs and DACs are discussed in some detail in Section 14.8, so look there for unfamiliar terms.

Or perhaps you need to sample a transient accurately, without worrying too much about its exact frequency content. In that case, the sampling theorem is more of a sanity check than a useful design criterion—you'll need to sample significantly faster than you might gather from Nyquist (see Section 17.4.4).

13.11.3 Designing with ADCs

Deciding on the resolution of your ADC is not necessarily trivial, as the following two examples demonstrate.

Example 13.9: CCD Digitizer. For instance, consider a CCD digitizer, to be operated at fairly low speed, say, 150k samples/s. Full scale on a commodity camcorder CCD is about 5×10^4 electrons, and the readout noise can be as small as 20 electrons. We saw above that an ADC contributes RMS noise voltage of $1/\sqrt{12}$ ADU, so in order for the readout noise to dominate the ADC contribution, we'd need

$$N_{\text{bits}} \geq \log_2 \left(\frac{5 \times 10^4}{20\sqrt{12}} \right) \quad (13.61)$$

or 10 bits, at least. Remember that it is not enough that the digitizer noise be smaller than the rest of the noise; allowing one stage to contribute as much noise as the rest of the instrument combined is a foolish cost/benefit choice unless that one is by far the most expensive. To reduce ADC noise power to 0.25 of the shot noise, so that the total SNR is degraded no more than 1 dB by the digitizer, we'll need an 11 bit ADC, which is unavailable. A 12 bit ADC will reduce the noise degradation to 0.26 dB, which is usually good enough. If you need to do this at high speed, that ADC will run into money. (If your photons are very expensive, it may pay to go to even higher ADC resolutions—a quarter of a decibel here and there adds up.)

Example 13.10: Soda Pop Bottles. Now let's say that you're building an inspection system for cola bottling plants. You have lots of light, so there's no reason to slow the sampling rate down, and you anticipate being near half-scale most of the time. Now the photon noise is not 20 electrons, but the square root of half the bucket capacity, or about 160, a factor of 8 higher than in the previous example. That means that an 8 bit digitizer is probably enough for that application. Video rate 8 bit ADCs are cheap enough to give away in cereal boxes, so you've saved some money and probably haven't lost any performance.

13.11.4 Choosing the Resolution

Rules of thumb for choosing a digitizer resolution are often simplistic and inadequate, and this one is no exception, because it assumes that the digitizer's performance is perfect. Ordinary problems like differential and integral nonlinearity are dealt with by mentally reducing the number of bits, but that's just the beginning. Real digitizers have noise and also exhibit serious time variations caused by the signal, so that they contribute spurious signals just as mixers do. If you are using a fast digitizer, and especially if your signal contains frequency components higher than 5% of the device's maximum rated sampling frequency, make sure you find out about its performance with large AC signals.

Many ADCs come with integral T/Hs; these are called sampling ADCs in the catalogs. The integrated T/Hs are often very good and well matched to the ADC behavior. Unless you have a pressing need to use a separate T/H, go for a sampling ADC. It'll just make your life easier. (A good example is putting the multiplexer between the sampling and digitizing steps to get simultaneous sampling of several channels.) Do keep an eye on the aperture uncertainty specification—it's critical but often overlooked. There's lots on that and other ADC specifications and problems in Chapter 14.

The bottom line of all this is that getting good performance out of your digitizer is nontrivial. To get the best measurement for your money, it is very important to make an end-to-end noise and error budget, so that if you can afford to give up 6 dB in sensitivity, you start out by spending most of it reducing the cost of the most expensive part (e.g., going to a lower power laser or reducing your collection NA), and then see how much is left for the rest of the system.

13.11.5 Keep Zero On-Scale

While we're on the subject of ADC resolution, we should note that even a 20 bit ADC has no resolution whatever beyond its voltage limits. Thus it is very important to ensure that all valid signals produce an on-scale reading. Normally people don't screw up the upper limit, but in these days of single-supply analog circuitry, an amazing number will put the zero-signal input to the ADC at 0 V, and not think anything of it. This is a disaster. First of all, you have no way of distinguishing between a zero signal condition and a fault (e.g., loss of the analog power supply). Also, everything has an offset voltage, and nobody can really pull down to 0, so there will be a region near zero signal where you can't get a reliable digital output.

Make sure your signal has a small, accurately known DC offset, enough to keep it on-scale over the worst case offset, supply, and nonlinearity variations (and then add a bit more for luck). Use 1% of full scale as a starting point.

Example 13.11: Error Budget for an Extinction Turbidity Meter. Turbidity meters are used in pollution monitoring, for example, to measure the nutrient concentration of water by looking for turbidity caused by bacteria and algae growth. A simple turbidity meter might use an inexpensive 600 nm LED driven with a 10 kHz square wave as a light source, a small solar cell as a detector, and a signal processor consisting of a transimpedance amp and active bandpass filter leading up to a synchronous AM detector to convert the AC signal into a voltage. The resulting voltage is to be digitized.

We anticipate needing high sensitivity only for small turbidities, since heavily contaminated sludge is not on the menu.

This arrangement should be relatively immune to etalon fringes, due to the very low temporal coherence of the LED source. Collimating the light and forcing it to pass through an aperture on the source side and a black tube baffle on the output side of the sample cell should eliminate false readings due to multiply scattered light. There should be no problems with pickup and $1/f$ noise, since it's working at a high audio frequency. Its time response will be set by the filter bandwidths but could easily be 1 ms. Let's say it is to be 10 ms, 10–90%. An RC lowpass filter has a rise time $t_r = 0.35/f_{-3dB}$, so the bandwidth should be about 35 Hz.

We'll use a CMOS analog gate, driven by the same square wave, as a phase-sensitive detector, followed by a 35 Hz RC lowpass filter. A single-pole RC lowpass has a noise

bandwidth $BW_{\text{noise}} = (\pi/2)f_c$, or 55 Hz. If the peak photocurrent detected at zero turbidity is $100 \mu\text{A}$, then the average is $i_{\text{DC}} = 50 \mu\text{A}$, and the shot noise is 4 pA, a SNR of 141 dB in 1 Hz, or 124 dB in 55 Hz, so that we need at least a 19 bit ADC to preserve the noise performance. With care, the shot noise is easily reachable in a system like this, so the choice appears to be using a 20 bit ADC (slow ones exist) or allowing ourselves to be ADC limited. In the ADC-limited case, we should vary the drive current a bit between samples, because unless the signal varies by at least a few ADUs, quantization noise may not be additive. Since the sensitivity level required is not all that high, being ADC limited may not be so bad here.

There is another way of looking at this problem, though. We'll need to normalize the detected signal. The simplest way to do this is to send part of the light through twice, perhaps by mirror-coating part of the flow cell; a better way, which will help cancel effects due to gradual fouling of the windows, is to make the flow cell wedge-shaped and send two beams through regions of different width. The two photocurrents can be divided, perhaps by using the once-dimmed beam to produce the ADC reference current for digitizing the twice-dimmed one,[†] yielding a ratiometric measurement in which zero turbidity produces a (nearly) full scale output. Alternatively, they can be subtracted, with the beam strengths adjusted to null the resulting signal at zero turbidity; a zero-based signal results. This zero-based signal can be amplified and then digitized at much lower resolution, because the ADC resolution limits the maximum signal, rather than the sensitivity. (Remember to keep zero and maximum signal slightly inside the ends of the ADC range.) Very turbid water could be dealt with by having two ranges or by digitizing the logarithm instead (see Section 15.5.9), but multiple scattering would start to dominate, so another approach would probably be required anyway.

[†]Why does this work?

Electronic Building Blocks

And then of course I've got this terrible pain in all the diodes down my left-hand side... I mean, I've asked for them to be replaced but nobody ever listens.

—Marvin the Paranoid Android, in *The Hitchhiker's Guide to the Galaxy* by Douglas Adams

14.1 INTRODUCTION

The subject of electronic components and fundamental circuits is a whole discipline unto itself and is very ably treated in lore-rich electronics books such as Horowitz and Hill's *The Art of Electronics*, Gray and Meyer's *Analysis and Design of Analog Integrated Circuits*, and (for a seat-of-the-pants approach to RF work) the *ARRL Handbook*. If you're doing your own circuit designs, you should own all three.

There's more to it than competence in basic circuit design. Electro-optical instruments present a different set of problems from other signal processing applications, and that means that there's lots of stuff—basic stuff—that hasn't been thought of yet or hasn't been taken seriously enough yet. That's an opportunity for you to do ground-breaking electronic design work, even if you're not primarily a circuit designer; not everybody with a circuit named after him can walk on water, after all, even in the more highly populated fields.[†]

If we can gain enough intuition and confidence about the basic characteristics of circuitry, we can try stepping out of the boat ourselves. Accordingly, we'll concentrate on some important but less well known aspects that we often encounter in electro-optical instrument design, especially in high dynamic range RF circuitry. You'll need to know about Ohm's law, inductive and capacitive reactance, series and parallel *LC* resonances, phase shifts, feedback, elementary linear and digital circuit design with ICs, and what FETs, bipolar transistors and differential pairs do.

The result is a bit of a grab-bag, but knowing this stuff in detail helps a great deal in producing elegant designs.

[†]Jim Williams's books have a lot of lore written by people like that and are worth reading carefully.

14.2 RESISTORS

Resistors are conceptually the simplest of components. A lot of nontrivial physics is involved, but the result is that a good resistor has very simple behavior. A metal resistor obeys Ohm's law to extremely high accuracy (people have mentioned numbers like 0.1 ppm/V, a second-order intercept point of +153 dBm for 50 Ω —don't worry about it). In addition, its noise is very well described, even at frequencies below 1 Hz, by the Johnson noise of an ideal resistor. Not bad for a one-cent part.

There are only a couple of things to remember about metal film resistors. The first is that the cheap ones have bigish temperature coefficients, as high as 200 ppm/ $^{\circ}$ C. This amounts to 1% drift over 50 $^{\circ}$ C, which may be a problem, especially if you're using trims. You can get 50 ppm/ $^{\circ}$ C for only a bit more money, and you can do much better (1–5 ppm/ $^{\circ}$ C) with expensive resistors. Clever use of cheap resistor arrays can achieve voltage divider ratios with this sort of stability as well. The other thing is that they are allowed to display some amount of $1/f$ noise when biased with high voltages: the specs allow rms noise from this source of 10^{-7} times the bias voltage applied, when measured in one decade of bandwidth. This is significantly more obnoxious than the nonlinearity; a 10 k Ω resistor with 15 V across it could have 10–100 Hz noise voltage of around 1.5 μ V rms, corresponding to a noise current of 150 pA in the same bandwidth, 22 dB above the Johnson noise. In practice, the noise is nowhere near this large. The author measured some 11 k Ω RN60C type resistors in a 2:1 voltage divider with a 9 V battery and found that the coefficient was lower than 5×10^{-9} per decade all the way down to 1 Hz, which was the resolution limit of the measurement.

Other resistor types, such as metal oxide, carbon film, conductive plastic, cermet, and carbon composition, are worse in one or more ways. For example, metal oxide resistors (the cheapest 1% units available) have huge thermoelectric coefficients, so that a temperature gradient will produce mysterious offset voltages and drifts. Like everything else, resistors have some parasitic capacitance; an ordinary $\frac{1}{8}$ W axial-lead resistor has about 0.25 pF in parallel with it.

14.2.1 Resistor Arrays

You can get several resistors deposited on a common substrate in an IC package. These are wired either completely separately or all having one common pin. Arrays typically come in a 2% tolerance, but the matching between devices is often much better than that, and the relative drift is very small indeed, because the temperature coefficients are very similar and the temperature is fairly uniform across the array. The matching becomes even better if you use a common centroid layout, as in Figure 14.1. From an 8 resistor array, we can make two matched values by wiring #1, 4, 5, and 8 in series to make one, and 2, 3, 6, and 7 to make the other. This causes linear gradients in both temperature and resistor value to cancel out. (Note: The absolute resistance isn't stabilized this way, only the ratios.) You can make four matched values using #1 + 8, 2 + 7, 3 + 6, and 4 + 5, or even a reasonable 3:1 ratio using 3 + 6 versus the rest, though these fancier approaches may not be quite as good, because there's liable to be a second-order (center vs. edge) contribution beside the gradients.




Figure 14.1. Common centroid layout for resistors in series: the two series strings will match closely because the layout ensures that linear gradients in surface resistivity cancel.

14.2.2 Potentiometers

A potentiometer is a variable resistor with three terminals: the two fixed ends of the resistive element and a wiper arm that slides up and down its surface. Because the resistance element is a single piece, the voltage division ratio of a pot is normally very stable, even though the total resistance may drift significantly.

The wiper is much less reliable than the ends. Try not to draw significant current from it, and whatever you do, don't build a circuit that can be damaged (or damage other things) if the wiper of the pot bounces open momentarily, because it's going to (see Example 15.2).

Panel mount pots are getting rarer nowadays, since microcontroller-based instruments use DACs instead, but they're still very useful in prototypes. Conductive plastic pots adjust smoothly—if you use them as volume controls they don't produce scratchy sounds on the speaker when you turn the knob—and have long life (>50k turns), but are very drifty— $\text{TCR} \approx 500\text{--}1000 \text{ ppm}/^\circ\text{C}$ —and have significant $1/f$ noise. Cermet ones are noisier when you twist them, and last half as many turns, but are much more stable, in the $100\text{--}200 \text{ ppm}/^\circ\text{C}$ range. Wire-wound pots also exist; their temperature coefficient is down around $50 \text{ ppm}/^\circ\text{C}$, but since the wiper moves from one turn to the next as you adjust it, the adjustment goes in steps of around $0.1\text{--}0.2\%$ of full scale, like a mechanical DAC. Their lower drift makes wire-wound pots much better overall.

14.2.3 Trim Pots

Trim pots, which are adjusted with a screwdriver, are much more common. They come in single-turn styles, where the shaft moves the wiper directly, and multiturn, which move the wiper with a screw. Stay away from the multiturn styles when possible. Their settability is no better in reality, and their stability is much poorer. If you need accurate fine adjustment, put low and high value pots in series or parallel, or use the loaded pot trick. Trimmers have much larger minimum resistances, as large as 5% of the end-to-end value, and won't stand much twiddling: typically their specs allow the resistance to change 15% in only 10 complete cycles.




Figure 14.2. Loaded pot.

14.2.4 Loaded Pots

You can do some interesting things by putting small resistors between the ends of the pot and the wiper, the so-called *loaded pot*, as shown in Figure 14.2. A $1\ \text{k}\Omega$ pot with a $10\ \Omega$ resistor across each side has the same adjustment sensitivity as a 51 turn, $20\ \Omega$ pot, but is dramatically more stable. Big twists of the shaft cause small changes in balance, until you get right near one edge, when the pot starts to take over. Loaded pots are great for offset adjustments, another example of the virtue of a vernier control. A popular parlor game is figuring out new nonlinear curves you can get from a pot and a few resistors; it's a fun challenge, like wringing new functions out of a 555 timer, and is occasionally useful.

14.3 CAPACITORS

Capacitors are also close to being perfect components. Although not as stable as resistors, good capacitors are nearly noiseless, owing to their lack of dissipative mechanisms. This makes it possible to use capacitors to subtract DC to very high accuracy, leaving behind the residual AC components. The impedance of an ideal capacitor is a pure reactance,

$$Z_C = \frac{1}{j\omega C}, \quad (14.1)$$

where the capacitance C is constant.

Aside: Impedance, Admittance, and All That. So far, we've used the highly intuitive idea of Ohm's law,[†] $V = IR$, for describing how voltages and currents interact in circuits. When we come to inductors, capacitors, and transmission lines, we need a better

[†]Purists often call the voltage E (for electromotive force) rather than V , but in an optics book we're better off consistently using E for electric field and V for circuit voltage.

developed notion. A given two-terminal network[†] will have a complicated relationship between I and V , which will follow a constant-coefficient ordinary differential equation in time. We saw in Chapter 13 that a linear, time-invariant network excited with an exponential input waveform (e.g., $V(t) = \exp(j\omega t)$) has an output waveform that is the same exponential, multiplied by some complex factor, which was equivalent to a gain and phase shift. We can thus describe the frequency-domain behavior of linear two-terminal networks in terms of their impedance Z , a generalized resistance,

$$V = IZ. \quad (14.2)$$

The real and imaginary parts of Z are $Z = R + jX$, where X is the reactance. A series combination of R_1 and R_2 has a resistance $R_1 + R_2$, and series impedances add the same way. Parallel resistances add in conductance, $G = G_1 + G_2$, where $G = 1/R$, and parallel impedances add in admittance, $Y = 1/Z$. The real and imaginary parts of the admittance are the conductance G and the susceptance B : $Y = G + jB$. (See Figure 14.3.)

Whether to use admittances or impedances depends on whether you're describing series or parallel combinations. The two descriptions are equivalent at any given frequency. Just don't expect X to be $-1/B$ or R to be $1/G$ for complex impedances, or that series RC circuits will have the same behavior as parallel ones over frequency.

The linear relationship between AC voltage and current in a reactive network is sometimes described as "Ohm's law for AC," but the physics is in fact quite different; Ohm's law arises from a large carrier density and many dephasing collisions in materials, whereas the linearity of Maxwell's equations means that reactance is still linear even in situations where resistance is not (e.g., in an electron beam or a superconductor).

14.3.1 Ceramic and Plastic Film Capacitors

The two most popular types of capacitor are ceramic and plastic film. Ceramic dielectrics come with cryptic names such as NPO, X7R, and Z5U. Table 14.1 should help sort out the confusion. Plastic film and low- k [‡] monolithic ceramic types are very linear and will




Figure 14.3. Series and parallel equivalent circuits for an impedance $R + jX$. Note that these only work at a single frequency due to the frequency dependence of X .

[†]Two-terminal is the same as one-port, that is, one set of two wires.

[‡]For historical reasons, we refer to the relative dielectric constant ϵ as the relative permittivity k when discussing capacitors. They're the same thing, don't worry.

TABLE 14.1. Common Dielectrics Used in Capacitors

Dielectric	Tol	TC	Range	(ppm/° C)
Ceramics				
NPO/C0G	5%	0 ± 30	1 pF–10 nF	Good for RF and pulse circuits; best TC, stable, low dielectric absorption.
X7R	10%	–1600	1 nF–1 μF	Like a cheap film capacitor; TC very nonlinear.
Z5U	20%	+10 ⁴	10 nF–2.2 μF	Good for bypassing and noncritical coupling, but not much else.
Silver mica	5%	0±100	1 pF–2.2 nF	Excellent at RF, bad low frequency dielectric absorption.
Plastic film				
Polyester (mylar)	10%	700	1 nF–10 μF	General purpose. Very nonlinear TC—it's cubic-looking.
Polycarbonate	10%	+100		Similar to polyester
Polypropylene	1–5%	–125	47 pF–10 nF	Best for pulses; very low leakage, loss, dispersion, and soakage
Teflon	5%			Unsurpassed low leakage and dielectric absorption, good for pulses
Electrolytics				
Aluminum	+80/–20		1 μF–1.0 F	big C, but slow, noisy, lossy, short-lived and vulnerable to frost.
Solid tantalum	20%		0.1 μF–220 μF	Better than wet Al; low ESR and ESL; big ones are very expensive.

cause you no problems apart from temperature drift and possibly a bit of series inductance. Polypropylene especially is an inexpensive material with really good dielectric properties and decent high temperature performance.

In order to pack a lot of capacitance into a small volume, you need many very thin layers of high- k material. Unfortunately, to get high k values, you have to trade off everything else you care about: stability, low loss, linearity, and so on. An X7R ceramic is allowed to change its value by 1%/V. This will cause serious distortion when X_C is a significant contributor to the circuit impedance. A Z5U ceramic's capacitance can go up by half between 25 °C and 80 °C. High- k ceramic capacitors can even become piezoelectric if mistreated, which is really fun to debug (look for weird reactance changes below about 1 MHz).

Capacitor tempco can also be a source of drift. An RC with bias voltage V_{bias} will produce a thermal offset of

$$\Delta V = V_{\text{bias}} R \frac{dC}{dT} \frac{dT}{dt}, \quad (14.3)$$

which can be serious, for example, if $dT/dt = 10$ mK/s, a 1 second TC using an X7R capacitor with 2.5 V across it can produce a 40 μV drift offset.

The wisdom here is to let high- k devices do what they're good at, so in high level circuits, use values large enough that you don't have to care how much they vary.

Since the big shift to surface mount, plastic film capacitors are not as common as they were. The reason is that the capacitor body gets very hot in surface mount reflow soldering. The good film caps come in through-hole only. Newer plastics such as polyphenylene sulfide (PPS) caps come in SMT, but they're not the most reliable things.

Metallized plastic film capacitors are especially useful in situations where transient overvoltages may occur, since if the dielectric arcs over, the arc erodes the metal film in the vicinity of the failure until the arc stops, which amounts to self-repair. On the other hand, film-and-foil capacitors have more metal in them, so their ohmic (I^2R) losses are much lower.

14.3.2 Parasitic Inductance and Resistance

A capacitor doesn't really discharge instantaneously if you short it out; there is always some internal inductance and resistance to slow it down. Inductive effects are mostly due to currents flowing along wires and across sheets, and are lumped together as *effective series inductance* (ESL). Resistive effects come from dielectric absorption and from ohmic losses in the metal and electrolyte (if any). Each charge storage mechanism in a given capacitor will have its own time constant. Modeling this would require a series *RC* for each, all of them wired in parallel; since usually one is dominant we use a single *RLC* model with *C* in series with a single *effective series resistance* (ESR) and ESL.

If you need really low ESR and ESL, e.g. for pulsed diode laser use, consider paralleling a whole lot of good quality film-and-foil capacitors.

14.3.3 Dielectric Absorption

Those other contributions we just mentioned are actually the main problem with capacitors in time-domain applications such as pulse amplifiers, charge pumps, and track/holds. These need near-ideal capacitor performance, and those parallel contributions (collectively known as dielectric absorption, or *soakage*) really screw them up. The imaginary (dissipative) part of the dielectric constant contributes to the ESR and thus puts an exponential tail on part of the discharge current. The tail can be very long; an aluminum electrolytic may still be discharging after 1 s, which would make a pretty poor track/hold. Choose Teflon, polypropylene, or NPO/COG ceramic dielectrics for that sort of stuff, and whatever you do, don't use high-*k* ceramics, micas, or electrolytics there. Polypropylene and silicon oxynitride (AVX) capacitors are available in 1% tolerances, so they're also very useful in passive filters and switched-capacitor circuits.

14.3.4 Electrolytic Capacitors

Electrolytic capacitors have a lot of capacitance in a small volume, which makes them useful as energy storage devices (e.g., power supply filters). Capacitors of a given physical size have a roughly constant *CV* product; thinning the dielectric increases the capacitance but decreases the working voltage, so that *CV* remains constant. The energy in the capacitor is $\frac{1}{2}CV^2$, so for energy storage you win by going to high voltage.

They work by forming a very thin oxide layer electrochemically on the surface of a metal; because the dielectric forms itself, it can be extremely thin. Electrolytics can be wet (aluminum) or dry (solid tantalum). Dry electrolytics are better but more expensive. Electrolytics exhibit leakage, popcorn noise, huge TCs, and severe dielectric absorption,

so keep them out of your signal paths. The ESR of an aluminum electrolytic climbs spectacularly when the electrolyte freezes.

Ordinary electrolytics are polarized, so that they must not be exposed to reverse voltages. Nonpolarized electrolytics exist but cost a lot and aren't that great. If you reverse an electrolytic capacitor in a low impedance circuit at appreciable voltage, it may explode. Old-style metal can electrolytics were known as *confetti generators* because of this behavior. Solid tantalum capacitors can be very touchy if mistreated, especially by reversal or excessively large current transients (e.g., on the input side of a voltage regulator) that cause hot spots. They go off like solid-fuel rockets, with the sintered metal being the fuel and the manganese dioxide electrolyte the oxidizer—they burn hot and generate shrapnel. Solid aluminum and polymer electrolytics don't share this problem. You needn't avoid tantalums altogether—they have very low impedances at high frequency and don't misbehave at low temperatures the way aluminum electrolytics can. Just treat them gently.

Electrolytics become electrically very noisy when they are reversed, which can be a good diagnostic test. They also can generate tens to hundreds of millivolts of offset on their own—far worse than the soakage in other types. (Of course, being electrochemical cells, they have more than a passing resemblance to batteries.) Don't use electrolytics in the signal or bias paths of any sort of low level, low frequency measurement.

Aside: Eyeball Capacitor Selection. Two capacitors that look identical may be very different in their properties, but there are a couple of rules of thumb: first, caps that are small for their value are probably high- k ceramics, and second, tight accuracy specs (J or K suffix) go with better dielectrics such as COG, polypropylene, or Teflon. All polarized electrolytics have a polarity marking, of course.

14.3.5 Variable Capacitors

You almost never see panel-mount variable capacitors any more, except in very high power applications (e.g., transmitters or RF plasma systems). This is something of a shame, since many of them were works of art, with silver-plated vanes beautifully sculptured to give linear frequency change with shaft angle. Tuning is done with varactor diodes, whose capacitance changes with bias voltage. Adjusting different stages for tracking is now done with DACs or trim pots, setting the slope and offset of the varactor bias voltage.

What we do still use is trimmer capacitors, based on single vanes made of metallization on ceramic plates, whose overlap depends on shaft angle. These have most of the same virtues and vices of NPO and N750 fixed capacitors. The figure of merit for tuning applications is the capacitance ratio C_{\max}/C_{\min} , since $\omega_{0\max}/\omega_{0\min} = (C_{\max}/C_{\min})^{1/2}$ for a fixed inductance.

14.3.6 Varactor Diodes

As we saw in Section 3.5.1, a PN junction has a built-in E field, which sweeps charge carriers out of the neighborhood of the junction, leaving a depletion region. It is the free carriers that conduct electricity, so the depletion region is effectively a dielectric. How far the carriers move depends on the strength of the field, so that applying a reverse bias causes the depletion region to widen; since the conducting regions effectively move

apart, the junction capacitance drops. Devices that are especially good at this are called varactors, but it occurs elsewhere; a PIN photodiode's C_d can drop by $7\times$ or more when the die is fully depleted.[†]

By steepening the doping profile in the junction region, C_{\max} can be increased, increasing the CR to as much as 15 in hyperabrupt devices, with Q values in the hundreds up to at least 100 MHz for lower capacitance devices (the Q gets better at higher voltages). NXP and Zetex are the major suppliers of varactors below 1 GHz—check out the BBY40 or MMBV609 and their relatives. The tuning rate is highest near 0 V—empirically, C goes pretty accurately as $\exp(-0.4V)$ for hyperabrupt devices and $(V + V_0)^{-2}$ for abrupt-junction ones[‡] (these approximations are good to a few percent over the full voltage range). Curve fitting can get you closer, but since you'll use digital calibration (and perhaps a phase-locked loop) anyway, this accuracy is usually good enough.

Keep the tuning voltage above 2 V or so for best Q and linearity, and watch out for parametric changes (where the signal peaks detune the varactor), which cause distortion and other odd effects. These parametric changes can be greatly reduced by using two matched varactors in series opposing (many are supplied that way), so when the signal increases one, it decreases the other, and the series combination stays nearly constant. Note also that the tuning voltage has to be very quiet.

The tuning curve of a varactor can be linearized by strategically chosen series and shunt inductances, as we'll see in Section 15.12.2, and its temperature coefficient is highish, about $+150 \text{ ppm}/^\circ\text{C}$ in ordinary ($\text{CR} = 2$) devices to $+400 \text{ ppm}/^\circ\text{C}$ in hyperabrupts.

14.3.7 Inductors

An inductor in its simplest form is just an N -turn helical coil of wire, called a solenoid. The \mathbf{B} fields from each turn of the helix add together, which increases the magnetic energy stored (the energy density goes as B^2) and so the inductance goes as N^2 . (Numerically the inductance is the coefficient L of the energy storage equation $\mathcal{E} = 1/2 LI^2$, and \mathcal{E} is proportional to the volume integral of B^2 , so inductance isn't hard to calculate if you need to.)

Inductance can be increased by a further factor of 1.2 to 5000 by winding the coil around a core of powdered iron or ferrite, with the amount of increase depending on how much of the flux path is through air ($\mu_r = 1$) and how much through the magnetic material ($\mu_r = 4$ to 5000). This reduces ohmic losses in the copper, as well as helping to confine the fields (and so reducing stray coupling). The hysteresis losses in the core itself are normally much smaller than the copper losses would otherwise be. Unfortunately, μ is a strong function of B for ferromagnetic materials, so adding a core compromises linearity, especially at high fields, where the core will saturate and the inductance drop like a stone. In a high power circuit, the resulting current spikes will wreak massive destruction if you're not careful.

Toroidal coils, where the wire is wound around a doughnut-shaped core, have nearly 100% of the field in the ferrite, and so have high inductance together with the lowest loss and fringing fields at frequencies where cores are useful (below about 100 MHz). They are inconvenient to wind and hard to adjust, since the huge permeability of the closed-loop core makes the inductance almost totally independent of the turn spacing.

[†]If you're using a tuned photodiode front end, you can use the bias voltage as a peaking adjustment.

[‡]This characteristic produces nice linear tuning behavior, if the varactor is the only capacitance in an LC circuit.

There is some stray inductance; due to the helical winding, an N -turn toroid is also a 1-turn solenoid with a core that's mostly air. Clever winding patterns can more or less eliminate this effect by reversing the current direction while maintaining the helicity.

Pot cores interchange the positions of the core and wire (a two-piece core wraps around and through a coil on a bobbin) and are good for kilohertz and low megahertz frequencies where many turns are needed. Pot cores are easily made with a narrow and well-controlled air gap between the pole pieces. The value of $1/B$ goes as the integral of $(1/\mu) \cdot ds$ along the magnetic path, so a $100 \mu\text{m}$ gap can be the dominant influence on the total inductance of the coil. This makes the inductance and saturation behavior much more predictable and stable, and allows the inductance to be varied by a small magnetic core screwed into and out of the gap region.

High frequency coils are often wound on plastic forms, or come in small ceramic surface mount packages, with or without magnetic cores. The plastic ones usually have adjustable cores that screw into the plastic. Make sure the plastic form is slotted, because otherwise the high thermal expansion of the plastic will stretch the copper wire and make the inductance drift much worse.

A perfect inductor has a reactance $Z_L = jX_L = j\omega L$, which follows from its differential equation, $V = L dI/dt$. Inductors are lossier (i.e., have lower Q^\dagger) than capacitors, due to the resistance of the copper wires, core hysteresis, and eddy currents induced in nearby conductors. Loss is not always bad; it does limit how sharp you can make a filter, but lossy inductors are best for decoupling, since they cause less ringing and bad behavior. Ferrite beads are lossy tubular cores that you string on wires. While they are inductive in character at low frequency, they have a mainly resistive impedance at high frequency, which is very useful in circuits. Ferroxcube has some nice applications literature on beads—note especially the combination of their 3E2A and 4B beads to provide a nice 10–100 Ω resistive-looking impedance over the whole VHF and UHF range.

If you need to design your own single-layer, air-core, helical wire coils, the inductance is approximately

$$L(\mu\text{H}) \approx \frac{a^2 N^2}{9a + 10b}, \quad (14.4)$$

where N is the total number of turns, a is the coil radius, and b is its length, both in inches.[‡] This formula is claimed to be accurate to 1% for $b/a > 0.8$ and wire diameter very small compared to a . It is very useful when you need to get something working quickly, for example, a high frequency photodiode front end or a diplexer (see the problems in the Supplementary Material). For high frequency work, the pitch of the helix should be around 2 wire diameters, to reduce the stray capacitance and so increase the self-resonant frequency of the inductor. You can spread out the turns or twist them out of alignment to reduce the value or squash them closer together to increase it; this is very useful in tuning prototypes. Make sure you get N right, not one turn too low (remember that a U-shaped piece of wire is 1 turn when you connect the ends to something). Thermal expansion gives air-core coils TCs of + 20 to + 100 ppm/°C depending on the details of their forms and mounting.

[†]Roughly speaking, Q is the ratio of the reactance of a component to its resistance. It has slightly different definitions depending on what you're talking about; see Section 14.3.9.

[‡]Frederick Emmons Terman, *Radio Engineers' Handbook*, McGraw-Hill, New York, 1943, pp. 47–64, and *Radio Instruments and Measurements*, US National Bureau of Standards Circular C74, 1924, p. 253.

Inductance happens inadvertently, too; a straight d centimeter length of round wire of radius a centimeters in free space has an inductance of about

$$L = (2 \text{ nH})d \left(\ln \left(\frac{d}{a} \right) - 0.16 \right), \quad (14.5)$$

which for ordinary hookup wire (#22, $a = 0.032$ cm) works out to about 7 nH for a 1 cm length.

14.3.8 Variable Inductors

Old-time car radios were tuned by pulling the magnetic cores in and out of their RF and LO tuning inductors.[†] Like capacitors, though, most variable inductors are for trimming; high inductance ones are tuned by varying the gap in a pot core, and low inductance ones by screwing a ferrite or powdered iron slug in and out of a solenoid.

You usually need trimmable inductors to make accurately tuned filters, but surface mount inductors are becoming available in tolerances as tight as 2%, which will make this less necessary in the future. This is a good thing, since tuning a complicated filter is a high art (see Section 16.11).

14.3.9 Resonance

Since inductors and capacitors have reactances of opposite signs, they can be made to cancel each others' effects, for example, in photodiode detectors where the diode has lots of capacitance. Since X_L and X_C go as different powers of f , the cancellation can occur at only one frequency for a simple LC combination, leading to a resonance at $\omega_0 = (LC)^{-1/2}$. If the circuit is dominated by reactive effects, these resonances can be very sharp. We'll talk about this effect in more detail in Chapter 15, because there are a lot of interesting things you can do with resonances, but for now they're just a way of canceling one reactance with another.

14.3.10 L-Networks and Q

One very useful concept in discussing reactive networks is the ratio of the reactance of a component to its resistance, which is called Q . In an LC tank circuit (L and C in parallel, with some small r_s in series with L), Q controls the ratio of the 3 dB full width to the center frequency ω_0 , which for frequencies near ω_0 is

$$Q \approx \frac{\omega_0 L}{R_s} = \frac{1}{\omega_0 R_s C} = \sqrt{\frac{L}{R_s^2 C}}. \quad (14.6)$$

If we transform the LR series combination to parallel, the equivalent shunt resistor R_p is

$$R_p = R_s(Q^2 + 1). \quad (14.7)$$

[†]For readers not old enough to remember, these were the first pushbutton radios—you'd set the station buttons by tuning in a station, then pulling the button far out and pushing it back in; this reset the cam that moved the cores in and out. It was a remarkable design for the time; you didn't have to take your eyes off the road to change stations.




Figure 14.4. In the high- Q limit, near resonance, an L-network multiplies R_s by $Q^2 + 1$ to get the equivalent R_p . The three forms have very differently shaped skirts.

This is an L-network impedance transformer, the simplest of matching networks. We'll meet it again in Chapters 15 and 18. (See Figure 14.4.)

Aside: Definitions of Q and f_0 . The competing definitions of Q and the resonant frequency f_0 make some of the literature harder to read. The most common is the reactance ratio, $Q = X_L/R = \omega L/R$ (assuming that the loss is mostly in the inductor, which is usually the case), but the center frequency to 3 dB bandwidth ratio is also commonly quoted ($Q = f_0/\text{FWHM}$). These are obviously different, because X_L is not constant in the bandwidth. Similarly, in a parallel-resonant circuit f_0 may be taken to be (a) the point of maximum impedance; (b) the point where the reactance is zero; or most commonly, (c) the point where $X_L = X_C$, which is the same as the series resonance of the same components. At high Q , all of these are closely similar, but you have to watch out at low Q .[†]

14.3.11 Inductive Coupling

It's worth spending a little time on the issue of inductive coupling, because it's very useful in applications. Figure 14.5 shows two inductors whose \mathbf{B} fields overlap significantly. We know that the voltage across an inductor is proportional to $\partial\Phi/\partial t$, where Φ is the total magnetic flux threading the inductor's turns, and the contributions of multiple turns add. Thus a change in the current in L_1 produces a voltage across L_2 . The voltage is

$$V_{2M} = M \frac{\partial I_1}{\partial t}, \tag{14.8}$$




Figure 14.5. Coupled inductors.

[†]See F. E. Terman, *Radio Engineer's Handbook* (1943), pp. 135–148.

where the constant M is the mutual inductance. The ratio of M to L_1 and L_2 expresses how tightly coupled the coils are; since the theoretical maximum of M is $(L_1 L_2)^{1/2}$, it is convenient to define the *coefficient of coupling* k^\dagger by

$$k = \frac{M}{\sqrt{L_1 L_2}}. \quad (14.9)$$

Air-core coils can have $k = 0.6$ – 0.7 if the two windings are on top of each other, ≈ 0.4 if they are wound right next to each other, with the whole winding shorter than its diameter, and ≈ 0.2 or so if the winding is longer than its diameter. (Terman has lots of tables and graphs of this.) Separating the coils decreases k . Coils wound together on a closed ferromagnetic core (e.g., a toroid or a laminated iron core) have very high coupling—typical measured values for power transformers are around 0.99989. The coefficient of coupling K between resonant circuits is equal to $k(Q_1 Q_2)^{1/2}$, so that the coupling can actually exceed 1 (this is reactive power, so it doesn't violate energy conservation—when you try to pull power out, Q drops like a rock).

14.3.12 Loss in Resonant Circuits

Both inductors and capacitors have loss, but which is worse depends on the frequency and impedance level. For a given physical size and core type, the resistance of a coil tends to go as L^2 (longer lengths of skinnier wire), so that even if we let them get large, coils get really bad at low frequencies for a constant impedance level; it's hard to get Q values over 100 in coils below 1 MHz, whereas you can easily do 400 above 10 MHz. The only recourse for getting higher Q at low frequency is to use cores with better flux confinement and higher μ (e.g., ferrite toroids and pot cores).

Capacitors have more trade-offs available, especially the trade-off of working voltage for higher C by using stacks of thin layers; the ESR of a given capacitor type is usually not a strong function of its value. The impedance level at which inductor and capacitor Q values become equal thus tends to increase with frequency.

14.3.13 Temperature Compensating Resonances

The k of some ceramics is a nice linear function of temperature, and this can be used for modest-accuracy temperature compensation. Capacitors intended for temperature compensation are designated N250 (-250 ppm/ $^\circ\text{C}$), N750, or N1500; the TCs are accurate to within 5% or 10%, good enough to be useful but not for real precision. A resonant circuit can be temperature compensated by using an NPO capacitor to do most of the work, and a small N750 to tweak the TC (see Problem 14.3 at <http://electrooptical.net/www/beos2e/problems2.pdf>). In a real instrument you'd use a varactor controlled by an MCU as part of a self-calibration strategy, but for lab use, N750s can be a big help.

14.3.14 Transformers

The normal inductively coupled transformer uses magnetic coupling between two or more windings on a closed magnetic core. The magnetization of the core is large, because it is

[†]Yes, we're reusing k yet again, for historical reasons.

the \mathbf{B} field in the core threading both the primary and secondary windings which produces the transformer action (see Section 14.5.4 for a different transformer principle). Because of the tight confinement of the field in the core, k is very high: around 0.999 or better in low frequency devices. At RF, μ is smaller, so the attainable coupling is reduced. The strong coupling allows very wide range impedance transformations with very low loss: an $N:M$ turns ratio gives an $N^2:M^2$ impedance ratio. For instance, if we have a 10-turn primary and a 30-turn secondary, connecting $50\ \Omega$ across the secondary will make the primary look like a $5.5\ \Omega$ resistor over the full bandwidth of the transformer. The lower frequency limit is reached when the load impedance starts to become comparable to the inductive reactance; the upper limit, when interwinding capacitance, core losses, or copper losses start to dominate. In an ordinary transformer the useful frequency range can be 10 or even 100 to 1. The phase relationships of transformer windings matter: if there are multiple windings, connecting them in *series aiding* will produce the sum of their voltages; *series opposing*, the difference.

14.3.15 Tank Circuits

A tapped inductor functions like a transformer with the primary and secondary wired in series aiding. It looks like an AC voltage divider, except that the mutual inductance makes the tap point behave very differently; a load connected across one section looks electrically as though it were a higher impedance connected across the whole inductor. A parallel resonant circuit with a tapped inductor or a tapped capacitor can be used as an impedance transformation device; the tapped capacitor has no M to stiffen the tap, of course, so for a given transformation ratio you have to use smaller capacitances and therefore a higher Q than with a tapped inductor. Such a network is called a *tank circuit* because it functions by storing energy. Terman discusses tanks and other reactive coupling networks at length.

14.4 TRANSMISSION LINES

A transmission line is a two-port network whose primary use is to pipe signals around from place to place. You put the signal into one end, and it reappears at the other end, after a time delay proportional to the length of the line. Coaxial cable is the best known transmission line, but there are others, such as metal or dielectric waveguide, twin lead, and microstrip (like PC board traces over a ground plane). Like light in a fiber, the signal can be reflected from the ends of the transmission line and bounce around inside, if it is not properly terminated.

Transmission lines can be modeled as a whole lot of very small series inductors alternating with very small shunt capacitors. The line has an inductance L and capacitance C per unit length. The two define a characteristic impedance Z_0 ,

$$Z_0 = \sqrt{\frac{L}{C}}, \quad (14.10)$$

which for a lossless line is purely real.[†] An infinite length of line looks electrically exactly like a resistor of value Z_0 connected across the input terminals. If we sneak in

[†]There is a good analogy between transmission line impedance and the refractive index in optics. The formulas for reflection and transmission at a transmission line discontinuity are the same as the normal-incidence forms of the Fresnel formulas, when Z is replaced by $1/n$.

one night and cut off all but a finite length, replacing the infinite “tail” with a resistor of value Z_0 , no one will notice a difference—there is no reflection from the output end of the line, so the input end has no way of distinguishing this case from an unbroken infinite line.

The voltage and current at any point on the line are in phase for a forward wave and 180° out of phase for a reverse wave—this is obvious at DC but applies at AC as well—which is what we mean by saying that the line looks like a resistor. We can find the instantaneous forward and reverse signal voltages V_F and V_R by solving the 2×2 system:

$$V_F = \frac{V + Z_0 I}{2} \quad \text{and} \quad V_R = \frac{V - Z_0 I}{2}. \quad (14.11)$$

14.4.1 Mismatch and Reflections

If the cable is terminated in an impedance Z different from Z_0 , the reflection coefficient Γ and the normalized impedance $z = Z/Z_0$ are related by

$$\Gamma = \frac{z - 1}{z + 1}, \quad z = \frac{1 + \Gamma}{1 - \Gamma}. \quad (14.12)$$

This transformation maps the imaginary axis (pure reactances, opens, and shorts) onto the unit circle $|\Gamma| \equiv 1$, and passive impedances ($\text{Re}\{Z\} > 0$) inside.

For example, if the far end of the cable is short-circuited, the voltage there has to be 0; this boundary condition forces the reflected wave to have a phase of π , so that the sum of their voltages is 0. At an open-circuited end, the current is obviously 0, so the reflection phase is 0;[†] the voltage there is thus twice what it would be if it were terminated in Z_0 —just what we’d expect, since the termination and the line impedance form a voltage divider.

Seen from the other end of the line, a distance ℓ away, the reflection is delayed by the round trip through the cable, so its phase is different;[‡] this leads to all sorts of useful devices. Using (14.12), with an additional phase delay $2\theta = 4\pi\ell/v$, where v is the propagation velocity of the wave in the line, we can derive the impedance Z' seen at the other end of the cable:

$$Z' = Z_0 \left(\frac{z + j \tan \theta}{1 + jz \tan \theta} \right). \quad (14.13)$$

The one-way phase θ is called the *electrical length* of the cable. The impedance at any point on the line derives from successive round-trip reflections, and thus repeats itself every half wavelength, as in Figure 14.6. We know that a short length of open-circuited transmission line looks like a capacitor, so it is reasonable that the same line would look inductive when short-circuited, which is what (14.13) predicts. However, it may be less obvious that when $\theta = \pi/2$, a shorted line looks like an open circuit, and vice versa. This *quarter-wave section* turns out to be very useful in applications.

[†]The fields actually stick out into space a bit on the ends, which gives rise to a slight phase shift. The situation is somewhat similar to total internal reflection at a dielectric interface, except that here the coupling between the transmission-line mode and the free-space propagating modes is not exactly 0, so there is some radiation.

[‡]We’re doing electrical engineering here, so a positive frequency is $e^{j\omega t}$ and a phase delay of θ radians is $e^{-j\theta}$.

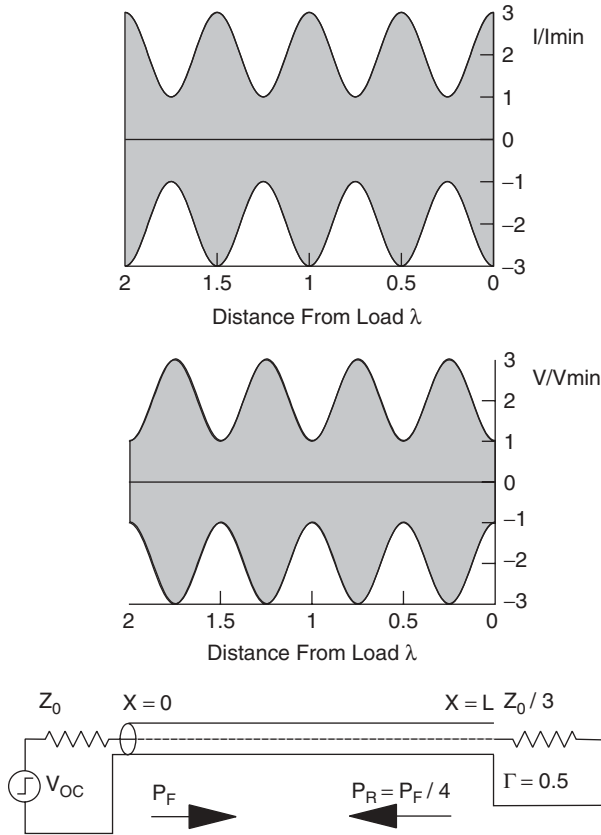


Figure 14.6. A transmission line can have forward and reverse waves propagating simultaneously.

The ratio of the incident to reflected wave powers is known as the return loss, $RL = 20 \log_{10}(1/|\Gamma|)$, and is an important parameter in RF circuit design. A well-matched load might have a return loss of 25 dB or so, and an amplifier is often only 10 dB. The other commonly heard parameter is the *voltage standing wave ratio* (VSWR, pronounced “vizwahr”), which is the ratio of the peak and valley amplitudes of the standing wave pattern produced by the reflection,

$$VSWR = \frac{1 + |\Gamma|}{1 - |\Gamma|}. \tag{14.14}$$

See Table 14.2 for VSWR/RL conversion values.

14.4.2 Quarter-Wave Series Sections

A quarter-wave section is just a chunk of line in series with the load, whose electrical length is $\pi/2$. It’s the electronic equivalent of a single-layer AR coating. Taking the limit

TABLE 14.2. Conversion Table from VSWR to Return Loss

VSWR	RL (dB)	VSWR	RL (dB)	VSWR	RL (dB)
1.01:1	46.1	1.25:1	19.1	5.00:1	3.5
1.05:1	32.3	1.50:1	14.0	3.00:1	6.0
1.10:1	26.4	2.00:1	9.5	10.0:1	1.7

of (14.13) as $\theta \rightarrow \pi/2$, we get

$$Z' = Z_0 \frac{1}{z} = \frac{Z_0^2}{Z}, \quad (14.15)$$

so that a short turns into an open, an open into a short, and reactances change in sign as well as in magnitude. This means that a capacitive load can be matched by using a quarter-wave section to change it into an inductive one, then putting in a series or parallel capacitance to resonate it out. If you have a resistive load R_L , you can match it to a resistive source R_S with a quarter-wave section of $Z_0 = (R_S R_L)^{1/2}$, just like the ideal $\lambda/4$ AR coating.

The AR coating analogy holds for more complicated networks as well; as in a Fabry–Perot interferometer, the discontinuity producing the compensating reflection to cancel the unwanted one can be placed elsewhere, providing that the electrical length is correct modulo π . Also as in a Fabry–Perot, the transmission bandwidth at a given order shrinks as the electrical length between the two reflections increases.

14.4.3 Coaxial Cable

Coax is familiar to almost everybody—a central wire with thick insulation, and an outer (coaxial) shield of braid, foil, or solid metal. You ground the shield and put your signal on the center conductor. Due to the dielectric, the propagation velocity is $c/\epsilon^{1/2}$ (just like light in glass). Current flows in opposite directions in the shield and center conductor, so there is no net magnetic field outside. The current in fact flows along the inside surface of the shield at high frequencies due to skin effect.

It's obvious that since the shield covers the inner conductor completely, there's no field outside, and so no crosstalk (i.e., unintended coupling of signals) can occur. Except that it does. This is another of those cases where failing to build what you designed produces gotchas. Most coax has a braided shield, which doesn't completely cover the inner conductor, and whose integrity depends on good contact occurring between wires merely laid on top of each other. You wouldn't trust that in a circuit, and neither should you here: single-shielded braided coax leaks like a sieve. RG-58A/U is especially bad—it talks to everybody in sight. If you're using coax cables bundled together, use braid-and-foil or double-shielded coax such as RG-223 for sensitive or high powered lines. For use inside an instrument, you can get small diameter coax (e.g., Belden 1617A) whose braid has been completely soaked with tin, forming an excellent electrical shield that is still reasonably flexible and highly solderable, and has matching SMA connectors available. (For some reason it has become very expensive recently—a few dollars per foot—but it's great stuff.) For microwaves, where the skin depth is small and therefore base-metal shields are lossy, there is semirigid coax (usually called hardline), made with a solid copper tube for an outer conductor. Special connectors are required for these lines,

although you can solder them directly to the board as well. Solid-shield coax also has much better phase stability than braided—at 50 MHz and above, jiggling a cable can cause a few degrees' phase shift, which is obnoxious in phase-sensitive systems. Even in nominally amplitude-only setups, the resulting change in the relative phases of cable reflections can cause instability reminiscent of etalon fringes.

If a cable has multiple ground connections (e.g., a patch cord strung between two connectors on the same chassis), the return current will divide itself between them according to their admittances, which will reduce the shield's effectiveness at low frequency and cause ground loops (see Section 16.5.2).

Even at low frequency, coax has its problems. Flexing it causes triboelectric noise, where charge moves from shield to insulation and back, like rubbing a balloon on your hair. And flexing causes the cable capacitance to change slightly, which turns any DC bias into noise currents as the cable capacitance changes.

14.4.4 Balanced Lines

One of the primary functions of a transmission line is to keep the signal fields confined to the interior of the line, or at least to its immediate neighborhood. The mathematical way of saying this is that the overlap integral between the transmission line mode and any propagating wave outside the line has to be 0. This can be achieved by 100% shielding, as in waveguide or coax, or more elegantly by using balanced line such as TV twin-lead or twisted pair. These balanced lines work by making sure that a current i flowing in one conductor is balanced by a current of $-i$ flowing in the other. By symmetry, this guarantees that the AC voltages are v and $-v$ as well. An arm-waving way of looking at this is to consider the asymptotic falloff of the signal. A finite length of wire with an AC current flowing in it is an antenna: it produces a far-field amplitude that falls off asymptotically as r^{-1} , so that its energy flux across any large sphere is independent of radius. Two wires with equal and opposite currents will produce a far-field amplitude proportional to $r_1^{-1} - r_2^{-1}$. If the wires are separated by a distance $2d$ in the X direction, we get

$$E \propto \frac{1}{\sqrt{(x-d)^2 + y^2 + z^2}} - \frac{1}{\sqrt{(x+d)^2 + y^2 + z^2}} = \frac{2d \cos \theta}{r^2} + O(r^{-3}), \quad (14.16)$$

where $x = r \cos \theta$. This falls off faster than a propagating wave and thus can have no propagating component whatever, since you can make the energy flux as small as you like by choosing a big enough sphere. This is relevant to the coupling of adjacent lines, as we'll see in Section 16.3.3. A crucial and usually overlooked precondition for this to work is that the currents must really sum to 0. Failing to obey this rule is the usual reason for failure with balanced lines. If the two currents do not sum to 0, you can decompose them into balanced and common-mode parts:

$$i_{\text{bal}} = \frac{i_1 - i_2}{2} \quad \text{and} \quad i_{\text{CM}} = \frac{i_1 + i_2}{2}. \quad (14.17)$$

The balanced part works as before, but the common-mode part flows in the same direction in both wires; the two terms in (14.16) add instead of subtracting, so from a far-field point of view, it's just like having only one wire—that is, an antenna. (In fact, folded

dipole antennas are commonly made from just this sort of balanced line—see the ARRL handbook.)

Since an appreciable field exists outside a balanced line, it is much more vulnerable to unintended coupling from very nearby objects than coax is. It doesn't like being near anything metal, in particular.

14.4.5 Twisted Pair

One particularly common balanced line is twisted pair. It is usually made in the lab by putting one end of a pair of hookup wires in the chuck of a hand drill, and spinning it while pulling gently on the other end (reverse the drill for a moment to avoid kinking when you release the tension). The twists serve to keep the pair at constant spacing and balance capacitive pickup; they also minimize the net loop area available for low frequency magnetic pickup, because the differential voltage induced in adjacent half-turns cancels. Twisted pair really works for that, provided it's truly used balanced. Too many people use it as a fetish instead and are surprised when it doesn't provide signal isolation.

Do not ground either side of a twisted pair. To reiterate: if you ground one conductor, you have built an antenna, not a transmission line. Use coax for unbalanced lines and twisted pair for balanced ones.[†]

14.4.6 Microstrip

An ordinary PC board trace of width w , with dielectric of thickness d between it and the ground plane, is an unbalanced transmission line called microstrip. You can look at it as a balanced line split in half, with the other side replaced by an image in the ground plane. Its characteristic impedance depends on the ratio of the width of the line to the thickness of the dielectric (the w/d ratio) and the dielectric constant of the material. You can make $50\ \Omega$ lines on ordinary G-10 or FR-4 fiberglass epoxy board ($\epsilon \approx 4.5$) by making $w/d = 2$. More detailed formulas are given in Section 16.2.2.

14.4.7 Termination Strategies

To avoid all reflections, both ends of the line must be terminated in Z_0 . The problem with this is that we only get half the open-circuit voltage at the other end, and we use a lot of power heating resistors. On the other hand, if only one end is properly terminated, any reflections from the other will go away after only one bounce, and so no resonances will occur. Accordingly, we can *series-terminate* coax, where only the driving end is at Z_0 . If we send a step function of open-circuit voltage V_{oc} down such a line, it initially sees a matched load of Z_0 , so the voltage is $V_{oc}/2$. When the pulse reaches the (open-circuited) far end, it is reflected with $\Gamma = 1$, so the voltage rises to V_{oc} , and the step function starts back toward the source. (Since the forward and reflected wave arrive simultaneously at the open end, there is a single smooth step there.) When the reflected wave gets back to the source, it is not reflected, and the whole line is at a constant V_{oc} . Thus a series-terminated line can transmit the entire open-circuit voltage without annoyance from the reflection, provided only that the driver can handle the reflected signal without difficulty. (See Figure 14.7.)

[†]Oh, all right. The magnetic pickup reduction still applies if you ground one side. Just don't ground it on both ends, and don't say you weren't warned.




Figure 14.7. A series-terminated (or back-terminated) line. Although points along the line see the effects of the reflection, the open-circuited end sees a nice clean step, as though it were connected right to the source.

Aside: What Is the Worst Case? In testing a transmission system, we often want to verify that it is stable and works acceptably with a “worst case termination.” However, what the worst case is may not be obvious. A line driver may work well into a mismatched 10 foot cable, and into a mismatched 1000 foot cable, but oscillate madly with a 250 foot cable. The loss in the very long cable makes it start to look like an infinite (i.e., resistive) line, whereas that 250 foot cable has both a huge delay and a strong reflection.

14.5 TRANSMISSION LINE DEVICES

14.5.1 Attenuators

Because a matched line looks like a resistor, we can use simple resistor networks to attenuate signals. They are three-element ladder networks (pi or T), so that we can set the attenuation, R_{in} , and R_{out} independently. Attenuators are usually called *pads*, because they reduce the effect of outside influences such as mismatches; if the source is isolated from the outside with an N dB pad, any reflected power has to cross it twice before reaching the source again; the return loss seen by the source thus goes up by $2N$ dB (provided the pad itself is that well matched—don’t expect a 60 dB pad to have a 120 dB return loss). This is often useful to stop an amplifier from oscillating when badly mismatched. Don’t use this idea on your low level circuitry, though, or the SNR will suffer; an N dB pad in front of an amplifier with a 3 dB noise figure will present an $N + 3$ dB noise figure to the world. (See Figure 14.8.)

14.5.2 Shunt Stubs

Since shorted pieces of coax can be made to look like capacitors and inductors, we can use them in matching networks. In particular, a shunt stub is very useful in resonating out




Figure 14.8. Three resistors give enough degrees of freedom to match specified input and output impedances as well as arbitrary attenuation values. This is a 10 dB tee-network attenuator for 50 Ω.




Figure 14.9. A shunt stub can tune out any reactance at a single frequency.

reactive loads, as shown in Figure 14.9. It's very convenient, too: put a coax patch cord on a tee connector, and stick a thumbtack through the shield into the center conductor, trying different places along the cable until your reactance goes away. If you'll be needing the setup for a while, cut the coax and solder the short. Thumbtack shorts are surprisingly stable: for temporary use, just leave the tack in there.

14.5.3 Trombone Lines

If you don't like thumbtacks, or need an adjustable series section, the line length can be adjusted with a telescoping coaxial section called a trombone line. New ones work best, as the contact gets flaky with age.

14.5.4 Transmission Line Transformers and Chokes

If you put a signal into a transmission line, it should emerge from the other end. The ends of the line have two terminals each, like resistors, and so can be wired in series or parallel with other circuit elements (e.g., termination resistors). The idea of a transmission line transformer is to wire two or more lines in parallel at one end and in series at the other. As shown in Figure 14.10, connecting two lines this way makes an input impedance of $Z_0/2$ and an output impedance of $2Z_0$, a 4:1 impedance transformation ratio. Series-parallel combinations can in principle achieve any rational impedance transformation value.

This sounds like getting something for nothing, because after all these supposedly independent Z_0 resistors are in fact connected together with wires, which will try to short out our nice idea; if you look carefully at it, our 4:1 transformer is a dead short at DC. On the other hand, if the sections are $\lambda/4$ long, the dead short at one end will look like an open at the other. (Should they be $\lambda/4$ in air or inside the line? Why?)




Figure 14.10. A 4:1 coaxial transformer.

Another thing to notice is that the current which acts to short out the transformer flows only on the shield. Its corresponding \mathbf{E} and \mathbf{B} fields therefore do not cancel in the region outside the line. We can prevent it from flowing by using a common-mode choke, which is nothing more than wrapping a few turns of line around a toroid core (both coax and twisted pair work well). Such a choke will kill the short circuit and allow us to build wideband transmission line transformers on this parallel-in, series-out principle. The differential mode signal produces no magnetization in the core, and so feels no inductance. The absence of core magnetization is their defining property, and because there are no eddy currents, hysteresis, or other magnetic losses to worry about, transmission line transformers are extremely stable, wideband, and low in loss.

14.5.5 Directional Couplers

As we saw in Section 8.3.3, if we couple two long lines together (e.g., by putting them side by side with slightly leaky shielding between them), a forward wave in one will couple into a forward wave in the other. Providing the interaction region is many wavelengths, it discriminates well between forward and reflected waves, because only the forward wave is phase matched. Such directional couplers are the easiest way to measure the forward and reflected power in your circuit. At frequencies where such long lines are inconvenient, we can take advantage of the opposite phase relationships of voltage and current for forward and reflected waves, making a balanced transformer device that does the same job.

14.5.6 Splitters and Tees

A tee connector wires two cables in parallel, creating an impedance mismatch and consequent reflection. A splitter is a transformer-based device that eliminates the problem and provides isolation between the two taps: as long as the input is matched to $50\ \Omega$, reflections at one output don't affect the other.

14.6 DIODES AND TRANSISTORS

14.6.1 Diode Switches

We all know that a diode conducts electricity in one direction and not in the other. Its conductivity is a smooth function of applied voltage, going from 0 to a large value as the bias voltage goes from large reverse bias to about 1 V forward (anode positive). The forward current of a diode is approximately predicted by the *diode equation*,

$$I_F \approx I_S(e^{V_F/V_\gamma} - 1). \quad (14.18)$$

It is sometimes maintained that V_γ is the thermal voltage $V_{th} = kT/e$, but that isn't so for real diodes; at 300 K, V_{th} is 25.7 mV but V_γ is 30–50 mV for most devices. Diode-connected transistors (base shorted to collector) are an exception, so if you need accuracy, use those instead of real diodes. In any case, the conductance of a diode is approximately I_F/V_γ , so with a few milliamps of forward current you can get impedances of 10 Ω or so—not that great compared with modern FETs, but with nice low capacitances.

Although (14.18) is wrong for $V_F \ll V_\gamma$ due to second-order effects, it is accurate enough around zero bias that we can calculate the zero-bias shunt impedance r_0 by differentiating,

$$r_0 = V_\gamma / I_S \quad (14.19)$$

which in general is far from infinite—and since I_S increases exponentially with temperature, it can get down to the 10 k Ω range at 125 °C—beware of your protection diodes in high impedance circuits (see Section 14.6.3).

Diodes are ideal current switches: they pass currents in the forward direction but not in the reverse direction. A reverse-biased diode looks like a small capacitor. To take advantage of this ideal behavior, design circuits that are not sensitive to the exact voltages involved (e.g., AC-coupled, current mode, or translinear circuits).

You can get cheap diodes (PIN and Schottky) that have very low capacitance when reverse biased, and which conduct hard in forward bias; these are good for switching AC signals.[†] Unlike electromechanical or FET switches, their control currents are not separate from the signal, but if you're using DC to switch an AC signal, you can separate them again with a capacitor, as in Figure 14.11. Don't worry about the DC voltage offset unless you're at DC or the switching waveform has components that land in your passband.

Another way of getting rid of the offset is with balanced circuits, for example, the series diode clipper and diode bridge of Figure 14.12. In understanding these circuits, think in terms of current, and of the diodes as ideal current switches; matched diodes




Figure 14.11. If you can separate the control current from the signal, diodes make excellent switches. Here two diodes make a gain switch for a current feedback amp.

[†]PIN diodes store lots of charge in their junctions, which has to be swept out before the diode turns off; this means that their conductance is modulated less on each half-cycle, and so they need less bias current for a given IP_2 .

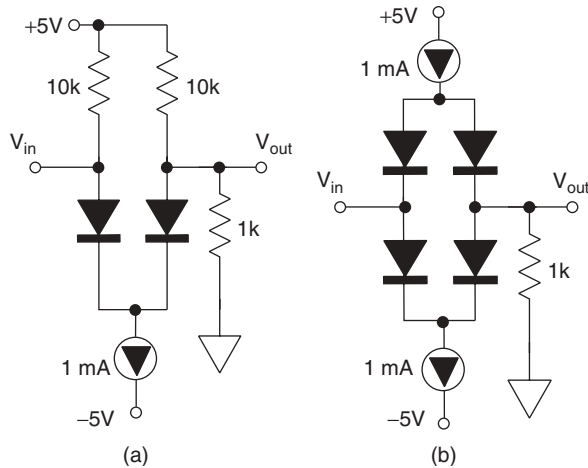


Figure 14.12. Another way of separating the switching current from the signal is to use balanced circuits: (a) series clipper and (b) diode bridge.

work best and are vital in the bridge configuration. This four-diode circuit is often used as an ultrafast sampling bridge; the diodes are biased off normally, but a fast pulse applied in place of the DC bias turns the switch on momentarily, charging up a capacitor on the output; this is how most sampling scopes work.[†]

For switching, Schottky diodes are best if your switch waveform may vary, as in diode mixers; PN diodes store charge, and so change speed with drive level, whereas Schottkys don't.

Aside: Diode Foibles. It's worth reemphasizing that odd diode pathology, their surprisingly low impedances when used near zero bias (e.g., as the protection diodes of a sensitive FET amplifier). We ordinarily think of a diode with zero bias as an open circuit, but in fact for small signals, it's a resistor of value $R_{d0} = \partial I_F / \partial V_F \approx V_F / I_0$; with gold-doped diodes such as 1N914s and 1N4148s, this can be as low as a few kilohms at elevated temperature, which starts to matter even with microvolt signals. Base–collector junctions of common transistors such as 2N3904s (or MMBT3904s) are much less leaky. For really demanding applications, wide bandgap diodes such as LEDs are enormously better for this use, because their much higher V_F for the same I_F means that I_0 is many orders of magnitude smaller. (Would you believe < 100 fA for biases between -5 V and $+0.5$ V?) Section 18.7.2 uses this nice property in a pyroelectric detector front end.

Diodes exhibit delays in turning on as well as off, and these delays vary widely among devices. A 1N4001 rectifier needs hundreds of microseconds, 1N914s a few nanoseconds, and 1N5711 Schottkys less than 1 ns. It is odd to see a diode circuit overshoot by a volt on turn-on, but some devices will do this. If you find one, complain to the manufacturer or change device types.

[†]Because of their very low duty cycles, sampling scopes precharge the capacitor to the previous value, using a feedback loop called the *sampling loop*, which greatly improves their accuracy.

14.6.2 Bipolar Transistors

The most flexible active device in electronics is the bipolar junction transistor (BJT). It is basically a three-terminal, voltage-programmable current source; the current between the collector and emitter terminals is controlled by the base–emitter voltage, V_{BE} , and is nearly constant over a very wide range of bias conditions. It has excellent characteristics and is highly repeatable from unit to unit, provided you use it properly. We can't devote nearly enough space to this remarkable device, so by all means look it up in Gray and Meyer. Instrument designers are always needing the occasional perfect circuit, but one that isn't available in the IC catalogs; thus we'll look at the BJT's virtues for signal processing.

The oldest, simplest, and still most useful mathematical description of a BJT is the Ebers–Moll model, which (in a somewhat simplified form, valid for normal bias[†] only) predicts that the collector current I_C is a simple function of V_{BE} :

$$I_C = I_S \exp \frac{eV_{BE}}{kT}, \quad (14.20)$$

where e is the electron charge, k is Boltzmann's constant, and T is absolute temperature. When computing the gain, we care more about the transconductance

$$g_M = \frac{\partial I_C}{\partial V_{BE}} = \frac{e}{kT} I_C, \quad (14.21)$$

which is $I_C/25$ mV at room temperature, a very high value.[‡] The exponential character is accurate to better than 0.1% over at least four decades of collector current in a good device, and the exponential constant is $V_T = kT/e$ to very high accuracy; what does vary unit-to-unit is I_S (and of course T , if we're not careful—another opportunity for common centroid design).

The collector and emitter currents are not exactly equal, because the base draws some small current. The current gain $\beta = I_C/I_B$ and is usually 30–500 in common devices, but (unlike g_M) varies widely between devices of the same type. It is thus a bad idea to design circuits that depend on β having some particular value. One other figure of merit for amplifier transistors is β linearity, that is, how much β changes with I_C —it may be flat to 5% over five orders of magnitude, or vary 2:1 over one order. We'll be very concerned with that in Chapter 18 when we talk about laser noise cancelers.

14.6.3 Temperature Dependence of I_S and V_{BE}

The Ebers–Moll model predicts that $V_{BE(on)}$ is

$$V_{BE(on)} = \frac{kT}{e} \ln \frac{I_C}{I_S}, \quad (14.22)$$

[†]Normal bias means reverse-biasing the CB junction and forward-biasing the BE junction.

[‡]No semiconductor device can have a higher transconductance than this, because it is limited by the thermal spread of the Fermi level in the semiconductor.

which looks proportional to T , but actually has a nearly constant -2 mV temperature coefficient from nitrogen temperature to 150°C . The reason for this is that I_S is a strongly increasing function of T , roughly proportional to[†]

$$I_S \propto T^{3.25} \exp\left(\frac{-eV_{G0}}{kT}\right), \quad (14.23)$$

where V_{G0} is the zero-temperature bandgap of the semiconductor (1.205 V for silicon), and the constant 3.25 is somewhat device dependent. Above nitrogen temperature, this produces a TC of -1.8 to -2.2 mV/ $^\circ\text{C}$ in V_{BEon} ,[‡] and since the zero-bias $r_0 = V_T/I_S$ decreases rapidly with T , we have to start worrying about shunt resistances at high temperatures. This is a problem mostly with diodes and diode-connected transistors. Note especially the extremely strong dependence of I_S on V_{G0} : high bandgap devices such as LEDs and SiC photodiodes have very high shunt resistances, whereas low V_{G0} IR diodes like InAs and InSb have very low r_0 and need cooling.

14.6.4 Speed

The β of a BJT tends to roll off as $1/f$ for large f (and therefore has a phase near -90°). The frequency at which $|\beta| = 1$ is called the *cutoff frequency*, f_T . As a first approximation, β goes as f_T/f down to where it flattens out at its DC value. The frequency f_T is a fairly strong function of collector current, going as I_C for low currents, up to a broad peak somewhere around a third of its $I_{C\text{max}}$ spec limit. It also improves with increasing bias voltages, which reduce the interelectrode capacitances.

14.6.5 Biasing and Current Sources

That -2.1 mV/ $^\circ\text{C}$ temperature coefficient is a nuisance sometimes. Since I_C goes up by a factor of e in 26 mV at 300 K, I_C goes up by 9%/ $^\circ\text{C}$ if V_{BE} is fixed (the increase in V_T reduces it by only 0.3%). If the increased power dissipation from that 9% extra current raises the temperature another 1°C , the bias will be unstable, and the resulting *thermal runaway* will melt the device.

BJTs also exhibit Early effect, where I_C increases with collector–emitter voltage V_{CE} , and furthermore, different devices have different I_S , which makes accurate current sources hard to build. In addition, the collector current predicted by (14.20) exhibits full shot noise, which has to be dealt with in low noise design.

These problems are easily fixed by negative feedback, as shown in Figure 14.13. Feedback can be applied from collector to base with the emitter grounded, but a more common way to do it is by putting a resistor in series with the emitter. This technique is known as *emitter degeneration*.[§] Without getting mathematical, if the emitter resistor drops 2 V, I_C drifts only 0.1%/ $^\circ\text{C}$ instead of 9%/ $^\circ\text{C}$. Problem 14.8 (at <http://electrooptical.net/www/beos2e/problems2.pdf>) has more on this. Similarly, the shot noise current produces an opposing collector–emitter voltage that reduces the shot noise current by the same factor.

[†]Paul A. Gray and Robert G. Meyer, *Analysis and Design of Analog Integrated Circuits*, 2nd ed. Wiley, Hoboken, NJ, 1984, p. 340.

[‡]The TC increases as collector current goes down, reaching as much as -3 mV/ $^\circ\text{C}$ at picoamp collector currents.

[§]Positive feedback is what keeps oscillators running, and so is said to be regenerative; negative feedback is then logically degenerative. No moral judgment is implied.

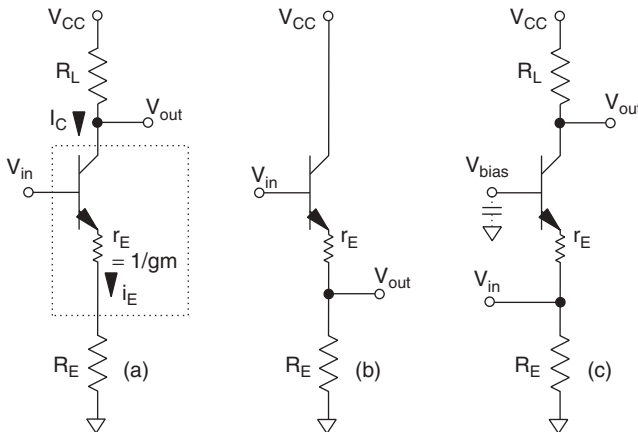


Figure 14.13. Transistor amplifier configurations: (a) common emitter, high A_V , medium Z_{in} ; (b) common collector or emitter follower, high Z_{in} , low Z_{out} , $A_V = 1$; and (c) common base, low Z_{in} , high A_V . Emitter degeneration due to R_E stabilizes the bias and suppresses nonlinearity in each case.

14.6.6 Cutoff and Saturation

We often run transistors in nonlinear modes: cutoff, where the collector is biased correctly but V_{BE} is small or negative, so no collector current flows; or saturation, where $I_C R_L$ is so big that the CB junction becomes forward biased (about 200 mV of forward CB bias is the real limit of normal bias conditions). Transistors can go in and out of cutoff very fast, but come out of saturation very, very slowly, so avoid saturating if you need speed. The origin of this effect is that pulling V_{CE} below V_{BE} reduces the E field in the base region so much that the carriers coming in from the emitter don't fly into the collector as normal, but hang around in the base until they recombine, which shows up as a huge increase in the base current. All the stored charge has to be swept out before the transistor can turn off again, which is why it's slow.

If you drive the base hard enough ($I_B \approx 0.1I_C$), you can get V_{CEsat} down to 50–100 mV in small signal devices, and around 0.4 V in power transistors. The CE bias necessary to avoid saturation is roughly proportional to absolute temperature (+0.33%/°C at room temperature), so check your circuits at high and low temperatures to make sure they still work properly.

Aside: Inverted Transistors. There is a sometimes-useful trick to get it lower: run the transistor upside down. That is, exchange the collector and emitter. The emitter is normally much more highly doped than the collector; this asymmetry makes the β of an inverted transistor very low (2–5, typically), but the V_{CEsat} can go as low as 10 mV.

14.6.7 Amplifier Configurations

Transistors are flexible devices; a BJT amplifier can use its base or emitter as input and its collector or emitter as output; since you need two terminals per port and there are only three available, one has to be common to both input and output; it can be any one of the three. (Being the control electrode, the base always has to connect to the input port.)

Figure 14.13a shows the collector as output, the common-emitter amplifier. Neglecting the Early effect, it has medium input impedance of $R_{in} = (1 + \beta)(R_E + r_E)$, high output impedance of $R_{out} = R_L$, and medium voltage gain, $A_V = R_L / (R_E + r_E)$. The BJT's shot noise contribution is reduced by a factor of $r_E / (R_E + r_E)$.

If the emitter is used as the output in Figure 14.13b (the common-collector or *emitter follower* connection), the same negative feedback that stabilized the bias stabilizes the output voltage, producing a very low output impedance ($R_{out} = R_L \parallel r_E$), higher input impedance (still $(1 + \beta)(R_E + r_E)$, but R_E is usually several times bigger). It has nearly unity voltage gain ($A_V = 1 - r_E / R_E$) since the emitter follows the base closely and (for the same reason) has a shot noise voltage contribution of $A_V i_{Nshot} R_E$. BJTs follow their noise models very closely (see Section 18.4.5).

Finally, as in Figure 14.13c, if we keep the base still (common-base), any signal current i_E we shove into the emitter will reappear at the collector; the collector swing is then $R_L i_E$, whereas the emitter moves only $(r_E \parallel R_E) i_E$. For frequencies where the source impedance remains high, there is no additional shot noise contribution from the BJT. The low noise and near-total lack of voltage swing at the emitter makes the common-base amplifier very useful in reducing the effects of capacitance, as in *cascode* amplifiers and photodiode front ends (see Section 18.4.4 for a great deal more on these points.)

Example 14.1: Capacitance Multiplier. Figure 14.14 shows what we've already adumbrated: a slow *RC* lowpass filter with an emitter follower makes a really good supply filter for low level circuitry. The input impedance of the follower is about $(\beta + 1)R_L$, which means that for a given sized capacitor, the *RC* corner frequency can be $\beta + 1$ times lower, and the filtering correspondingly better. Considering that an MPSA-14 Darlington has $\beta \approx 20,000$, this is a big deal. This improvement can be traded off for reducing the drop across the resistor, so a β of 300 can be apportioned as 30 times lower f_c and 10 times lower *IR* drop. If you need real help, try splitting the input resistor in two, and putting another capacitor to ground from the middle, as in Figure 14.14c. You can reach 100 dB suppression at low kilohertz frequencies with a couple of 22 k Ω resistors and 1

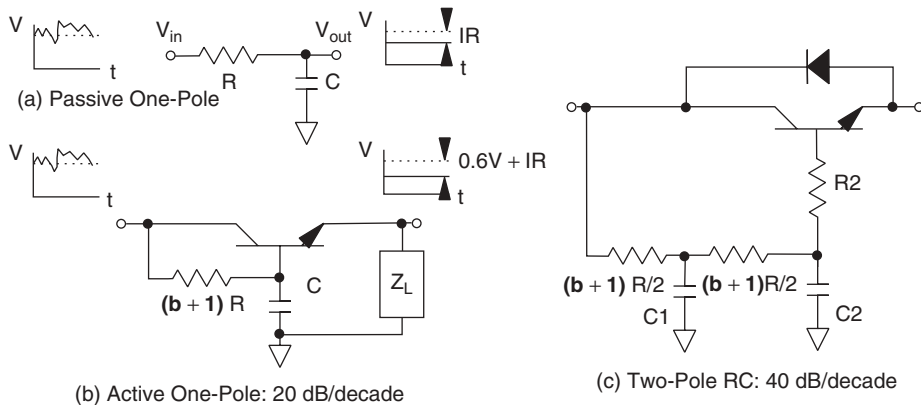


Figure 14.14. Capacitance multipliers. (a) *RC* lowpass, whose rejection is limited by the allowable voltage drop. (b) one-pole capacitance multiplier: the rejection improves $\beta + 1$ times. (c) Realistic two-pole version: the diode and R_2 are for protection from input shorts, which otherwise will destroy the transistor (make sure the bypasses on the circuits you're driving are at least 10^6 times bigger than the diode capacitance).

μF capacitors. For the highest suppression, you can replace the single diode with two diodes with a big bypass capacitor in between.

Somewhat similar tricks can be played with currents. For instance, by taking the circuit of Figure 14.14b and connecting the capacitor between the base and emitter of the transistor, you can make a simulated inductor.

14.6.8 Differential Pairs

One of the most important elementary circuits is the differential pair, formed by connecting the emitters of two (notionally identical) BJTs, as shown in Figure 14.15. The Ebers–Moll model predicts that collector currents I_{C1} and I_{C2} obey

$$\frac{I_{C2}}{I_{C1}} = \exp\left(\frac{e\Delta V_{BE}}{kT}\right), \quad (14.24)$$

which specifies the ratio of I_{C1} and I_{C2} but doesn't depend on how big they are: an ideal bipolar diff pair is a perfect voltage-controlled current splitter. (The TC of I_S cancels out.) This comes in very handy, especially in noise canceling front ends (see Section 18.6.3). Since ΔV_{BE} need not be large (± 60 mV gets you a 90:10 current ratio at 300 K), the currents can be switched without large voltage swings. If the input swing is a bit bigger than this (e.g., 500 mV p-p), the diff pair is an excellent current switch; this is how emitter-coupled logic (ECL) works.

When building differential amps, we usually care more about the difference between I_{C1} and I_{C2} , since that's what the output voltage depends on; it is

$$I_{C2} - I_{C1} = I_{\text{Bias}} \tanh\left(\frac{e\Delta V_{BE}}{2kT}\right). \quad (14.25)$$

Usually you bias a diff pair with a current source, because r_E of each transistor serves as the other one's emitter resistor, and using a current source keeps the magnitudes of I_{C1} and I_{C2} independent of the common-mode base voltage $V_{CM} = (V_{B1} + V_{B2})/2$. You can also use a large resistor if you don't need good common-mode rejection away from $\Delta V_{BE} = 0$.




Figure 14.15. A BJT differential pair makes a near-ideal current splitter, since I_{C2}/I_{C1} depends only on ΔV_{BE} .

14.6.9 Current-Mode Circuitry

The diff pair is the simplest example of a *translinear circuit*. Translinear circuits such as the Gilbert cell multiplier of Figure 14.16 use bipolar transistors for current steering, which is what they're best at. The cross-connected differential pairs at the top of the Gilbert cell commute the output currents of Q_1 and Q_2 , making a nearly ideal multiplier.

Like the differential pair, translinear circuits switch at very low input voltages, that is, at a fundamentally minimal circuit impedance for a given power dissipation. This greatly reduces the effect of device capacitances, making them very fast. They're not that easy to debug, but they work very well and are especially suitable for low supply voltages. A typical example is an operational transconductance amp with linearizing diodes on its inputs (e.g., the LM13700).

You have to watch out for shot noise in these circuits, because the transistor junctions generate it, and it takes feedback to suppress it. In the limit of infinite β , if the emitter current of a differential pair has exactly full shot noise, each collector current also has exactly full shot noise, independent of the splitting ratio. Thus if you don't keep copying the current with undegenerated current mirrors,[†] there's no serious buildup of shot noise in the signal path.




Figure 14.16. Gilbert cell multiplier.

[†]A current mirror uses a diode-connected BJT to provide V_{BE} for another identical device; assuming $\beta \gg 1$, the second BJT's collector current will be a replica of the first one's, with additional full shot noise.

Aside: FETs. It may seem strange that FETs haven't been given equal time with BJTs, especially with the number of FET integrated circuits around. The reason is that apart from isolated instances, e.g., using power FETs for switching, dual-gate GaAs FETs for isolation amplifiers, or MOSFETs for zero input current buffers, small-signal discrete FETs have little to recommend them for circuit design. Besides low transconductance, FETs have disgracefully poorly specified biasing properties; a typical discrete N -channel JFET has a threshold voltage $V_{gs(\text{Th})}$ specified as -2 to -5 V (a typical BJT might have a range of 50 mV). That means that your source follower circuit has an offset that may vary by 3 V from device to device, and it only gets worse if you want voltage gain. Not many circuits can conveniently accommodate that sort of sloppiness, especially on a single 5 V supply. Dual matched JFETs are a partial exception, but good luck finding them these days.

There do exist some nice quiet FETs that are good for some things. For audio frequencies, the Toshiba 2SK368 is very good, with 1 Hz v_N below 0.8 nV/Hz^{1/2}, but it has input capacitance $C_{gs} = 80$ pF and feedback capacitance $C_{dg} = 15$ pF, which is fairly horrible. At RF, the BF862 makes good bootstraps and discrete input stages for TIAs. The author has designed with them several times, but they never seem to make the final cut.

Heterojunction FETs (HJFETs) such as the NE3508 series have much better DC specs than JFETs and are about 100 times faster than a BF862, even—ten times the g_m and a tenth the C_{dg} . At 2 GHz, the NE3509 can reach noise temperatures of 30 K at room temperature in a properly tuned amplifier.

14.7 SIGNAL PROCESSING COMPONENTS

We'll talk about the ins and outs of ICs and modules shortly, but first of all it's worth making a high level tour of how to choose components wisely.

14.7.1 Choosing Components

Make sure you design your circuit around widely available components, which are likely to be obtainable and not be discontinued for the life of your product. The author has had his favorite unique components discontinued a number of times, especially the TL011 series of three-terminal current mirrors, which were hard to replace, the MRF966 dual-gate GaAs FET, the MRF9331 micropower RF transistor, the AD639 sine converter, the MC13155 FM IF chip. . . . Especially dubious are devices intended for specific fast-changing markets, for instance, read channel amplifiers for hard discs or DVD player motor control ICs. The likelihood of that chip being available in five years is next to zilch.

14.7.2 Read the Data Sheet Carefully

There's a lot of data in a data sheet. Some absolutely crucial things are usually specified in small print in one corner of it; for example, the variation of phase margin versus feedback resistance in current feedback amplifiers, the tendency for multiplexers to connect all their inputs together if even one is overdriven, FET op amps' habit of pegging their outputs when their input common-mode range is exceeded, the maximum allowable voltage drop

between analog and digital ground on an ADC, that sort of thing. Read the data sheet carefully, and don't ignore the stuff you don't understand; look it up. Bob Pease of National Semiconductor once wrote an article on reading data sheets, which it's worth getting.[†]

14.7.3 Don't Trust Typical Specs

Try to design your circuits so that parts meeting only the minimum spec will work without difficulty. If you need something a bit more special, for example, the typical spec as a guaranteed minimum, you can do that one of two ways; first, get the manufacturer to test them for you, which is expensive, or order a couple of months' stock ahead of time, and test them yourself. The extra inventory is to cover you when you get a batch that is good, but doesn't meet the typical spec; keep a month's worth of known good parts in a safe somewhere for emergencies.

Above all, don't trust any spec in a data sheet marked "Preliminary" or "Advance Information." By way of example, the MC34085 quad JFET op amp had its typical GBW and slew rate numbers reduced by nearly a factor of 2 between the advance information and product release data sheets. This is embarrassing for the manufacturer but may be catastrophic for you.

14.7.4 SpecsmanSHIP

Don't expect that all the nice specifications of a given component will necessarily apply at the same time. It is reasonable for the flatness and accuracy of a switched-capacitor filter chip, for example, to degrade somewhat at the extremes of its specified frequency range. Dig through the data sheet to find out how it works under the actual conditions anticipated, and don't try to push more than one limit at a time. The classical example of this is analog lock-in amplifiers, which have accuracies of 0.1% in amplitude and 0.1° in phase, with a bandwidth of 1 Hz to 100 kHz. You have to read the manual carefully to find out that the near-perfect accuracy specs only apply near 1 kHz, and that the error at the band edges is more like 2 dB and 1 radian (most digital ones are better).

14.7.5 Watch for Gotchas

The second most important reason to learn about discrete circuit design is to be able to understand IC data sheets, especially the "simplified schematics" that farsighted manufacturers like National print in their data sheets (at least for bipolar parts). You can spot a lot of circuit funnies that way, for example, how the output swing will vary with the supply voltage, and whether you can do anything about it. In data sheets especially, if you notice something strange, stop and figure it out before using that part—treat it as though it were a large machine with unfamiliar controls.

[†]Bob Pease, How to get the right information from a datasheet, Appendix F of *Operational Amplifiers Data Book*. National Semiconductor Inc., 1988–1993. Reprinted in *Troubleshooting Analog Circuits*, Butterworth-Heinemann, Woburn, MA, 1991.

14.7.6 Mixers

The main application of diode switching is in diode ring double-balanced mixers (DBMs, not to be confused with dBm). This simple but remarkably effective combination of four matched Schottky diodes and two toroidal transformers is shown in Figure 14.17.

The diodes are arranged differently than in a bridge rectifier; both ports conduct on both half-cycles. The RF port transformer produces 0° and 180° phases of the weak RF input, and the strong LO signal switches the diodes to commute between these two phases. The output is the resulting current; since it flows 180° out of phase through both halves of both transformer secondaries, it suffers no inductance and no coupling to the RF or LO port—that's why the mixer is called double-balanced.[†]

How well a mixer does this separation trick is specified in its LO–RF and IF–RF isolation numbers. Typically, a good double-balanced mixer will have LO–RF isolation of 30–45 dB, and LO–IF isolation of 20–35 dB. This is of more than cosmetic importance. If we have unintended signal paths, such as the LO of one mixing stage going out its RF port and getting into the previous mixing stage, the signals will get intermodulated in the other stage and cause spurious signals.

A diode bridge mixer typically has a conversion loss of about 5 dB, so that the desired IF is 5 dB below the RF input level. Its noise figure is usually approximately the same as its conversion loss, which means that it contributes very little noise of its own. The upper end of its dynamic range is set by intermodulation; its input-referred $P_{1\text{dB}}$ is typically $P_{\text{LO}} - 7$ dB.

14.7.7 LO Effects

Mixers typically work well only over a narrow range of LO powers (e.g., the Mini Circuits SRA-1 likes +7 dBm, although it will work from +4 to +10 dBm). The third-order nonlinearity of diode bridges and FET mixers is dominated by the nonlinear conductance of the devices during turn on and turn off. Using square wave drive makes the active devices switch faster, and so improves the mixer's linearity. At frequencies for which good logic devices exist, use gates or inverters to drive the mixer. (Use analog power for that stage, and don't use the output of a clocked logic block such as a counter or flip-flop for anything analog.)




Figure 14.17. Diode bridge double-balanced mixer (DBM).

[†]Single-balanced mixers keep the LO signal out of the IF and RF, and double-balanced mixers also keep the RF out of the IF.

The symmetry of the mixing characteristic prevents LO amplitude noise from being converted into IF phase noise, and if the mixer is driven hard, such noise won't amplitude modulate the IF, either. On the other hand, decreasing the LO amplitude will cause the diodes to switch more slowly, since it takes more of a cycle to get enough LO current to saturate them. When the diodes are only partly turned on, the RF input can significantly change the diode conductance, causing nonlinearities and attendant signal distortion. This effect is symmetric, so it ideally causes no phase shift, but the increased nonlinearity can make LO noise cross-modulate the signal components, and will certainly increase the intermodulation products. This is not too obnoxious unless the LO is weak and the RF has more than one strong signal.

14.7.8 Op Amps

You've probably designed at least one op amp circuit already, and know the ideal op amp rules, virtual ground, and so forth, so we won't belabor them. (If you haven't, you can learn about them in Horowitz and Hill.) Decent op amps have noise voltages of 0.9 to 30 nV/Hz^{1/2}, a 30 dB range. Current noise ranges much more widely, from a few dozen electrons/s/Hz^{1/2} up to tens of picoamps (a 120 dB range). Very low noise voltages come at the price of large noise currents, large bias currents in bipolar amplifiers, and, most insidiously, high input capacitance. This input capacitance loads the summing junction (where feedback is applied) and can cause all sorts of mysterious instabilities. This severely limits the impedance levels you can use; in a front end amplifier, it hurts your SNR with small photocurrents. Chapter 18 has an extended practical example.

The other thing to mention is that the inverting input of an op amp is not really virtual ground, not at AC. If the output swings ± 5 V at frequency f , the inverting input will be moving by about $(5 \text{ V}) f/f_c$, where f_c is the closed-loop 3 dB bandwidth (the exact value depends on the frequency compensation, but it won't be less than this). If f is 1 MHz and f_c is 10 MHz, the "virtual ground" will be swinging by a good volt peak-to-peak. Whether an op amp can manage that depends on its input circuit—a BJT diff pair can only manage ± 60 mV. (This also shows the limits of "virtual ground" in providing isolation between signal sources connected to the summing junction.)

In our current enthusiasm for low supply voltages, op amps whose outputs swing from rail to rail are popular. Their output stages cannot be emitter followers, which lose at least 0.6 V on each side, and so they come from the collectors or drains of the output devices. As we saw earlier, these are high impedance points, and indeed a rail-to-rail op amp has an open-loop output impedance at least an order of magnitude higher than a good emitter-follower output. This makes them much more sensitive to weird loads and funky reactive feedback networks, so beware.

14.7.9 Differential Amps

Voltage is a relative measure, like altitude. All amplifiers measure a voltage difference between the input terminals, and turn it into another voltage between the output terminals. A single ended amplifier has a common reference point for both input and output, whereas a differential amplifier allows both input terminals to swing back and forth, and provides a separate terminal for the output reference voltage.

The figure of merit for a diff amp is its common-mode rejection ratio (CMR or CMRR), which is notionally its differential gain divided by its common-mode gain (inputs shorted). CMR is a subtle thing to measure; Pease gives a good discussion.

Differential pairs make good diff amps, since their outputs move symmetrically and ignore common-mode voltages, but their outputs are up near the supply rail. Another stage is normally required to convert their differential output to a ground-referred single-ended one.

Besides rejection of common-mode signals, a truly differential amplifier should not load down the measurement points, and certainly should not load them asymmetrically. Operational amplifiers (op amps) are good differential amps but have the drawback that their feedback networks are connected directly to their inverting inputs, resulting in low and unbalanced input impedances. An instrumentation amp has good differential performance (120 dB rejection of common-mode signals near DC), together with high input impedance. It works by buffering each input of an op amp with its own noninverting stage, with a clever circuit trick to greatly improve CMR at high gain (see Horowitz and Hill).

At RF, differential amps are often used to reduce coupling via grounds, and for their inherently better overload behavior due to their symmetry.

Gotcha: Floating Sources. Beginners are often frustrated by differential measurements. If you put a floating voltage source such as a transformer secondary or a thermocouple across the inputs of a high impedance differential amplifier such as an instrumentation amp or the inputs of a balanced A/D converter board, it won't work. This is because the amplifier inputs will always source or sink some bias current. In a single-ended measurement, this current flows through the source to ground, causing few problems unless the source resistance is extremely high. In a floating measurement, this current has nowhere to go, because the bias currents of the two inputs will never sum to zero. The result is that the floating source floats toward one of the supply rails until the net current going into it becomes zero, and of course the amplifier ceases to work properly. This leads to even more mysterious problems when there is significant hum on the source (e.g., a remote thermocouple input). The solution is simple, fortunately: put a big ($\approx 1\text{ M}\Omega$) resistor from one side to signal ground.

14.7.10 RF Amps

Radio frequency processing is usually done in the context of transmission lines, which have very low characteristic impedances (50–300 Ω , nearly always 50 or 75). RF amplifiers are designed to match into these lines reasonably well, so that they can be strung together with cables and other 50 Ω devices such as power splitters, mixers, attenuators, and directional couplers, to form a complete RF signal processing system. This is a very convenient way to proceed if you don't require low power consumption. These devices come packaged with connectors, or in transistor packages for PC board mounting. Most are *cascadable*, meaning that you can string them together without having them oscillate (unless you do something dumb, like folding the string in a U shape so that the output is next to the input).

14.7.11 Isolation Amps

Ordinary RF amps have resistive feedback, meaning that inside the package, a low resistance is connected from the input to the output. As a result, if you put a signal into the output of an amplifier, some will appear at its input. As a corollary, mismatching

the output will change the input impedance. This is usually not too objectionable, but sometimes the reverse signal path will introduce artifacts into your measurement. For example, if you need to produce two signals whose relative phases are accurately specified, having the two leak into one another can cause noticeable phase errors, as we saw in Section 13.6.9. This is often a concern when we are doing time and frequency measurements, which are attractive because of the extreme accuracy that is possible with current technology, or in a buffer stage following a sensitive resonator such as a SAW device or quartz oscillator.

One way to get more isolation is to cascade amplifiers with pads in between. An amplifier with 25 dB isolation followed by a 10 dB pad has 35 dB isolation, at the price of reduced gain and output power. Two sets in a row will get you 70 dB, which is usually enough. This approach is easy and is sure to work, but it is expensive and eats power and board space. It's great for prototyping, because the right combination of connectorized amplifiers and pads can be bayoneted together in a few minutes. Watch out for cable crosstalk, supply coupling, radiation, and other sneak paths when you're doing this—isolation amps plug one leak, but there are always more.

It is possible to build good amplifiers that have almost perfect isolation of the input from the outputs. A garden variety common-source amplifier using a \$0.75 dual-gate GaAs FET with no feedback to the gate can achieve 70 dB isolation in one stage at 100 MHz, with very small power consumption and board area. Apply feedback using a source degeneration resistor.

14.7.12 Radio ICs

There are a lot of wireless devices sold, and although a lot of wireless ICs are very specialized, a fair number are useful in instruments, too, especially the FM IFs, mixer/oscillators, and VCO/dividers; they are intended for communications use, and since communications puts surprisingly stringent requirements on RF signal processing, they have to perform pretty well. For a dollar or two, you get a miniature signal processing toolkit that runs from 3 V at 10 mA. Especially useful are the quadrature FM detector and logarithmic meter (RSSI) outputs, which have excellent performance and quick response. Before designing your signal processor, make sure you have the latest RF and wireless product data from NEC, Maxim, On Semiconductor, NXP (formerly Philips), National, and Texas Instruments at least, and that you've spent some time comparing the block diagrams of the parts with the one for your system.

Apart from their meter outputs, which are amazingly useful, the performance of these parts is never as good as you'd get from a collection of Mini Circuits stuff. Their noise figures tend to be around 6 dB, and their BJT mixers don't have as good dynamic range as a high level Schottky diode mixer. This matters a lot in optical instruments, since our dynamic range is often enormous. On the other hand, radio ICs tend to be much cheaper, and they chew up less power and board space. There are often circuit hacks to get round the dynamic range problem; for example, since the radio ICs are so cheap, use two complete systems, one with a 30 dB attenuator in front of it. Digitize both, and use the low gain system when the high gain one rails, as in Section 15.2.1.

14.7.13 Stability

Cascadability and isolation are related to the whole notion of stability of RF amplifiers. Most amplifiers have a significant amount of feedback (intentional or not) from output

to input. This will usually have a phase lag. If you put a capacitive load on the output, the feedback will be further delayed, until it starts to have a negative conductance. If this negative conductance gets too large, the net input resistance will go negative at some frequency, and the amplifier will potentially become an oscillator. Amplifiers whose feedback is sufficiently small that this cannot occur are said to be unconditionally stable. Isolation amplifiers are always unconditionally stable. If your amplifier is not unconditionally stable, don't use it for an input or output stage where somebody might hang an unterminated cable, or you'll be sorry. A pad on the input will usually solve the problem at the expense of gain and noise. There's more in Sections 15.4.1 and 18.4.1.

14.7.14 Slew Rate

The output of any device can move only so fast. The maximum dV/dt (or *slew rate*) is set by how fast the device's internal bias currents can charge its internal capacitances. Thus, as shown in Figure 14.18, a large-signal transient consists of a straight-line slew followed by an exponential settling from the step response, and possibly thermal and dielectric contributions later.

14.7.15 Settling Time

When the output of any device moves, it takes a while (the *slewing time*) to get there, and another while to settle down to a given accuracy; the time from the input transient to the final entry into the output error band is called the *settling time*.

For sufficiently small signals, the settling time of a linear device is easily deducible from the step response, as we saw with filters in Chapter 13. For larger signals, or with




Figure 14.18. Slew rate and settling time.

intrinsically nonlinear or time-dependent devices such as track/holds or DACs, this is not the case—the settling time must be characterized carefully, and surprises may lurk; it is easily possible for important settling transients to occur on time scales 10^4 times longer than the initial slew of the output.

Dielectric absorption and thermal transients within the IC are two of the most common causes of this, but there are also transmission line reflections and load effects to worry about. Make sure you scrutinize the settling behavior of your finished circuit very carefully—those long-tailed effects are not that easy to see, and almost never come out in simulations (SPICE is even worse at multiple time scales than a digital oscilloscope, because it can't skip stuff—it must laboriously follow every wiggle).

14.7.16 Limiting Amplifiers

A limiter is an amplifier that is intentionally severely overdriven, so that its output bangs between its high and low limits on each cycle. A typical limiter stage is one or more differential pairs, as in Figure 14.19, where the collector loads are small enough that the transistors cut off rather than saturating on each half-cycle. Heuristically, a limiter preserves only the timing of the signal's zero crossings, which is related to its frequency.

The relation is not trivial, however; if two signals at slightly different instantaneous frequency are fed into the limiter (e.g., two different FM stations on the same channel), the limiter exhibits a *capture effect*: the frequency of the signal coming out of the limiter is that of the stronger signal, even if it's only a very small amount stronger (1 dB, say). This behavior is a gift, because it allows us to coexist with our spurious signals unusually easily in FM systems. The phasor diagram in Figure 14.20 should make this more plausible (see Section 13.3.2). The two analytic signals are represented by phasors, which are vectors in the complex plane representing the instantaneous signal voltages. As the longer phasor spins around the origin, the smaller one cannot change the average

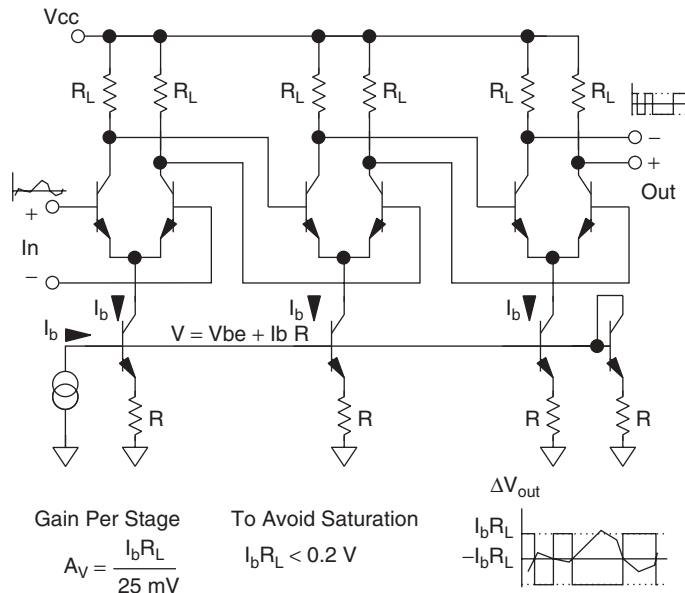


Figure 14.19. Three-stage BJT limiter with ultrasimple biasing.




Figure 14.20. Phasor diagram of two signals entering a limiter. The smaller is too short to reach the origin, so the larger sets the output frequency.

rate at which the total signal voltage goes around, because it's too short to reach. Thus the only way it can cause extra zero crossings is to make the total voltage briefly turn round and go backwards enough to cross and recross the imaginary axis. This requires the smaller signal to have a steeper slope than the larger one, and so cannot occur as long as

$$\omega_{\text{sig}} V_{\text{sig}} > \omega_{\text{spur}} V_{\text{spur}}. \quad (14.26)$$

We saw in Section 13.6.9 that additive spurs and noise cause phase jitter, which is preserved by limiting, so do your limiting after the bandwidth setting filter. Bandwidth limiting also prevents extra zero crossings, as shown by (14.26). Note that this is a bit of a two-edged sword; as the SNR in the IF bandwidth drops below 15 dB, you'll get momentary dropouts due to noise peaks bigger than the signal, and when it drops below 0 dB, the limiter will completely suppress the signal in favor of the noise.

Limiters suffer from AM–PM conversion (see Section 13.5.6), so limit at low frequency if possible, where the slowing down of the limiter with strong signals doesn't cause a large phase error. If you build your own, use bipolar diff pairs with small collector loads, as shown in Figure 14.19, so they cut off instead of saturating, which slows them down abysmally. You need overall DC feedback for stable operation, which is not shown but is easily added with an RC network.

Majority-carrier devices like GaAs FETs are somewhat better than bipolars here, if they're well enough matched. It's often worthwhile to convert, limit, and convert back if the phase accuracy must be high. Alternatively, a calibration of the AM–PM conversion can be made, and the phase data adjusted afterwards. It is also possible to compensate for it in real time with a phase shifter and some circuitry, if the data rate is too large or the latency requirements too tight for postprocessing. Section 15.7.2 has a good on-the-fly calibration method applicable to limiters.

Linear rectifying amplitude detectors also exhibit a weaker capture effect. Except near the strong signal's zero crossings, the weaker signal mostly just puts wiggles on the DC from the stronger one. Assuming that the two signals are at sufficiently different frequencies that their beat note is outside the baseband frequency response of the system, the effect of the weaker signal is reduced. The weaker signal can dominate the detector's output near the strong signal's zero crossings, so this capture effect isn't as pronounced as with a limiter.

14.7.17 Lock-in Amplifiers

A lock-in is basically a radio that measures the phase and amplitude of its input using two phase detectors driven in quadrature from a reference signal you supply (see Section 13.7.5).

Lock-ins have good dynamic range and contain circuitry that verifies that no stage of the amplifier is being overdriven, provided that the signals are slow enough (most lock-ins will be driven berserk by 2 volts of 100 MHz, but few will inform you of the fact). The value of lock-ins, and their fatal allure, is their apparent ability to conjure signal-to-noise improvements out of nowhere, and to provide phase measurements of high accuracy (0.1°) at frequencies up to 100 kHz. Users who do not expect more than their front ends will provide, and who read the spec sheets carefully, will find these devices very useful. Pay attention to the way the specified accuracy varies with frequency and amplitude, and remember that narrowing the bandwidth by a factor $1/\alpha$ in a continuously swept measurement increases the measurement time by a factor of α^2 (so that 10 dB better SNR costs $100\times$ in measurement time), because you've got α times as many points, each of which takes α times as long to settle). See Sections 13.10.4 and 17.11.5 for a good alternative.

One thing you have to be careful about in lock-in amplifiers is their calibration. Lock-ins using switching phase detectors produce an output corresponding to the average value of the commutated signal voltage, including both the signal frequency and all its harmonics. This means that you aren't just measuring the amplitude of the fundamental signal, and that your noise bandwidth may be quite different from the $1/(4\tau)$ of a continuous-time RC rolloff. Lock-ins using analog multipliers really measure just the fundamental. Digital lock-ins may do a better job, but there are no guarantees.

14.8 DIGITIZERS

Nearly all electro-optical instruments use digital data acquisition or digital control for something. We saw how to calculate the noise degradation caused by an ideal digitizer in Chapter 13; here we'll see what a real one looks like.

14.8.1 Digital-to-Analog Converters

The simplest place to begin is with DACs. A DAC in its purest form is a binary-weighted array of current sources or resistors, switched in and out bitwise in order to produce a current or a voltage ratio corresponding to the binary number input. Errors caused by switch impedance are solved by switching currents instead of voltages, or by scaling the die areas of the switches in the MSBs so that the switch resistance terms in the voltage

ratio cancel out. A subtler method is the $R-2R$ network, where all the resistors have the same value, and the switch impedances therefore don't need to be scaled. The $R-2R$ network can be inverted as well, with identical current sources connected to the taps and the output taken from the V_{ref} pin. A third method, used mainly in 8 bit applications, is to take 256 identical resistors wired in series and select a tap using a gigantic multiplexer and a noninverting buffer. This has the advantage of inherent monotonicity and stability, but tends to be slow.

Besides speed, resolution, and accuracy, the main figure of merit is *glitch energy*, which is the area on the $V(t)$ graph between the actual curve and some (poorly specified) smooth monotonic curve—basically it tells you how much total charge the glitch is going to deliver to the load.[†]

To get high speed and good linearity with low glitch energy, modern DACs often use combinations of these techniques, rather than being pure $R-2R$ or binary-weighted; for example, the Analog Devices AD9760 (12 bits, 160 MS/s), which uses 31 identical current sources for the 5 MSBs, another 15 identical current sources of 1/16 the value for the next 4 bits, and binary-scaled sources for the 3 LSBs. (See Figure 14.21.)

Since the binary input is dimensionless, the scale for the DAC output must be set by an independent reference; thus a DAC inherently multiplies the reference by the given binary fraction to produce the output. A *multiplying* DAC is one that is especially good at it, with good linearity and reasonable bandwidth for V_{ref} .

Get the Analog Devices and Maxim catalogs to see the wide variety of DAC types, and pay close attention to the octal serial DACs such as the AD8801; they're cheap and very useful for trimming, auto-zeroing, and low-speed control applications.

DACs are pretty trouble-free devices, providing you give them a clean, low- Z reference, a jitter-free clock, well-bypassed supplies, and really solid analog and digital grounds (see Chapter 16). There are audio-type 24 bit dual DACs available for a dollar or two, as well as much better ones (e.g., the Cirrus CS4398), which will solve most problems of DAC dynamic range, if not linearity.

One partial exception to this rosy picture is noise. An ideal DAC is a totally noiseless device. Unlike an ADC, a DAC contributes no quantization noise, since its output voltage corresponds exactly to the input digital code. Thus a DAC does not add the




Figure 14.21. $R-2R$ current-output DAC.

[†]You're right, it's not really an energy—it's specified in picovolt-seconds, so it's an impulse.

$1/\sqrt{12}$ ADU quantization noise of an ADC. Of course, since no DAC has perfectly spaced steps, switches instantaneously without jitter or glitches, or produces noiseless levels, the real situation is somewhat worse.[†] Even with constant DC level outputs, voltage-mode CMOS DACs can be significantly noisier than current-mode or bipolar voltage DACs—100–200 nV/Hz^{1/2} noise, with $1/f$ noise corners as high as 10 kHz, which is a disaster in high accuracy situations since it can't be filtered out. Some DACs also exhibit bad popcorn noise. Serial DACs have a lot of clock feedthrough—for best performance, use a parallel-access DAC and an external op amp.

14.8.2 Track/Hold Amplifiers

Track-and-holds[‡] are the workhorses of A/D conversion. The track/hold does just what it says: tracks the signal and then holds its last value when told to. Good track/holds (T/Hs) are very linear, settle rapidly in both states, don't load the signal source unduly, have low and stable output impedances, don't produce big glitches on their inputs or especially their outputs, and of course hold the level without drooping at all. We're still waiting for one like that, but modern T/Hs are very good.

Track/holds are used to hold a voltage so that it can be digitized by an A/D, or to deglitch a DAC. Deglitching requires only modest timing accuracy in the T/H, but of course its own glitches must be small. ADC front ends require excellent timing accuracy and low droop. We expect high slew rates and fast settling during track mode, and good isolation of the stored voltage in hold mode. The crucial timing parameter in a track/hold is the *aperture uncertainty*, or *aperture jitter*.[§] This is the uncertainty in just when the sampling occurs. It's usually random, but not necessarily, as it may depend on signal swings. Manufacturers aren't very good at specifying how they measure this, so we'd better be conservative and assume that the published spec is the rms error, rather than the peak to peak.

We can estimate the voltage noise caused by an rms aperture uncertainty t_{ap} with a full scale signal at f ,

$$\left\langle \frac{\Delta V}{V_{\text{FS}}} \right\rangle = \frac{t_{\text{ap}}}{V_{\text{FS}}} \sqrt{\left\langle \left(\frac{dV}{dt} \right)^2 \right\rangle} = \frac{\pi f t_{\text{ap}}}{\sqrt{2}}. \quad (14.27)$$

In order to hold this to less than $1/\sqrt{12}$ ADU, so that it doesn't dominate the noise for large signals, the jitter must be less than

$$t_{\text{ap}} < \frac{2^{-N}}{\pi \sqrt{6} f} \approx \frac{2^{-N}}{8 f}. \quad (14.28)$$

For 12 bits at 1 MHz, $t_{\text{ap}} < 30$ ps, even allowing aperture jitter to contribute as much noise as the quantization itself. If you know that you will never have signals varying that

[†]There are also digital signal processing issues such as the use of zero-order holds—see Chapter 17.

[‡]In a simpler age, these used to be called *sample/hold* amps. Now we have to worry more about what happens during acquisition.

[§]Aperture uncertainty is occasionally confused with the *aperture time*, which is the time window in which the sampling occurs. This should stay still, so it isn't the same as jitter.

fast, you can get away with poorer jitter, but don't push it too far—bad performance there can look harmless in an FFT spectrum (merely raising the noise floor a bit), but still have horrendous time-domain artifacts (e.g., glitches and missing codes). Extensive digital postprocessing puts a huge premium on the quality of the data.

The maximum aperture uncertainty is much less than the rise time of the sampling clock edge, and slow or noisy clock edges make it worse. Clock jitter of course adds in RMS fashion as well; if you're using PLL clock recovery, for instance, that jitter counts too. Aperture uncertainty is an almost impossible limitation in fast, high resolution converters; if our 12 bit ADC were 16 bit instead, the jitter spec would be 2 ps. There are T/Hs available that are this good, but remember that this spec has to be met by your entire clocking system. This is really impracticable without great care. It may be sufficient to resynchronize the T/H control line directly from the master clock oscillator using a fast flip-flop (e.g., a 74AC74A), but in some instances the clock has to be filtered first. The timing accuracy of the rest of the system has to be considered too, especially if your high frequency component is ultimately derived from another clock somewhere.

Gotcha: T/H Outputs Are Also Inputs. Track-and-holds tend to be vulnerable to transients at their outputs. To achieve simultaneous sampling of several inputs, it's common to use several T/Hs with a multiplexer following, and a single ADC. If the multiplexer is charged up to a very different voltage, its capacitance will have to discharge very rapidly through our T/H's output pin, and some of this current is liable to find its way onto the hold capacitor, producing a spurious step in the stored voltage. A resistor between the T/H and mux will help, but in any such system, make sure you test for this problem.

14.8.3 Analog-to-Digital Converters

There are lots of different kinds of ADCs, but the ones you'll probably use most are half-flash for high speed and $\Delta\Sigma$ for high accuracy at low speed.

Flash ADCs. The easiest kind of ADC to understand is the flash converter (Figure 14.22), where you have $2^N - 1$ separate comparators (why not 2^N ?), with all their + inputs connected together and their - inputs connected to a $2^N - 1$ tap voltage divider. The output is constructed as the binary address of the highest comparator that goes active. Flash converters are very fast, but their complexity increases exponentially with N . You might think that they would have excellent aperture jitter, but it isn't as good as all that, due to a fixed-pattern effect; because of chip gradients and the different source impedance of each tap, the comparators don't all have the same delay. Accordingly, you should use a T/H anyway.

Half-Flash ADCs. A half-flash converter (Figure 14.23) uses a two-stage process where a very accurate $N/2$ bit flash converter gets the MSBs. An equally accurate DAC subtracts their contribution off, leaving a residual that is amplified by $2^{N/2}$ times and digitized by the same or another flash converter, yielding the LSBs. This is complicated and error-prone, but it's worth it for the reduced circuit complexity and power consumption. Modern half-flash units use an extra bit or two in the flash section, plus some on-chip smarts for self-calibration, relaxing the accuracy requirements. As an aside, the DACs inside ADCs are often based on capacitive voltage division instead of resistive, which has precision and power consumption advantages, but can't hold a fixed level.




Figure 14.22. Flash analog-to-digital converter.

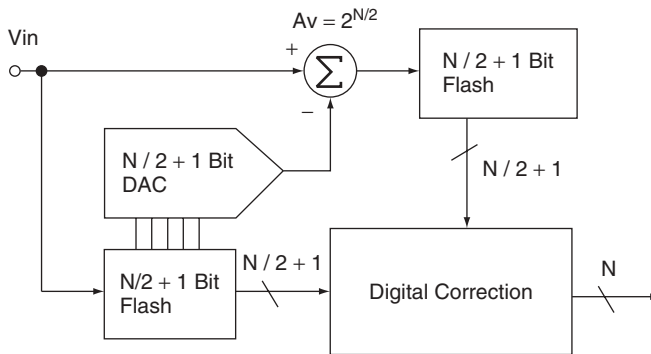


Figure 14.23. Half-flash ADC.

Successive Approximation ADCs. Occasionally, you’ll need to use successive approximation, which uses a DAC, a single comparator, and some logic (a successive approximation register, or SAR). In $N + 1$ clock cycles, the SAR constructs the output word one bit at a time (MSB first) by doing a binary search of the DAC’s code space to find the exact point at which the comparator switches. The beauty of SAR circuits is that since the step size varies with the uncertainty (cf. slew rate in an analog loop) they are inherently nulling without being exponentially slow for large N as $\Delta\Sigma$ or integrating converters are. This combination makes them good for nulling-type phase and amplitude digitizers, which let us avoid two-dimensional calibrations.

Delta-Sigma ($\Delta\Sigma$) ADC. A $\Delta\Sigma$ ADC, also called a delta-sigma, bitstream, or “1 bit” ADC (Figure 14.24), uses a single comparator in a digital feedback loop, to continually

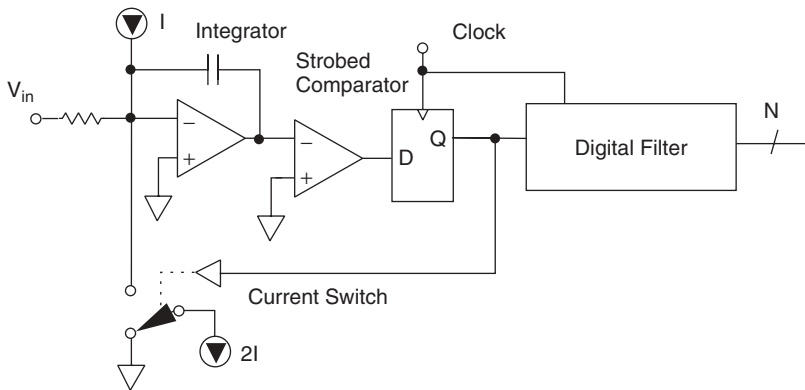


Figure 14.24. Delta-sigma ADC (bitstream ADC).

null out the input signal with a single switched current source.[†] Its output is the duty cycle required, expressed as a binary fraction after digital filtering. Only one threshold and one current source are involved, so there are no matching problems, and the capacitor and comparator voltages are always near 0. Thus its linearity is excellent. Besides, they come with up to 24 bits—8 more than any other technique—so what’s not to like?

Due perhaps to their association with “high end” audio, or to the ease of adding lots of extra bits—useless, but always irresistible to marketers—delta-sigmas have attracted about as much hype and specsmanship as anything in electronics, so you have to be really careful in reading their data sheets. This is a pity, because they’re amazingly good for some things. It’s partly explained by the fact that the theory of $\Delta\Sigma$ s is much richer (read, more difficult) than that of ordinary quantizers because of the quantized, and hence nonlinear, feedback. This means that not all the usual theorems about quantization noise being equally distributed and white apply.[‡]

They aren’t very fast; $\Delta\Sigma$ ADCs slow down exponentially with N , needing 2^N clocks per independent measurement. SARs need N clocks, and an analog feedback loop settles to N bits in about $0.7N$ time constants, because the slew rate is proportional to the size of the error, whereas in a $\Delta\Sigma$ or slope converter it is constant. A $\Delta\Sigma$ integrates continuously, so that each measurement merely samples its output. Self-calibration is usually used to null out the current source errors and voltage offsets here too. The slowness of $\Delta\Sigma$ converters limits their usefulness, because you generally can’t put a multiplexer in front to sample several different sources, as you can with devices whose conversion time is shorter and better defined. (Some $\Delta\Sigma$ s use multibit front-end converters, almost like a half-flash; if more than 1 bit feedback is used, this trades off accuracy for a bit more speed.)

Their linearity is very good; $\Delta\Sigma$ s have INLs about $\pm 10^{-5}$ (and typically the rest of your circuit won’t be this good). However, a 24 bit A/D has 16 million codes, so 10^{-5} of that is ± 160 ADUs. Similarly, they’re usually claimed to be monotonic to their maximum resolution, but due to their noise, typically they won’t even *sit still* to that accuracy, even

[†]Another oldie in new clothes: H. Inose et al., New modulation technique simplifies circuits. *Electronics* **36** (4), 52–55 (1963).

[‡]Robert M. Gray, Quantization noise in $\Delta\Sigma$ A/D converters, in S. Norsworthy et al., *Delta-Sigma Data Converters*, Wiley-IEEE, Hoboken, NJ, 1996.

with their inputs shorted. (Don't expect a $\Delta\Sigma$ to get anywhere near $\text{ADU}/\sqrt{12}$ noise.) So basically your 24 bit $\Delta\Sigma$ is a very nice 18 bit converter with six nearly useless extra bits attached by the marketing department.

Their popularity is explained largely by their low prices, good accuracy, freedom from range switching, and (secretly) the effortless ease with which they allow a designer to finesse a stupidly written accuracy specification by waving their apparent 60-parts-per-billion resolution. This makes them the circuit equivalent of switching to double precision floating-point without changing your algorithm.

14.8.4 DAC and ADC Pathologies

DACs and ADCs never have the nice even staircase characteristic we'd like, and the characteristic degrades somewhat with fast signals. For slowly varying signals, the main types of error are the static *differential* and *integral nonlinearity* (DNL and INL, respectively), defined visually in Figure 14.25 for a DAC. The INL is the maximum deviation of the output voltage from a perfect staircase connecting the upper and lower reference voltages, in units of ADU. The DNL is the deviation of each individual step from being exactly 1 ADU tall, which is easily expressed in finite differences:

$$\text{DNL}(n) = \frac{V(n+1) - V(n)}{2^{-N} V_{FS}} - 1 \quad (14.29)$$

A DNL of ≤ -1 ADU makes the DAC nonmonotonic, which is very undesirable in control applications.

The ADC case is a bit more complicated, because we're looking at Figure 14.25 sideways, coming in with a voltage, and coming out with a number. From what we know about inverse functions, we expect problems if the DAC isn't monotonic; here, a DNL of ≤ -1 ADU at code M will make the inverse function multivalued, and that makes M a *missing code*; a slowly varying input will produce an output that jumps directly




Figure 14.25. DAC pathologies: integral and differential nonlinearity.

from $M - 1$ to $M + 1$ (3 to 5 here), which is even worse than DAC nonmonotonicity, especially in feedback systems, where a set point of M leads to the loop hunting forever. (Big positive DNL leads to deadbands, which are just as bad.) Many of the highest resolution ADCs have many ADU of input noise, which causes them not to sit still, even with their inputs shorted; this smears out the DNL and may make it less significant in servo applications.

INL and DNL have effects identical to putting a nonlinear amplifier in front of a perfect digitizer (or after a perfect DAC). Broadly speaking, DNL causes time-domain glitches and raises the noise floor, while INL causes harmonics and intermod; thus you should test for DNL in the time domain and INL in the frequency domain.

You test ADCs in the frequency domain with two sine waves, just like a two-tone IMD test for an amplifier. Do the frequency analysis with an unwindowed FFT—don't use a DAC plus a spectrum analyzer. Make sure the test signals and sampling clock are very stable, and placed so the tones and IM products land exactly at sample frequencies (e.g., 65,536 samples, with f_1 going through 100.0000 cycles and f_2 103.0000 cycles in that interval—see Section 17.4.7). Phase-lock them together if possible—these conditions force all the products to concentrate where they're easily measured. INL is usually gentle and creates low-lying harmonics that show up well.

Check very carefully for DNL in the time domain, by subtracting off what you ought to have got (samples of two sine waves) and looking at the residuals. It's often a very illuminating test.

The same tests run backwards are good for testing DACs; use a spectrum analyzer to look for the intermodulation products and harmonics. The DNL test is best done with a good oscilloscope and a second DAC driven with the twos complement; you can sort out which DAC is being nonlinear with a digital adder in series with one of the DACs' inputs; adding 1 to the DAC code will shift its glitches over by a clock period. Subtracting the two traces leaves only the DNL of the shifted DAC, which is very convenient. Another point is that the dynamic performance of ADCs will depend on the relative phases of the clock and test tones; if you have phase-shiftable synthesizers, experiment with different phases to see if it makes any difference.

14.8.5 Differential Nonlinearity and Histograms

Differential nonlinearity is especially obnoxious when we're doing histograms, since the expectation value of the count in each histogram bin is proportional to the bin width, so that any differential nonlinearity introduces a systematic error into the histogram. This is not even slightly subtle— $\pm \frac{1}{2}$ ADU DNL corresponds to a 3:1 range of bin widths. There are ADCs with intrinsically identical bin widths, such as voltage-to-frequency or time-to-amplitude converters and $\Delta\Sigma$ ADCs, which are a better match for histogramming applications; an inferior expedient is to use a higher resolution converter and sum adjacent bins. Another approach is *dithering*: use a fine-resolution DAC to add pseudorandom noise to the signal, and subtract it again numerically after digitizing. This homogenizes out much of the DNL at some sacrifice in complexity and dynamic range (since you don't want to clip the signal + dither waveform). (Hint: Since DNL is localized, it takes *lots* of noise, not just a few ADU, and requires very close attention to the impulse response so the undithering can be done accurately—test the daylights out of it.)

If you're oversampling significantly, the error-smearing approach can be used in a much simpler way: add out-of-band noise at the ADC input, and filter it out again

digitally afterwards (see Section 17.6). Again, you need to use a fair amount—at least several ADU, and maybe 1/4 of full scale.

14.8.6 Dynamic Errors

The two-tone test we used earlier only tests for low speed errors. High speed testing is done just the same way, but aliased (see Chapter 17). The measurement conditions are typically something like this: $f_{\text{sample}} = 10.01$ MHz, $f_{\text{in}} = 10.00$ MHz, $V_{\text{in}} = 95\%$ of full scale. The desired output shows up at 10 kHz, and the spurious signals wind up at harmonics of 10 kHz. Two-tone testing might have $f_2 = 10.031$ MHz. (See Figure 14.26.)

ADCs intended for signal processing usually have AC specs quoted in terms of spurious-free dynamic range (SFDR), which is the power ratio of the largest sinusoidal signal the ADC can handle to the largest of these spurs (usually the second or third), under given testing conditions. Another way of specifying it is the effective number of bits (ENOB), which is based on the SINAD (Signal to Noise And Distortion, a generalized SNR); you take the largest AC swing you can accommodate, $V_{\text{rms}} = 1/(2\sqrt{2})$ times the reference voltage, and divide it by the $1/\sqrt{12}$ ADU additive noise of a perfect digitizer, take 20 log to get decibels, and you get the mysterious-looking formula

$$\text{SINAD}|_{\text{ideal}} = 20 \left[N \log 2 + \log \left[\frac{\sqrt{12}}{\sqrt{8}} \right] \right] = 6.02N + 1.76 \text{ dB.} \quad (14.30)$$

The ENOB is what you get by plugging in the observed SINAD and solving (14.30) for N . It isn't that useful for calculations, but the number is easily memorable and gives a good seat-of-the-pants feel for how much worse the ADC gets at high frequency. Data sheet preambles and signal processing books will give you precise definitions for all these things, but when you look at the actual specifications, you find that the specs are uncertain by factors of 3 to 10 in SFDR anyway, so don't sweat it too much.

These aliased tests are pretty stringent, although they don't tell you a lot about the time-domain performance; all those spurs might add up to one great big time-domain




Figure 14.26. Dynamic errors: spurious-free dynamic range.

glitch due to 3 missing codes right at the zero crossing, for example; or it might have horrible cross-modulation due to an internal node that can't slew rapidly; or it might have lots of jitter, which will raise the noise floor. You care about these things, even if the SFDR looks OK, because the effects are nonlinear and the test signal isn't the same as your real data. Make sure that you test for what you care about, including at least the noise floor, missing codes, intermodulation, offset, and harmonics. Leave a big fat safety factor if you can possibly afford to—ADCs can surprise you in a lot of unpleasant ways. Ignore their idiosyncrasies at your peril.

14.8.7 Dynamic Range

The SINAD formula (14.30) is appropriate for ADCs working on AC signals, where both half-cycles have to be digitized, and the maximum rms signal is $1/(2\sqrt{2})$ of full scale. A converter dealing with unipolar signals (e.g., the output of an AM detector) has things considerably better; an ideal N bit ADC has a dynamic range (DR) of

$$\text{DR(dB)} = 20 \left[N \log 2 + \log \sqrt{12} \right] \approx 6.02N + 10.78 \text{ dB}, \quad (14.31)$$

which is 9 dB better than the SINAD number (14.30) due to not needing to accommodate the peak-to-peak swing of a sine wave.[†] Keeping these straight is not difficult if you bear in mind what's really going on at the ADC input (see Section 15.2.1).

The dynamic range of a DAC is somewhat thornier philosophically, because an ideal DAC introduces no error whatever; its output voltage exactly corresponds to the input word.[‡] The quantization is what introduces the noise, and that has to have occurred elsewhere before the DAC word could have been generated. Accordingly, we usually talk about the *resolution* of a DAC instead. Measurements of the output spectrum of a DAC, as in the two-tone SINAD measurement above, inherently include the quantization noise contributed by the computation of the approximate sinusoids used as the input. It's worth keeping the two distinct in our minds, however, because in error propagation calculations we don't want to add in the quantization error twice. Once again, *ADCs add quantization noise, DACs do not.*

14.8.8 ADC Noise

The noise of a DAC is easily measured—sit it at various levels and look at the AC part of its output. ADC noise is a bit more problematical; it is a random motion of the comparator thresholds, causing inputs near the edge of a bin to show up randomly in the two adjoining bins. The easiest way to see it is with a low amplitude ramp signal; sweep repetitively and look for fuzziness on the edges. Remember that the conversion itself contributes $1/\sqrt{12}$ ADU of noise, so noise at least 10 dB below that is probably OK, since we aren't letting the quantization noise dominate in the first place. Sometimes we even add noise intentionally, to smear out the staircase and allow signal averaging.

[†]Another way of looking at this is the full-scale DC signal of a converter in a unipolar circuit has $2\sqrt{2}$ times the rms value of a full-scale sine wave, for the same converter used in bipolar mode.

[‡]Real DACs, of course, do produce output noise, but it's easy to measure, and usually not highly relevant.

14.8.9 Ultrafast ADCs

Really quick ADCs, such as 25 MHz at 14 bits, 50 MHz at 12 bits, or 100 MHz at 10 bits, are about as easy to use as a recipe for unicorn stew. Their performance changes dramatically with tiny changes in conditions.[†] If you can't possibly avoid these monolithic slabs of misery, get the manufacturer's evaluation board, and bang on it. Copy his layout and component choices slavishly, except *don't split the ground plane* (see Section 16.4.4). Do not attempt to economize on the circuit card or decoupling components when using one of these. There exist even more ridiculously fast ADCs—like 16 bits at 130 MHz—but their SNR levels are far worse than the quantization limit; they're mostly used in *software-defined radio*, where digitizing is done early in the signal processing chain and the ADC and clock jitter problems are spread out over a wide bandwidth and subsequently filtered out digitally. These usually have optional dither and noise injection to cover up the bad ADC behavior.

14.9 ANALOG BEHAVIOR OF DIGITAL CIRCUITS

14.9.1 Frequency Dividers

Frequency dividers are usually digital counters, though there are other kinds (parametric dividers and injection-locked oscillators, in particular). You can use a binary counter for sequencing your data acquisition and control tasks; a walking-ring (Johnson) counter for generating equally spaced phases for an SSB mixer; or a loop divider in a phase-locked loop frequency multiplier. There are a couple of analogish things to worry about with counters. The first is the usual sensitivity of the logic transition to supply noise and poor clock edges, which lead to timing jitter and to fuzz on the output voltage levels that will get into subsequent stages—pay attention to your bypassing if you care about that, and consider resynchronizing with an additional flip-flop stage at the end of the logic chain. (Microprocessor and FPGA outputs are especially bad for this, with poorly defined logic levels and lots of spurious junk; zero-power programmable logic and CMOS gates and flip-flops are much better.) The second is more mathematical: a divide-by- N counter has an N -fold phase ambiguity. This is seriously annoying in situations where separate units have to agree on the phase (e.g., in time-of-flight ranging). The reason is that the counter can wake up in any of N states, and even if you reset it right away, state 0 is no more likely to be correct than any of the other states. One thing to remember about clocked logic is that it is vulnerable to cable reflections at its outputs, which can cause extra clocks, missed clocks, and metastability.

14.9.2 Phase Noise and Jitter of Logic

Digital circuits have noise just like analog ones, but (as in comparators) it shows up as timing jitter. A good HCMOS system can have 10 ps jitter, 10KH ECL a few picoseconds, and ECLinPS Lite can be in the 100 fs range. Chaining functions together makes this worse, generally faster than the rms sum since variations in temperature and supply

[†]For carefully documented fast ADC horror stories and discussion, see Bill Travis, Demystifying ADCs: EDN Hands-On Project. *EDN* 42, 7 (March 27, 1997), and its sequel, Remystifying ADCs. *EDN* 42, 21 (October 9, 1997).

voltage cause systematic changes. Resynchronizing with a flip-flop running from a quiet supply and a good clock will help a lot. In bad cases, two stages of resynchronization can be better than just one.

14.9.3 Analog Uses of Gates and Inverters

CMOS gates especially have lots of analog uses. One of the best is in driving diode bridge mixers. A lot of the intermodulation spurs from diode mixers are generated during the transition times, when the diodes aren't fully on or off, so using a square wave LO improves the intermod performance. Most mixers have transformer inputs and like to have their ports properly terminated; drive them via a 47 ohm resistor and a capacitor in series. They're also good for driving MOSFET switches. Section 18.7.2 shows a multiplexer for a pyroelectric imager, in which diode switches are driven by CMOS shift register outputs. Three-state outputs can be used to switch channels, for example, by grounding one of N photodiode anodes, when the cathodes are all connected to single transimpedance amplifier. Zero-power programmable logic swings reliably from rail to rail, which makes it good for analog jobs.

Electronic Subsystem Design

When art critics get together, they talk about style, composition, and technique. When artists get together, all they talk about is where to get the best turpentine.

—Pablo Picasso

15.1 INTRODUCTION

In Chapter 14, we talked about electronic components and fundamental circuits. Chapter 18 is about building fast, quiet photodiode amplifiers. Here in the middle, we'll discuss design philosophy and strategies—how to approach the problem of synthesizing an electronic subsystem, with emphasis on the analog back end: RF signal processing, pulse detection, and analog to digital conversion. There are some digital bits, but they're simple—counters, registers, and state machines.

The most beautiful circuit and instrument designs come out of a really deep knowledge of undergraduate physics and engineering. Beauty in this sense consists not only in conceptual simplicity and high performance, but in robustness, manufacturability, and low cost.

15.2 DESIGN APPROACHES

The right approach to a design depends a lot on where you're starting from. If you have a clean sheet of paper in front of you, you have choices not available to the guy who's just putting more buttons on an existing design.

There is no substitute for experience here, but fortunately it doesn't all have to be your experience—there's a lot of design lore available. Really good analog circuit designers are fairly rare but have a subculture of their own. Have a look in some of Jim Williams's books, the ARRL handbook, and the linear applications books of National Semiconductor, Linear Technology, and Philips Semiconductor (see the Appendix). The Usenet group `sci.electronics.design` is a good place for detailed design questions. For help in getting started, try `sci.electronics.basics`.

Some of the circuit elements in this section, especially the details of the SA605 IF chip and its RSSI (meter) outputs are discussed later in the chapter, so turn there if you're unfamiliar with them.

15.2.1 Describe the Problem Carefully

The first thing you need is a good, brief description of the problem. If you're in charge of a big chunk of the design, it need not be elaborate; something like the following, which is a leaked excerpt from the system spec for an acousto-optically and galvanometrically scanned laser bug zapper (coming soon to a store near you). The system concept is to use a HeNe beam to sense flying insects, and a 50 mJ Q -switched Nd:YAG laser to vaporize them.[†] The HeNe beam makes two passes through the A-O cell, so that its frequency offset is doubled.

Section 2.23: Detector/Digitizer The detector/digitizer must detect Gaussian-envelope tone bursts, 1 μ s wide at $1/e^2$, arriving in a Poisson process of average rate less than 1 kHz. The peak amplitude of the bursts is 0.2 mV to 1 V (50 Ω), and the carrier frequency f_C is 100–200 MHz. The value of f_C depends on the laser tuning and system time delays, and so must be software-adjustable, but at any given time it is known in advance to an accuracy of ± 5 MHz. The tone bursts are immersed in additive white Gaussian noise of -160 dBm/Hz. The system is to produce a 1 μ s trigger pulse beginning within 1.5 μ s of each event. Its output is to be three 14-bit serial digital words per event, corresponding to the peak amplitude $\pm 0.2\%$ (1σ), scan position and laser power. Efforts should be made to exceed these minimum dynamic range specifications as far as practicable, since improvements there will translate directly into improved instrument performance. The digitizer must accept external control lines DDEN' (negative-true), which enables triggering the digitizer on detection of a target, and XTRIG' (falling edge active), which initiates an immediate conversion cycle regardless of the target status. XTRIG' will be asserted not less than 1 μ s after DDEN' goes active.

15.2.2 Systems Engineers and Thermodynamics

Now it's time for a back-of-the-envelope calculation to determine how hard it is, and in this case we immediately find that the spec is full of problems.

Noise Floor. As we showed in Section 13.6.2, the 1 Hz Johnson noise power in an impedance-matched circuit is kT , which is -174 dBm/Hz at room temperature. The specified noise floor is already 14 dB above there. A garden variety low noise amplifier has about a 3 dB noise figure (i.e., $T_N = 300$ K). That kT per hertz and the incoming noise power add, so the fractional change is

$$\Delta P_N(\text{dB}) = 10 \log \left(\frac{10^{1.4} + 10^{0.3}}{10^{1.4}} \right) = 10 \log 1.079 = 0.33 \text{ dB}. \quad (15.1)$$

(This sort of move is how you should think of most things in signal processing, because the absolute levels generally aren't as significant as the ratios.) Our part of the system should thus be reasonably quiet, but needn't be as quiet as a front end amplifier—so far, so good.

[†]The optical designer is having his problems too—distinguishing bugs from aircraft, for example, and deciding what to do about a mosquito that's just landing on someone.

Frequency Plan. The wide frequency range sounds like a job for a superhet, like the AM radio of Example 13.6, because the measurement bandwidth is so much narrower than the tuning range and because the pulse envelope is what we care about, just as in AM. We need to get the LO from somewhere, and a tuning dial obviously won't do it—perhaps a DAC-controlled VCO or perhaps a low resolution frequency synthesizer, depending on how fast that “known in advance” changes. It isn't immediately obvious how we're supposed to know what the frequency is to be—do we get a reference signal, a digital word, or what? The instantaneous measurement bandwidth is at least 11 MHz, counting ± 5 MHz for the carrier frequency uncertainty and a 1 MHz δf for the tone bursts ($\delta f \approx 1/\text{FWHM}$). For best SNR, we'd like to use a matched filter, whose bandwidth will be $\sim \delta f$ (see Section 13.8.10). Because the noise power goes as the bandwidth, we're giving away at least 10 dB in sensitivity by accepting that huge f_c uncertainty, at least if we do our detection in one band, and it'd be worth hearing the reason it has to be so big.

Dynamic Range. That +10 dBm (10 mW) maximum signal level looks a bit ugly. Remember from Section 13.5 that the 1 dB compression level is typically $P_{\text{LO}} - 7$ dB, and that the amount of compression is approximately linear in P_{RF} for small compressions—that is, the 0.5 dB compression point is 50% of the 1 dB point ($\approx P_{\text{LO}} - 10$ dB). That 0.2% is 0.017 dB, which is 1/60 of 1 dB. By our linear rule, the compression point has to be about 60×10 mW, a little matter of 600 mW—even ignoring the first stage gain. Using the $P_{1\text{dB}} = P_{\text{LO}} - 7$ dB rule, we'll need a diode bridge mixer with about a 3 W LO level. You can't get mixers that strong, and even if you could, there's just something wrong with having to put a heat sink on a mixer. Even if we perpetrated such an atrocity, we'd still be in the soup because (due to their insertion loss) diode mixers have about a 6 dB noise figure. Because of that, we'd need at least 10 dB of front end gain to keep the noise floor where it should be, so the formula predicts an impossible 30 W of LO injection. You might want to verify that assertion, as in (15.1). Things are definitely looking a bit stickier.

One possibility would be to attenuate the signal before mixing it, because a signal N dB smaller signal needs N dB less LO power. Unfortunately, that makes the problems worse at low P_{in} , because an N dB attenuator has an N dB noise figure if it's connected to a matched load at 300 K—which means that the weakest signals will disappear into the Johnson noise of the attenuator. Might be a job for two subranges (using range switching or two complete receivers), or possibly even direct detection at the RF frequency. Let's press on.

Digitizer. The signal power level inside a tone burst ranges from 400 pW to 10 mW (−64 dBm to +10 dBm), a 74 dB range. As we saw in Section 13.6.9, the amplitude error $\langle \Delta A \rangle / A = 1/(2 \times \text{CNR})^{1/2}$, so if we have a spec on $\langle \Delta A \rangle$, we can convert that to a CNR spec: $\text{CNR} \geq 1/(2(\langle \Delta A \rangle / A)^2)$. Meeting our specified relative accuracy of 0.2% thus requires a carrier-to-noise ratio of $1/(2 \times 0.002^2) \approx 125,000$ or 51 dB, so the required digitizer dynamic range is 74 dB for the signal level plus 51 dB for the accuracy, or 125 dB. Assuming that “1 σ ” means the $1/\sqrt{12}$ ADU quantization noise, we know from (14.31) that in a unipolar measurement, an ideal N bit ADC has a dynamic range of $6.02N + 10.78$ dB, which is 95 dB for 14 bits and 107 dB for 16 bits. Thus the signal makes it inside the 14 bit nominal dynamic range, but not by enough to give 0.2% relative accuracy anywhere near the lower limit—and that's assuming we really

have that low a noise floor in the ADC, which is starting to be dubious at that resolution. Perhaps the spec means *0.2% of full scale*, in which case we're fine. On the other hand, that needs only 9 bits, so it isn't obvious what the lower 5 bits are for in that case. Alternatively, we might need to use an ADC with enough bits to span that 125 dB, but that's a 21 bit part; these exist but are extremely slow, more like 10 ms than 1 μ s, and never ever get near those dynamic range numbers in real life (see Section 14.8.3). We'll have to clarify the spec there, by thinking some more or asking the system designer.

Noise Budget. Here's where the real bite comes. With a 51 dB CNR spec at $P_{in} = -64$ dBm, the noise power in the measurement bandwidth has to be less than -64 dBm $- 51$ dB = -115 dBm. Our input noise floor is -160 dBm/Hz, remember, so in our 11 MHz bandwidth, the noise power is -160 dBm/Hz + $10 \log(11 \text{ MHz}/1 \text{ Hz})$, which comes to -89 dBm; we're off by a good 25 dB from being feasible here, regardless of the ADC. Now it's really time to go talk to the systems person (or possibly to the mirror, depending on the size of the engineering team).

* * * *

That somewhat awkward interview is over, and we have a revised accuracy spec of "0.2% of full scale or ± 1 dB, whichever is less." It also became clear that that 0.2% is a bit squishy, because there's going to be an online, end-to-end calibration facility in the instrument, so what we really need is 0.2% repeatability and stability over times of a few minutes. ("So *that's* what DDEN' and XTRIG' were for.") The frequency uncertainty has been tightened to ± 2 MHz by making some other poor slob work harder, but that's the limit, because of the Doppler shift of the scatterer.

It turns out that the need for the wide tuning range comes from scanning back and forth. Because of the signal detection delay, by the time the detector/digitizer has produced its pulse, the HeNe beam will no longer be on the bug, so the pulsed YAG laser has to trail it slightly. The angular offset required is opposite on the two half-cycles of the scan, so we have to use an acousto-optic deflector (AOD) on the HeNe beam and change the acoustic frequency as we go. This doesn't have to track very fast (1 ms or so), so we can probably derive our LO from the AOD drive signal, although it'll have to be frequency-doubled and offset; the HeNe beam is diffracted on both transmit and receive, which doubles the frequency shift (see Section 7.10.7).

ADC Resolution Again. At the 0.2 mV end, the ± 1 dB accuracy spec means that the CNR has to be $\geq 1/(2 \times (10^{0.05} - 1)^2) \approx 15$ dB, at the very least (see Section 13.6.9), which is still pushing it a little, since it implies a dynamic range of 74 dB + 15 dB = 89 dB. Furthermore, we want each measurement to be accurate to 1 dB at worst, which means we can't use the rms uncertainty of $1/\sqrt{12}$ ADU, because that's an average over an ensemble of measurements. Perfect quantization has a maximum error of 0.5 ADU, so even with a very good DNL spec of ± 0.5 ADU DNL, we're looking at ± 1 ADU at best. Since we have to deal with the inaccuracy of the calibration, it's really more like ± 2 ADU. We're making a unipolar measurement, so we don't have to accommodate the peak-peak range of an AC signal in our ADC range (a factor of $2\sqrt{2} \approx 9$ dB). Our 14 bit A/D has 16384 ADU signal/2 ADU error ≈ 78 dB dynamic range on this basis. We'll have to compress the signal in the lower range a bit, but that isn't too hard to do

at 1 dB accuracy. The door is open to a nonlinear pulse height detector, and probably to our multiple subrange idea.

We couldn't get the front end gain cranked up any higher to loosen our noise figure spec. The maximum signal is already +10 dBm, and making that much bigger will start to cause linearity, EMI, and power consumption problems. We conceded the point reluctantly. It isn't that the expensive mixer with the 3 W LO power wasn't in our minds too, but that end-to-end calibration will greatly ease the compression problem, allowing a much more sensible LO level—we can comfortably go right up to the 1 dB compression point. We can use a low noise amplifier on the low range, where we have noise floor problems, and an attenuator on the high range, where it's saturation we're worried about.

Aside: The Old Way. In the palmy days (say, 1982), when accurate ADCs were very expensive, we'd have been thinking along different lines: "Well, allowing a factor of 2 for calibration errors, $\pm 0.2\%$ of full scale needs 500 codes all told, and ± 1 dB for the remaining 40 dB or so needs another 40. We can do it in 10 bits, with a carefully tweaked DLVA summed with an AM detector, and maybe 9 if we pay close attention to calibration errors." Now that 14 bit ADCs are cheaper than tweaking, and computing power is cheap too, using one of those plus software calibration is a lot more reasonable.

15.2.3 Guess a Block Diagram

The next thing is to make a guess as to a good block diagram of the system. Not only does it describe how the system works, but it gives you a seat-of-the-pants feeling for how complicated it's getting. As we expected, this is a basic superhet problem, like the AM radio of Example 13.6, except we have to get the LO from somewhere. Since we can use the 50–100 MHz AOD drive to get our frequency reference, we have two basic choices: use mixers and a frequency doubler, or a PLL with a frequency offset and a $\div 2$ counter (see Section 15.6.2). The direct method of Figure 15.1a is fast and has no PLL acquisition funnies, but the PLL of Figure 15.1b is simple and has far fewer spurs, because we can filter its control voltage vigorously (the only way spurs can get into the oscillator output is by coming in on the control voltage line, a key PLL advantage in frequency synthesis). A 50 kHz loop bandwidth will get us fast acquisition and a settling time of less than 100 μ s, which is fine. Let's draw both (Figure 15.1).

We need at least a few cycles of the carrier to get a decent envelope. A rule of thumb is that 4 is the minimum, 10 is comfortable if the timing constraints are loose, but the tighter the timing the more cycles you need. The input pulse will have at least 100 cycles, so any IF above 10 MHz or so will be fine. We could use a 70 MHz TV IF filter, which avoids the RF and LO frequencies, but the 50–100 MHz reference input crosses that IF, which will lead to nasty spurs since mixers have only so much RF rejection. Thus we're better off staying below 50 MHz. Let's choose 21.4 MHz, which is a standard frequency for wideband IFs and leaves some room for the filter to roll off.

The direct method needs three mixers, two oscillators, two filters, and a doubler, which is on the complicated side. It will also generate a weak spur exactly at 21.4 MHz ($2f_{LO1} - f_{LO2}$), which will get into our IF eventually unless we filter really well between the first mixer and the doubler. On balance, the PLL method is superior, but harder to get chips for; the highly integrated amp/mixer/oscillator chips almost always have the VCO

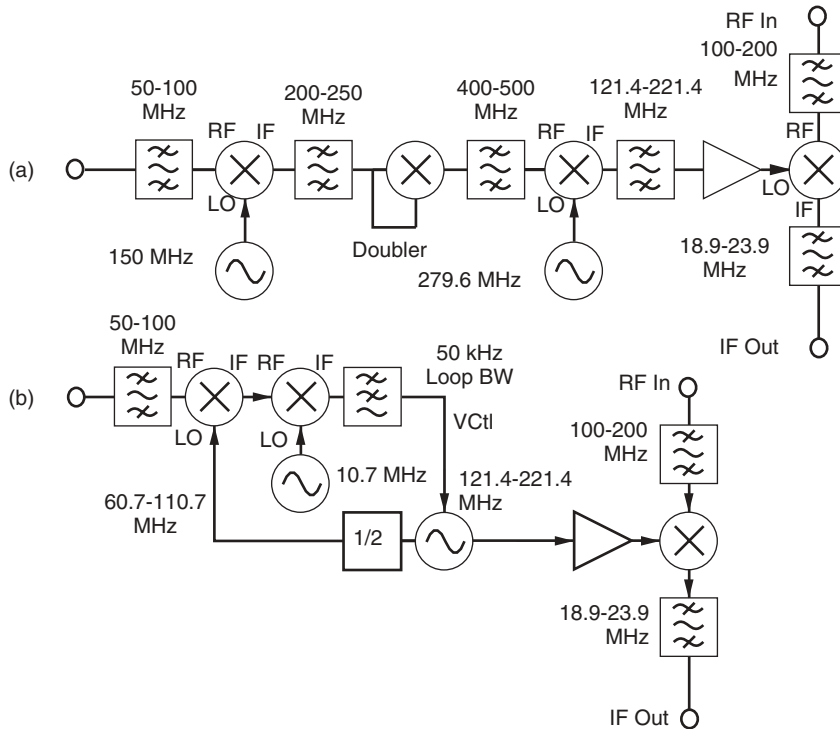


Figure 15.1. Pulse detection signal processing: (a) direct synthesis and (b) offset phase-locked loop.

connected right to the mixer, which we can't tolerate here. We could change it around to frequency-double the 50–100 MHz and then use a PLL to offset that, but that would require a tracking filter because the frequency range is a whole octave. What we'll do is to use an offset PLL at 1:1, with a doubler afterwards, as shown in Figure 15.2. The advantage of this is that as long as the offset frequency is well outside the PLL's loop bandwidth, and the VCO tuning range is restricted so that it can't lock onto the wrong sideband, we won't have any spurious problems. Figure 15.2 also shows the detection strategy: a single LO chain, with two mixer/detectors; one has a 30 dB attenuator in front of it. We sum the RSSI outputs from the two sections and take our AM output from the attenuated one. For a single digitizer, we can just sum them all, since the result will be monotonic and the (fairly serious) temperature drift will be taken out by the end-to-end calibration.

Aside: Sneaky Sideband Selection. At the offset frequency, where the PD operates, the upper sideband corresponds to a phase of 0 and the lower sideband to $\pm\pi$, due to the phase subtraction. This gives us another way to select sidebands. The idea is to use a phase–frequency detector like the one in a 74HC4046, which ideally has a sawtooth $V(\phi)$ characteristic. Although propagation delays prevent the step at $\pm\pi$ from being infinitely steep, its slope (and hence the loop gain) is a good 500 times larger than at 0. Jacking up a PLL's loop gain by over 50 dB is a sure-fire way to make the loop unstable, so the offset loop is unable to lock up at the wrong sideband. Thus we can choose sidebands by

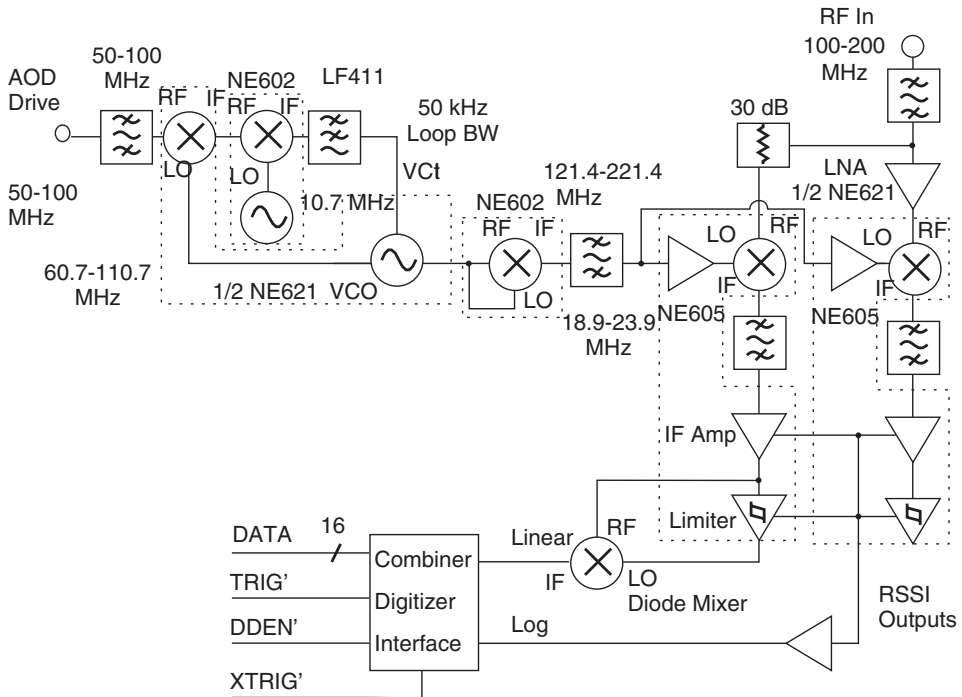


Figure 15.2. Final signal processing block diagram: a PLL similar to that in Figure 15.1 generates the LO, and two parallel single-conversion superhet receivers handle the wide dynamic range. The logarithmic outputs are summed, and the large-signal channel produces a linear output as well, which will satisfy the amplitude accuracy spec once an end-to-end calibration is done.

changing the sign of the loop gain, for example, with an exclusive-OR gate or a ± 1 -gain amplifier—a pretty useful trick, especially in a loop with a small offset where the other method is difficult.

15.2.4 Getting the Gains Right

When we get down to an actual schematic with resistors, capacitors, and pin numbers, we'll need to make sure that the signal levels throughout the system are correct. In general, this will require a sprinkle of amplifiers and attenuators in various places. The radio ICs we're using have lots of gain in all their functional blocks, so we'll ignore the problem since we can put in the required attenuation very easily at the end of the design.

15.2.5 Error Budget

Once we have a block diagram, we can start apportioning the pain between the functional blocks. The way to do this depends on your application. If you're building a high-volume commercial product, start with the data sheets of the cheapest jellybean parts that look like they might do the job. Use their guaranteed worst-case specs as a starting point, and go to more expensive parts only when driven to it. This approach takes more engineering talent but reduces the per-unit cost (it's also the shortest road to wizardhood). In a research

TABLE 15.1. Bug Zapper Signal Levels^a

Position	Channel 1					Channel 2			
	Floor (dBm)	Min Sig	Max Sig	S/N _{Min} (dB)	CC (mW ⁻¹)	Floor (dBm)	Min Sig (dBm)	Max Sig (dBm)	CC (mW ⁻¹)
Input	-92.2	-64.	+10	28.2	0	-92.2	-64.0	+10.0	0
Filter	-93.0	-64.8	+9.2	28.2	0	-93.0	-64.8	+9.2	0
30 dB Pad	-107.0	-94.8	-20.8	12.2	0				
LNA out						-75.0	-46.8	(-7.0)	2.0×10^3
Mixer in	-102.5	-94.8	-20.8	7.7	0.001	-75.0	-46.8	(-7.0)	
Mixer out	-87.5	-79.8	-5.6	7.7		-60.0	-31.8	(0.0)	1.6×10^3
<i>Total compressibility coefficient</i>					<i>0.001</i>	<i>3.6×10^3</i>			

^adBm unless noted; levels unreachable due to saturation are in parentheses.

lab, we mostly use what's in the drawer, or pick some fairly gold-plated parts that we're sure will work and which we can get in a hurry. Since no home should be without a bug zapper, Table 15.1 was calculated with jellybean parts in mind. (A real analog wizard could do most of this with discretes for less money.)

The draft design uses only six ICs, which isn't bad, and they're inexpensive FM radio chips. The BGA2011 LNA (low noise amplifier) has a noise figure of 1.6 dB, which means that the noise floor is degraded only a tiny amount. As usual, the mixer's dynamic range is the lowest in the string—it has a 4.5 dB noise figure and -13 dB input IP₃ (it's easier to make something linear than to make it be nonlinear in just the right way). We're keeping the LNA gain lowish to take maximum advantage of the mixer's dynamic range. Putting a noiseless 10 dB amp ahead of the mixer is as good as reducing the mixer noise by 10 dB, except that it reduces the system's input IP₃ by 10 dB as well—so it's easy to get too much of a good thing.

To avoid losing too much dynamic range on the high side, we cascade stages so that their third-order intercept points grow at least as fast as the accumulated gain; other things being equal, a 20 dB amplifier should have an output-referred IP₃ at least 20 dB above that of the previous stage. This is fairly permissive, because we can lose 3 dB of dynamic range per stage. We always lose something, but this rule at least avoids catastrophes. Low noise and especially high power amplifiers generally can't do this, but they're often needed for best system performance; as elsewhere in this book, rules are intended as a place to start, not as a substitute for thought.

Aside: Cost-Effective Error Budgets. Remember from Section 13.5.4 that the compressibility coefficients add, so scaling the IP₃s along with the stage gain means that each stage will contribute the same amount of compression—each will degrade linearity the same amount. This makes sense only if they cost the same. LNAs and power amps are the most expensive in general—we've paid for their performance, so we don't want to degrade it with cheap and cheesy intermediate stages (see Sections 1.7.1 and 18.8).

We assume a 6 MHz system noise bandwidth, a mixer with 4.5 dB noise figure, 15 dB (SSB) conversion gain and an output $P_{1\text{dB}}$ of 0 dBm, and an LNA with 1.5 dB NF, 18 dB gain, and -14 dBm $P_{1\text{dB}}$. The right-hand column in Table 15.1, marked CC, is the compressibility coefficient (see Section 13.5.4), which allows us to calculate that the input 1 dB compression point of channel 2 is $1/(3.6 \times 10^3)$ mW or -35.4 dBm.

Channel 1 (with the pad) has a calculated $P_{1\text{ dB}}$ of +30 dBm, and so is just beginning to compress at $P_{\text{in}} = +10$ dBm [from (13.21), $\delta = 0.23 \times (10\text{ mW})/(1000\text{ mW}) = 0.23\%$] and maintains the required 51 dB SNR down to $P_{\text{in}} = -20.5$ dBm, which far exceeds the 0.2% of full scale requirement (at least from the noise point of view). Due to gain variations, it would probably be worth spending some of that extra thirty-odd decibels to increase the headroom, perhaps going to a 40 dB pad. Going too far will make the circuit more vulnerable to pickup. Channel 2 (with the LNA) maintains its required 10 dB SNR (for the ± 1 dB standard deviation) down to $P_{\text{in}} = -83$ dBm, 19 dB better than spec. Incidentally, none of these ICs is more than \$3.

The main thing in the detectors (the two SA605s at right) is to make sure that the RSSI outputs are not just detecting their own noise, and that there isn't a flat spot between the end of the linear range and the beginning of the logarithmic one. Making sure of that will probably require a few voltage dividers (or another two digitizer channels and a bit of software), but it's a straightforward exercise with a few pots, an ohmmeter, and an oscilloscope. In a moderately complex signal processor like this, where the conversion gains, intercept points, and other features (such as the output duty cycle and DC offset of a severely overdriven limiter) are not well specified, a good prototype is absolutely essential, so get five or so boards to hack up, besides however many you anticipate needing for the system prototype and spares. A ring-down calibrator (Section 15.7.2) will help by making the desired logarithmic characteristic into a straight line.

15.2.6 Judicious Gold-Plating

Slightly more expensive parts (e.g., Mini Circuits mixers and cascadable amplifiers) offer far better controlled gains and intercept points than these ICs. Diode mixers are also very *strong*, that is, have high intercept points, low noise, and consequently very high dynamic range; on the other hand, they need extra gain stages to provide high level LO injection and to overcome their conversion loss.

Here, where we have an online calibration, logarithmic detection, and enough dynamic range that there's room for some sloppiness, we can get away with cheap radio ICs. In a really long string of signal processing components, you'll probably want to use those better quality mixers and amplifiers rather than the IC versions, or failing that, put a bunch of DAC-controlled attenuators in the signal path and run the calibration several times with different settings to estimate what the gains and intercept points of the different stages are. That level of self-calibration probably isn't worthwhile unless you're building a lot of systems, and at that point it's worth seriously considering sampling the IF directly and doing the rest in software or in an FPGA.

15.2.7 Interface Design

We used a very simple problem description in Section 15.2.1 because there was just one designer. If more people are involved, you'll need to prespecify all the interfaces between subsystems: the protocols, data formats, timings, logic levels, power supplies, and so forth. Do this in exact detail and spill some sweat over it; more system integration nightmares come from mistakes here than you'd easily believe (Mars Climate Orbiter wasn't an isolated case). You will be amazed at the number of different ways a given spec can be interpreted too, if you don't keep talking with the designers of all the pieces you have to interface with. Contracts are very precise documents too, but lawyers spend a lot of time arguing about them even so. Neighborly cooperation is the way to go.

15.3 PERFECTION

It's often said that nothing is perfect, but components and circuits do some jobs effortlessly. These perfections come in two flavors: near-ideal component properties, such as the zero input current of MOSFETs, and conservation laws (e.g., charge conservation). Electronic designers are fond of saying "well, that may be true in *theory*, but..." Don't fall for that one; there's gold to be mined out of things you know have to be perfect. Here's a short list.

Fundamental Laws

1. Things wired in parallel see exactly the same voltage, if the loops are small enough (Kirchhoff's voltage law, i.e., Faraday's law).
2. Two things in series see exactly the same current, except when there's significant radiation or capacitive coupling (Kirchhoff's current law, i.e., charge conservation: $\nabla \cdot \mathbf{J} = 0$).
3. A passive component with no voltage swing across it will draw no AC current (Ohm's law and the linearity of Maxwell's equations).
4. Devices with no dissipation exhibit no thermal noise (fluctuation–dissipation theorem).

Component Properties

5. MOSFETs draw no gate current whatever.
6. Metal film resistors contribute only Johnson noise (not shot or $1/f$).
7. A BJT's collector is a good current source; a cascoded BJT's collector is a near-perfect current source.
8. Feedback can improve the frequency response but not the SNR.
9. An RF filter can completely remove spurious products in a fixed-frequency AC measurement.

Let's look at a couple of related points in some detail.

1. *Flying capacitors can add and subtract perfectly.* Good maps have a scale of distance down in one corner, which is nice but not enough by itself, because the two points whose separation we care about never lie on the scale. Accordingly, we use chart dividers (like sharp-pointed calipers) to transfer a replica of the line segment down to the scale, and read it there.

Voltage measurements are the same way; usually neither terminal of the voltage source is exactly at the amplifier's idea of ground. Instrumentation amps transfer the voltage difference to ground with precision resistor strings and some op amp circuit hacks, but we can use the transfer-a-replica trick here too. A good quality capacitor functions the same way as the chart dividers: we charge it up across the voltage source, then disconnect it and reconnect it across the input of the amplifier. We do it with CMOS analog switches, and it really works; CMRs of 120 dB are typical, even with $\pm 10\%$ capacitors. There are special circuits (e.g., the LT1043) for this, but if you don't mind a constant offset voltage caused by charge injection, you can use ordinary switches too, as in Figure 15.3. It is

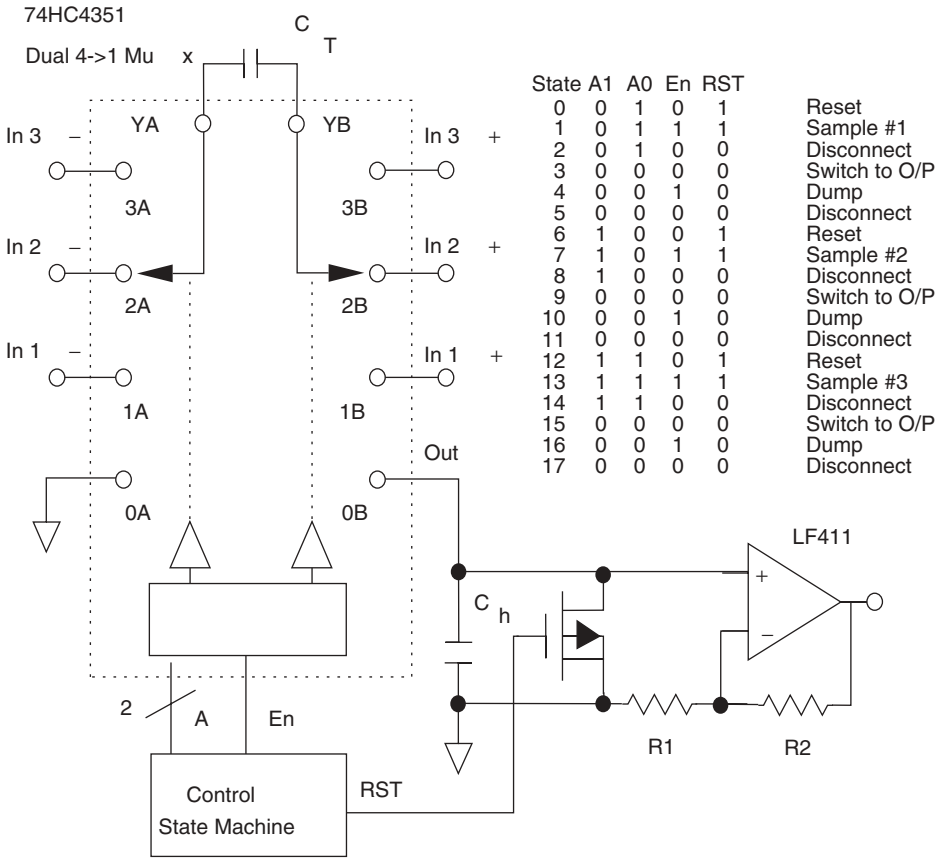


Figure 15.3. Three-way flying capacitor sampler made from a dual 4-into-1 analog multiplexer.

very important that the switches finish disconnecting before they reconnect to the other set of terminals—the *break-before-make* condition. Most inexpensive switches can't be relied on to do this, so we use a simple state machine such as a 4017 CMOS 1-of-10 counter and a few gates to do the sequencing, leaving the switches open-circuited for a whole clock cycle between disconnecting from one set and reconnecting to the other. In the state table in Figure 15.3, it may look tempting to eliminate disconnect states 2, 8, and 14, so that the state table can be implemented with 4 bits. You might be able to get away with it, but that depends on the inner workings of the 4351. The 4351 is not guaranteed to be break-before-make, so it may cause transient shorts between inputs, leading to hard-to-find errors. Building that sort of glitch into your circuit is a mistake, especially when an extra flip-flop will guarantee it doesn't happen.

2. *Independent noise sources are really uncorrelated.* There's another example in Section 17.11.6: the two-point correlation technique, where the power in a small AC signal is measured with two noisy voltmeters in parallel; their voltage noises are uncorrelated, so if the voltage measurements are $v_1(t)$ and $v_2(t)$, the signal power is $\langle v_1(t) \cdot v_2(t) \rangle / R$. Here's where the perfect component comes in: i_N is a real current that comes out on the input lead, so the current noises of the two voltmeters add, and you can't

separate that out by cross-correlating. Fortunately, we can easily get MOSFETs with femtoampere input currents and noises of a few hundred electrons/s/Hz^{1/2}, so we just use two MOSFET buffers in front of the voltmeters, and we're in business. (That kind of MOSFET doesn't have gate protection diodes, so we have to be careful not to blow them up.)

15.4 FEEDBACK LOOPS

Negative feedback is one of those concepts that is so deeply imbedded in all our technologies that it may come as a shock to realize that the feedback amplifier was invented as late as 1927.

There are earlier examples, for example, James Watt's flying ball governor for steam engine speed, or the tank valve of a flush toilet, but they are rare and isolated by comparison. The inventor, Harold S. Black of Bell Labs, had a great struggle to get it through the patent office. Why would anyone want to make an expensive amplifier with 40 dB more gain than it needed—two extra tubes in those days—and then just throw it away by closing the loop? Patentable inventions have to have utility, and that didn't seem so useful, it seems.

As Black probably explained to the patent examiner (in words of one syllable or less), the reason was that doing so gave an amplifier of unprecedented linearity and flat frequency response. That was good enough then, and it's good enough now. A cascade of N identical amplifiers whose gain ripple is m dB exhibits mN dB ripple overall. This was requiring the telephone network to use unwieldy cascades of stagger-tuned amplifiers, so that the peaks of one would coincide with its neighbor's valleys, reducing the ripple. Furthermore, the distortion due to nonlinearity accumulated regardless of such maneuvers. (Readers using optical amplifiers as fiber repeaters will recognize both problems.)

Given an amplifier like an LF356, whose gain varies from 100 dB near DC to 1 at 5 MHz, subject to big swings with temperature, supply voltage, and common-mode voltage, negative feedback using two resistors will produce a 30 dB amplifier flat from DC to 150 kHz; the essence is that when there's lots of excess gain, and certain stability criteria are met, a feedback system's behavior is governed entirely by the feedback network, which is made of extremely well-behaved components such as metal film resistors.

Our interest in this section is not so much learning how to build op amp circuits. You can read about the intuitive method in Horowitz and Hill and the rigorous method in Dostal. What we're aiming at here is learning how to make more complicated feedback systems that use op amp circuits as loop amplifiers and filters—things like temperature controllers, feedback laser stabilizers, and phase-locked loops. Accordingly, we won't spend a lot of time on what happens up near the op amp's unity gain crossover f_T , because we almost never build complicated loops that fast. There's more on pushing those speed limits in Chapter 18.

15.4.1 Feedback Amplifier Theory and Frequency Compensation

A feedback network is shown schematically in Figure 15.4. For now, we'll assume that the components are linear and time invariant, and the amplifier's output impedance is zero and its input impedance infinite; as usual, we can fix things up afterwards when

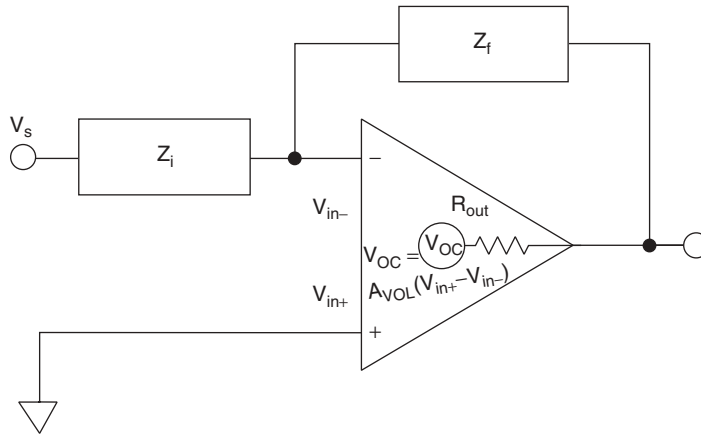


Figure 15.4. A feedback amplifier.

they aren't, quite. The amplifier has a frequency-dependent gain $A_{VOL}(f)$, and the input and feedback impedances are Z_i and Z_f . Since this gain is always spoken of in the frequency domain, it's usually called the transfer function, which is what it really is. All amplifiers are differential, but some are better than others. This one produces an output voltage

$$V_o(f) = A_{VOL}(f)(V_{IN+} - V_{IN-}). \tag{15.2}$$

The feedback network forces V_i to obey

$$V_i = V_s \frac{Z_f}{Z_i + Z_f} + V_o \frac{Z_i}{Z_i + Z_f}. \tag{15.3}$$

Requiring that both of these hold together, we get the closed-loop gain A_{VCL} of the whole circuit,

$$A_{VCL-}(f) \equiv \frac{V_o}{V_s} = \frac{-Z_f}{Z_i} \left(\frac{\beta A_{VOL}}{\beta A_{VOL} + 1} \right), \tag{15.4}$$

where the *feedback fraction* β is the gain from the output back to the inverting input,

$$\beta \equiv \frac{\partial V_{in-}}{\partial V_o} = \frac{Z_i}{Z_i + Z_f}. \tag{15.5}$$

If we use the noninverting pin as the input, we get the noninverting closed-loop gain,

$$A_{VCL+}(f) = \left(1 + \frac{Z_f}{Z_i} \right) \left(\frac{\beta A_{VOL}}{\beta A_{VOL} + 1} \right). \tag{15.6}$$

Since the amplifier is linear, and there's nothing special about ground, the total output voltage is the sum of the two contributions,

$$V_o = A_{VCL+}V_+ + A_{VCL-}V_-. \tag{15.7}$$

If $|\beta A_{VOL}| \gg 1$, these simplify to

$$A_{VCL-} \approx -\frac{Z_f}{Z_i}, \quad |\beta A_{VOL}| \gg 1, \quad (15.8)$$

$$A_{VCL+} \approx 1 + \frac{Z_f}{Z_i}, \quad |\beta A_{VOL}| \gg 1, \quad (15.9)$$

which is what we mean by saying that the feedback network takes over control of the amplifier's performance. Departures from this ideal condition are expressed by the error factor, the last term in parentheses in (15.4) and (15.6), and are clearly controlled by the *loop gain* $A_{VL} = \beta A_{VOL}$. There's nothing special about voltage dividers—transformers work just as well at AC, and there are more complicated networks too. They're cheap and convenient, and they stabilize the DC bias as well as the gain, but it's β that matters.

The amplifier A_1 doesn't have to be an op amp, of course. In a complicated system such as an AGC loop, where the gain of a receiver's IF amplifier is controlled to keep the DC voltage from the detector at a constant value, A_1 is the IF amplifier with its input signal, the IF filters, detector, and AGC error amplifier; any and all of these may have their own local feedback loops, too. The open-loop gain is then the product of all the gains of the string. Its frequency dependence of A_{VOL} is a bit complicated in that case, but finding it is a straightforward matter of calculation or measurement, after which these equations apply.

15.4.2 Loop Gain

The response of a feedback loop to an external disturbance, for example, somebody sneaking in and putting a battery in series with its output (inside the loop), is also governed by the loop gain, which is the amount of gain thrown away when we close the loop. If we put a signal source in series with the amplifier output, we find that its effect is reduced by a factor of A_{VL} . Since we assume that the amplifier is stable when the loop is open, and $\beta \neq 0$, there is no opportunity for instability unless $A_{VL} \approx 1$, when the denominator can go through 0.

The open-loop gain has to roll off somewhere below the frequency of daylight, however, so the high gain limit is not the only interesting place. We saw in Chapter 13 that a stable network has no poles in the lower half f plane, so the denominator of (15.4) has to be nonzero there.

The simplest way to be sure we satisfy the stability condition is to ensure that $A_{VOL}\beta$ has less than 180° phase delay at the frequency where $|\beta A_{VOL}| = 1$. The difference between the actual phase delay and 180° is called the *phase margin* and shouldn't be allowed to get below 45° in most cases, because gain peaking, overshoot, and ringing get much worse for phase margins less than that. We do this by constructing a *Bode plot* like the one in Figure 15.5, which is a plot of the asymptotes to the open-loop transfer functions and (especially) A_{VL} on a log-log plot with phase on the other Y axis. In Chapter 18, we'll come back to this point in more detail.

15.4.3 Adding Poles and Zeros

If you're making a fast amplifier, then all you usually need are resistors R_i and R_f , and perhaps a small capacitor C_f across R_f to cancel the effects of the input capacitance of the op amp. The manufacturer's data sheet will tell you what you need to know.

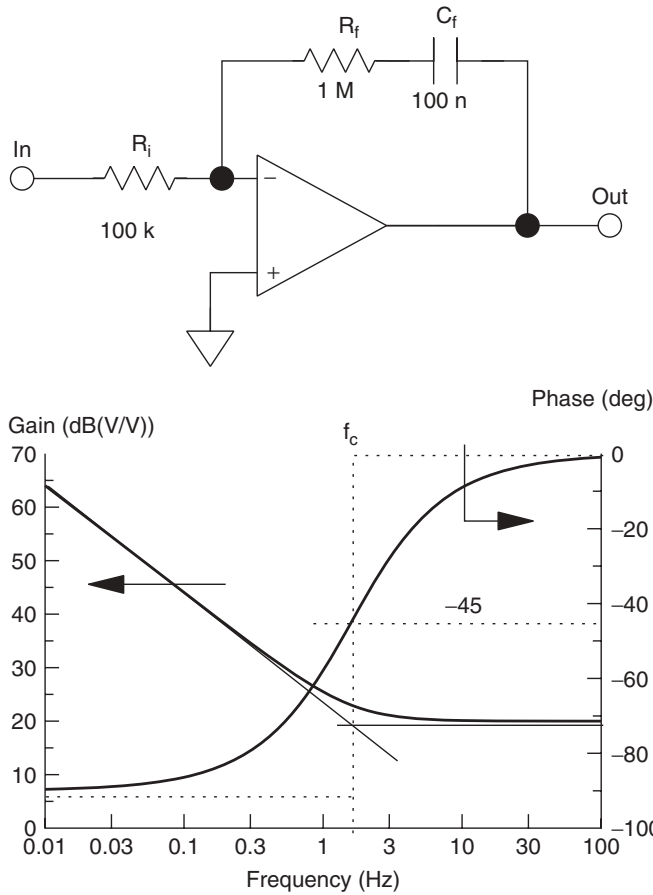


Figure 15.5. Op amp lead-lag network and its transfer function. Due to C_F , the amplifier's bias is unstable, so this circuit fragment has to be wrapped in another feedback loop.

Right now, as very often, we're building some much more complicated thing such as a motor control loop. The loop bandwidth is then usually far below f_T of the op amp, so the high gain equations (15.8) and (15.9) apply.

At that point, it becomes duck soup to put poles and zeros just where you like, because the zeros of the transfer function are at poles of Z_i and zeros of Z_f , while its poles are at poles of Z_f and zeros of Z_i . We chuck in a few series or parallel RC networks, and we can make our loop gain do nearly anything we want.

Example 15.1: Lead-Lag Network. The op amp circuit of Figure 15.5 has a *lead-lag* characteristic. At low frequency, series combination $R_f + C_f$ looks capacitive, so the circuit looks like an integrator. At f_c the corner frequency of R_f and C_f , it starts looking like a flat 20 dB amplifier with resistive feedback. When used as part of a complicated loop such as a PLL, this allows a two-pole (-40 dB/decade, 180° phase lag) rolloff where the loop gain (of the big loop) remains high, changing to a one-pole rolloff (-20 dB/decade, 90° phase lag) before the loop gain gets to 0 dB, to preserve a good phase margin. This allows us to use a much higher low frequency gain, for much

lower static errors, at the price of a longer settling time. We'll use a similar trick in Section 18.4.1 to get more bandwidth from a transimpedance amp.

Aside: Nonlinear Instability. High order feedback networks can be stable with $|A_{VL}| > 1$ and huge phase shifts near DC (270°). The problem is that when a nonlinearity occurs (e.g., on signal peaks or on power-up), the nonlinearity will reduce $|A_{VL}|$ and the stupid thing will oscillate. Good overload recovery circuitry is a vital necessity there.

15.4.4 Integrating Loops

Some circuit elements have an integrating response already built into them; the canonical example is a PLL's VCO and phase detector, where the control voltage sets the frequency, which is $d\phi/dt$, but the detected output is ϕ . From the formula (1.40) for the Fourier transform of an integral, the feedback fraction β is then proportional to $1/f$, instead of being a constant ratio as we're used to from op amps.

A very tightly coupled heater and thermal mass are a close approximation under some circumstances. A given heater power \dot{Q} will eventually raise the sensor temperature by $\dot{Q}R_{TH}$ degrees above ambient (where R_{TH} is the thermal resistance), so the response is constant at large t (near DC). The heat takes a while to diffuse from the heating element to the sensor, and this thermal diffusion dominates the response at short times (i.e., high frequencies). If the sensor is intelligently placed, there's a region in the middle where the transfer function goes as $\dot{Q}/(M_{TH}f)$, where M_{TH} is the thermal mass. Another example is a current-sensing motor speed controller, where the motor current controls the torque Γ , which controls the angular acceleration $\dot{\Omega} = \Gamma/I$, where I is the moment of inertia.

In that sort of region, increasing the loop gain moves the loop corner frequency f_c up, and so increases the speed of the loop without greatly affecting the shape of its transient response. You can play this game up to the point where the excess phase shift due to other poles, to thermal diffusion, or to losses in the motor becomes important.

Integrating loops such as PLLs are intrinsically more difficult to compensate, because your loop filter has only 45° to work with; you really have to calculate these ones carefully, because it's virtually impossible to get a good result otherwise. When you think you've finished, apply Pease's Principle: Bang On It. Do it electrically, by looking at the turn-on transient under all sorts of load conditions, from no load to above maximum; apply an adjustable-frequency square wave to the control voltages, and look for overshoot or odd peaks and nulls in the response curve. Put a brake on the shaft, cold spray on the temperature sensor.

If the loop gain is nonlinear, as it is in heaters and usually in PLLs, test it over the entire range of gains, and put in a safety factor for unit-to-unit variations.

15.4.5 Settling and Windup

Besides loop stability, we usually care about the transient response. Depending on the application, a step response that overshoots may be harmless; it may cause measurement errors, as in a front end amplifier; or it may cause destruction of the whole system, as in a heater or a motor controller. Systems with slow transducers, such as servomotors, are prone to *windup*. At power-up, or when a command to move arrives, there is a large error signal. The control loop starts ramping the motor speed, which increases until the

position encoder reaches the set point. Unfortunately, due to the motor's inertia, it goes flying past the set point and has to be brought back again—at which point it winds up again, and has to be brought back again, and so on. It's basically a disease of slow transducers with integrating control loops. There are two main ways of dealing with this: adding some speed feedback (i.e., a derivative term) into the error signal, and using some nonlinear element for windup control, such as an analog switch that freezes the integrator contribution when the error signal is large.

Good windup control is important for another reason: the long-time settling behavior of any amplifier is dominated by its lowest frequency open-loop pole or zero, so to get good settling, we don't want the the open-loop transfer function to have corners. Integral-only loops settle to high accuracy faster than if there's a lead-lag or derivative network. That's where a well-designed nonlinear network can be a big help—it controls the windup during slew, but goes away during settling. Tim Westcott's book on control systems[†] has lots of good advice on windup control among many other things.

15.4.6 Speedup Tricks

After seeing how long it can take a feedback loop to recover from saturation when there's a slow integrating time constant in it, the temptation to limit the range of the integrating part is nearly irresistible—we inevitably try shoving in some back to back diodes or perhaps a cleverly connected FET across the capacitor, across the resistors, wherever they look nice. Done right, anti-windup and speedup tricks can help a lot, but they won't do anything good unless you're using some simulations of how the loop will behave. One exception is the so-called baby-scale loop, a slow one-pole, unity gain noninverting lowpass filter, with high and low value input resistors controlled by an analog switch and a window comparator; whenever $0.95V_{in} < V_{out} < 1.05V_{in}$, the high value resistor is in, for wriggle rejection, and otherwise the low value one is used, to help the slew rate, and keep the babies from getting cold waiting. This is pretty problem-free in first-order loops, but not so easy otherwise.

15.4.7 Output Loading

There are two major effects of nonzero output impedance in the op amp: feedthrough and sensitivity to output loading. Feedthrough is signal current going through the feedback network to the output and causing a voltage drop in the output resistance that appears as output voltage. A rail-to-rail op amp with a $1000\ \Omega R_O$, and a feedback network made up of a $1\text{k}\ R_i$ and $10\text{k}\ R_f$ will have a small-signal gain of $1000/12,000$ ($-21.5\ \text{dB}$) even with the output railed, due to finite output impedance alone. This happens often in op amp based analog switching applications, so it isn't an imaginary problem. If the output is connected to some noisy place like a motor, junk can arrive at the input from the output, which is sometimes pretty serious, especially if it can get into other circuit nodes from there.

A finite R_O makes the insertion phase of the amplifier depend on the reactance of the load; a capacitor on the output puts a pole in the open-loop transfer function at $f = 1/(2\pi R_O C_L)$, which adds phase shift that can destabilize the amplifier. The classical way to fix this is to isolate the amp from the capacitance with a series resistor, with

[†]Tim Westcott, *Applied Control Theory for Embedded Systems*. Newnes, 2006.




Figure 15.6. Modified current source for low impedance loads.

feedback taken directly from the op amp output. A slightly modified version of this uses an RC network to take low frequency feedback from the load and high frequency feedback from the amplifier output, which will improve the accuracy for low impedance resistive loads, at the price of significantly poorer settling time.

For cable capacitance, we can use a beefy op amp or buffer, and terminate the cable in its characteristic impedance—either series terminated at the buffer’s output or (if you can stand the power dissipation) shunt terminated at the cable’s output. We can even be model citizens and do both, but then must be prepared for a 6 dB loss as well as high dissipation.

Another strategy, in cases where some loss of speed is acceptable, is to use an LM7332 or equivalent op amp that can drive lots of capacitance. A slightly unbalanced op amp current source (see Figure 15.6) can be made to have a 50 Ω output impedance without needing a 50 Ω series resistor, which helps with the signal loss. Watch the high frequency behavior of this circuit, though, and don’t make the series resistor *too* small.

As we saw in Chapter 14, rail-to-rail output op amps have very high open-loop output impedances, and so are far more susceptible to instability due to output loading.

Aside: Oscillating Integrators. Given the 90° phase lead from the feedback network in an integrator, it takes real talent to make one oscillate, but it’s possible. Use an op amp with a high output resistance, a low R_i , and a huge C_f , and you can make an integrator that rings like a bell.

15.5 SIGNAL DETECTORS

We spent a fair amount of time on the basic principles of signal detection in Chapter 13, and one thing we concluded was that detectors are much of a muchness at large SNR, but differ substantially at low SNR, which is what limits our measurement performance. Here we’ll delve a bit deeper into how to make good ones.

15.5.1 AM Detection

We’ve already encountered the two basic types of AM detector, namely, rectifying and synchronous, in Section 13.9.2.

The simplest AM detector is a half-wave rectifier made out of a diode with an RC lowpass on its output, but as we saw there, it doesn’t work too well, because diodes

are not great voltage switches. For small signals, the output voltage is roughly proportional to the input power (the so-called square-law region), but then as the RF amplitude rises, it gradually changes to being proportional to the RF voltage, as you'd expect from a rectifier. The level at which this happens depends on loading, temperature, and frequency. As a linear detector, the half-wave rectifier has about a 20 dB dynamic range. This may be adequate for some purposes, for example, AGC detectors, where a servo loop controls the IF amplifier gain to keep the detected voltage constant.

Backing up to what we know about diodes, they are nearly perfect current switches. It thus might occur to us to drive the diode detector with a current instead, as in Figure 15.7; any current imbalance will cause the collector of Q_1 to swing as far as it has to, to send that current through one of the detector diodes. This looks like a complete solution, but of course it isn't—the slewing becomes slow at small signals, and the collector impedance isn't infinite. Still, it's good for a 40 dB dynamic range up to 50 MHz or so, and maybe 60 dB at 10 MHz, which is a great improvement.

15.5.2 Emitter Detector

Figure 15.8 shows an emitter detector, which is nothing more than an emitter follower with a slow RC in series with the emitter. It functions as a half-wave detector that doesn't load the circuit the way a diode does. This detector is mainly useful when two or more of them are strung together, to make full-wave emitter detectors and DLVAs. For small signals, where the transistor stays conducting, its rise and fall times are $r_E C$, while for a big negative excursion, the transistor cuts off and leaves only the slower RC falloff.

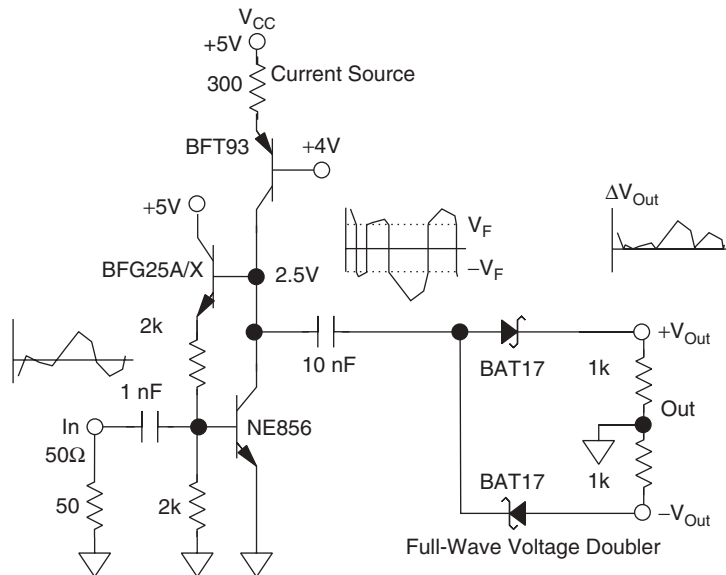


Figure 15.7. Linear rectifying detector. The common emitter amplifier and current source load function as a bipolar current source, so that its output voltage skips across the diodes' dead zones to source current to the load resistors.




Figure 15.8. The emitter detector works like a combination of a half-wave diode and an emitter follower; the rise time and the small-signal fall time are set by r_E and the large-signal fall time by R . (Watch out for oscillations if the source impedance is very low.)

15.5.3 Synchronous Detectors

A better approach is the synchronous detector, where you put the signal into a balanced mixer whose LO has the same frequency and phase as the signal. If you're providing the carrier signal, this is how you should do it almost every time (it's how we did it in the bug zapper example). If all you have is the input signal, you can produce the LO by shoving the signal into a limiter or a comparator. This works pretty well at all signal-to-noise ratios, providing the signal is always large enough for good limiting. At low SNR you wind up detecting the noise as well as the signal, because of course the circuit cannot distinguish the two. AM–PM conversion can be controlled by slight adjustment of the phase, so that the slight slowing of the limiter for large signals doesn't move the phase too far from 0° .

Limiters with AC coupling between stages work better than barefoot comparator ICs, because DC offsets and possible comparator oscillations become less troublesome. Comparator ICs usually need some sort of hysteresis to prevent oscillation at low signal levels, and this will seriously degrade the small-signal operation of the detector.

15.5.4 High Performance Envelope Detection

We encountered vector (I and Q) modulation in Section 13.8.7 as a way to convert a signal to baseband without folding the sidebands. In detection problems, it allows us to use a fixed oscillator to drive the phase-sensitive detector instead of having to track the phase of the incoming signal. The problem is then to convert the resulting quadrature pair into amplitude and possibly phase too. One technique is to digitize both and use a lookup table, which is a good way to proceed if you need phase information, but it isn't the only possibility.

In radar, they use a sleazy trick to get the envelope magnitude from I and Q , as shown in Figure 15.9: take absolute values, then pick the largest out of $|I|$, $|Q|$, $(|I| + |Q|)/\sqrt{2}$. This is equivalent to using eight phases, spaced evenly at 45° ; the worst-case scallop loss is only 0.9 dB. Using more phases works even better (you win quadratically), at the cost of one op amp per time: using $|I|$, $|Q|$, $|I| \sin(\pi/6) + |Q| \cos(\pi/6)$ and the same for $\pi/3$ (leaving out the $\pi/4$ one) gives a scallop loss of only 0.3 dB.

Using absolute values saves parts, because it confines us to the first quadrant. Those rectifying amplifiers limit our speed, though. By omitting them, we can go much faster,




Figure 15.9. High performance envelope detection. Here three phases are shown; the six transistors correspond to positive and negative half-cycles of all three.

tens to hundreds of megahertz. We use resistor networks and an inverting amp or two to generate the phases independently all round the circle: I , $-I$, Q , $-Q$, $(I + Q)/\sqrt{2}$, $(I - Q)/\sqrt{2}$, and so on, and put an emitter follower on each one. The phase selection is done by wiring all the emitters together, wire-OR fashion, to a single current source load. Although this scheme shares the small-signal nonlinearity of the emitter detector (see Section 15.5.2), this actually helps a bit at larger signals, since it occurs only when two phases are producing nearly the same output, and it will thus help to fill in the scallops.

15.5.5 Pulse Detection

The classical peak detector is an op amp with a diode in its feedback loop and an RC lowpass network on its output, with feedback taken from the hold capacitor (Figure 15.10 shows an improved version). It's a bit like a track/hold, except not as good. Peak detectors are slow, inaccurate, or both, due to overshoot and slew limiting of the amplifier used. As shown in Figure 15.11, a peak detector that overshoots is frozen forever at the peak of the overshoot waveform, which depends on time, temperature, frequency compensation, and the fine details of the pulse shape. You can do a decent job on a pulse whose top is smooth and whose rise time is at least $1 \mu s$, provided that the range of pulse heights does not exceed 30 dB.




Figure 15.10. Peak detector.




Figure 15.11. Peak detector pathology.

Aside: Constant Fraction Triggering. Getting a low jitter trigger from pulses with variable amplitude is a pretty generic problem, especially with equivalent-time (stroboscopic) sampling. Unless you already have a really low jitter timing reference, you have to design around the problem.

A common error is to set the trigger voltage somewhere on the leading edge and hope for the best. The reason this doesn't work is that the leading edge isn't infinitely sharp, so if your amplitude uncertainty is $\langle dA \rangle$, the resulting timing jitter is

$$\langle dt \rangle \approx \langle dA \rangle / (dV/dt). \quad (15.10)$$

If the pulse is round-looking at all, this is pretty bad. The trick is shown in Figure 15.12: without slowing down the edges too much, modify the pulse so that it has positive and negative lobes, and trigger on the zero crossing in the middle, which is independent of pulse height. We aren't quite done, because a trigger set at 0 V will keep triggering on noise. If you're using a scope with fancy triggering, you can tell it to look for a threshold crossing followed by a falling edge. If you're building it into an instrument, you can combine the zero crossing with an ordinary threshold signal, as in Figure 15.12b, which will reject the noise. The combination will give a stable, low jitter trigger signal more or less independent of amplitude variations. (You can also apply the threshold signal to the trigger comparator in analog, using different delays; also, there's nothing sacred about the coax stub—all you need is an impulse response with one positive and one negative lobe and fast edges.)

15.5.6 Gated Integrators

Most of the time, the shape of the pulse does not change much with its amplitude, so that the area under the pulse is a reasonable surrogate for the pulse height. Gated integrators

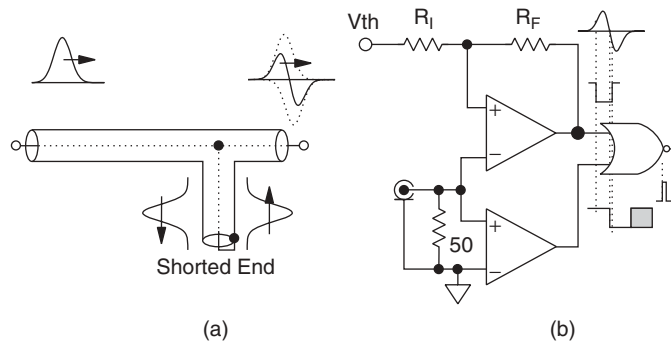


Figure 15.12. A constant-fraction trigger eliminates trigger jitter caused by pulse height variations: (a) a shorted coaxial stub produces an inverted copy of the input pulse, and the sum of the two produces a zero crossing somewhere near the peak of the input pulse; (b) two comparators and a NOR gate produce a stable trigger edge. Hysteresis provided by R_I and R_F ensures that the upper comparator remains active long enough to generate a good trigger pulse but not long enough to let through random junk from the lower one when the input has returned to 0.

are much easier to make than peak detectors. Because the integrator slews much more slowly than the peak detector, and can recover from overshoot rather than being frozen, gated integrators make much better pulse height detectors than peak detectors do. You do need a trigger, though, which as we just saw can be a nuisance. An excellent alternative is the *ping-pong gated integrator*: use two gated integrators running continuously at about a 75% duty cycle, but 180° out of phase (so that half the time, both integrators are active at once), which ought to guarantee that any signal pulse that shows up will be completely inside one of them. You can pick which output samples you want more or less at your leisure. Ping-ponging is useful in all sorts of places.

15.5.7 Peak Track/Hold

If you really must detect peak height, use a peak track/hold (T/H) instead of one of those ugly things with the diodes in its feedback loop. A peak track/hold is a track/hold whose sampling signal is controlled by a comparator, as in Figure 15.13. The comparator compares the input and output of the T/H, and as soon as the output drops below the input, it switches to hold mode. Ideally the comparator should be faster than the T/H, or sample the signal somewhere slightly upstream, so that errors caused by the signal delay are reduced. It's still hard to do better than 1% with this arrangement, and it will give the wrong answers systematically for pulses with long tails (slower than the T/H droop).

15.5.8 Perfect Rectifiers

The exponential dependence of $I(V)$ for a forward-biased diode leads to serious nonlinearity in simple rectifier–diode detectors. Because the conductivity of the diode depends on the drive level, low load impedances make them even worse. This effect is worse than you would expect, because the conduction occurs on the signal peaks, and this small duty cycle magnifies the effect. The spec for this is the video resistance, which is usually around 1–10 k Ω .

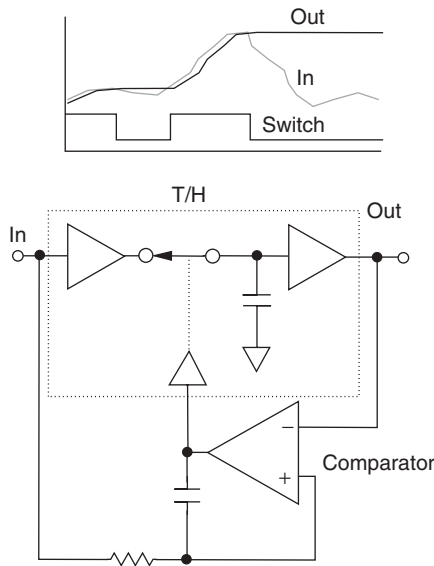


Figure 15.13. Peak track/hold circuit.

One approach to fixing it is to put it in a feedback loop, just as we did with the peak detector. This can be done well, or it can be done fast, but not both—the perfect rectifier problem has a lot in common with peak detection, since after all an AC waveform is just a string of closely spaced pulses. The fast approach is to use an emitter follower stage instead of the diode; several can be wire-ORed together to make a maximum detector, as we saw in Section 15.5.4. This is not a bad deal but does tend to have glitches caused by the finite slew rate; the output has to slew very rapidly across the forward voltages of the two diodes (or diode and transistor). They tend to be worse at small signal levels, where there is less overdrive available to make the output slew rapidly. Thus for signals of any speed, we're back to a quadratic problem.

You can also use the emitter follower with the op amp output connected to its base and the inverting input connected to a voltage divider between base and emitter. This is a diodeless version of the perfect rectifier plus catch diode trick. This works pretty well, a lot better than the catch diode version, but is also nonlinear at small signals because the output impedance of the emitter follower changes with its load current, so that the switchover between the two transistors is gradual at the level of 20 millivolts or so.

The closest approach to a *perfect* perfect rectifier is a switching detector (e.g., a Gilbert cell) driven by a pure current, and switched by a really high gain, fast limiter, whose gain is high enough to switch cleanly on the noise of the input in the absence of signal.

15.5.9 Logarithmic Detectors

Optical measurements can reach shot noise limited sensitivities on the order of 1 part in 10^8 intensity change in 1 second, so their dynamic range is often enormous. That 160 dB dynamic range cannot be accommodated by any A/D converter whatever; unless we can afford to throw away dynamic range, we have to use range switching or logarithmic detection schemes to fit it all into our ADC.

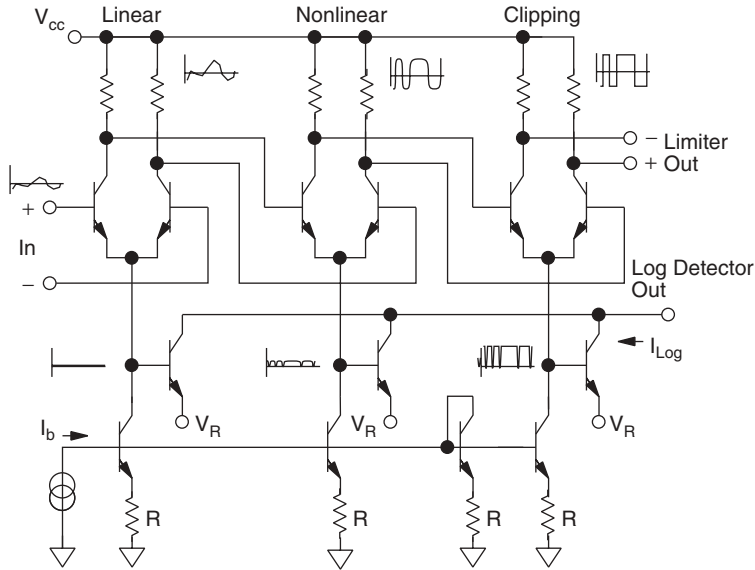


Figure 15.14. Successive detection logarithmic AM detector.

Logarithmic detection has several advantages: a level change of x dB translates into the same number of ADU anywhere in the dynamic range, the sensitivity of subsequent stages to interference and noise is greatly reduced, and the ease of achieving an accurate logarithmic characteristic makes them stable and predictable. Furthermore, we don't lose any data by being caught on the wrong range setting. The trade-offs are lower resolution at higher signal levels and moderate errors in the logarithmic characteristic (0.5–1 dB usually).

The classical method of logarithmic signal detection is *successive detection*, as shown in Figure 15.14. In this clever scheme, N amplifiers of equal gain A are cascaded, each with some sort of AM detector, whose outputs are summed. Here it's differential pairs, doubling as amplifiers and full-wave emitter detectors; as you can see, the DC voltage at the emitters of each stage goes up with the signal, until the *previous* stage starts clipping.

As the input signal level increases from 0, the last detector produces an output proportional to $A^N v_{in}$, until the last stage clips. When that happens, the last detector's output remains constant, and the output slope is proportional to $A^{N-1} v_{in}$. Every time the input level goes up by a factor A , the gain drops by a factor A , so that

$$\frac{\partial V_{out}}{\partial V_{in}} \propto \frac{1}{V_{in}}, \quad (15.11)$$

whose solution is proportional to $\log(V_{in})$. Purpose-built successive detection circuits are often known as detector-log video amplifiers, or DLVAs. These can be very fast and do not suffer from the linearity and speed limitations of diode detectors at low signal levels. There are also logarithmic amplifiers, made from bipolar transistor pairs, but these tend to be slow and have serious parametric effects such as signal-level-dependent slew rates and bandwidth.

You can get dedicated DLVA parts, such as the MAX4003 (which has a 45 dB range), or use FM IF ICs for consumer radio receivers as we did in the bug zapper. These use DLVAs for the signal level meter output (called RSSI, for received signal strength indicator). The meter outputs are often very useful in instruments, since they typically cover an 80 dB range with ± 1 dB accuracy, and the ICs cost only \$2 or so. There aren't as many of these as there used to be, but NXP still has several. "High speed" parts have op amps on the RSSI output, but the older ones just bring out the detected current. You can actually get quicker response from the "slow" current-output RSSIs, as fast as 40 ns if they're loaded with a low enough impedance (see Section 18.4.4). Besides that ± 1 dB error, these devices are normally not temperature compensated too well, so calibration and compensation will be required. We saw in Section 14.6.6 that BJT saturation is slow and badly behaved, so if you're designing your own, make sure that clipping occurs due to cutoff instead.

15.5.10 Phase-Sensitive Detectors

Diode bridge mixers are the phase-sensitive detectors of choice at high frequencies, from a few megahertz up to 30 GHz or so, because they work extremely well for such simple devices. The output swing is about ± 300 mV for the garden variety to ± 1 V for dedicated ones like the Mini Circuits RPD-1. Their amplitude fidelity is pretty good; until the signal is lost in the DC offset of the mixer, or starts to saturate it, a diode bridge phase detector will provide good linearity in amplitude for a fixed phase anywhere in its range. It works best at $\theta = 0$ or π .

Its phase accuracy is mixed. On one hand, the AM-PM performance of a phase detector is usually excellent near null; if θ is near $\pm\pi/2$, turning the signal amplitude up and down causes only very small changes in the position of the null until the RF level becomes comparable to the LO and this starts controlling the DC offset. On the other hand, though, the detailed shape of the function $V_{IF}(\theta)$ for constant A_{RF} is not all that accurately sinusoidal, and its exact shape depends fairly sensitively on the RF signal amplitude. Another way of looking at this is that all those spurious responses we spent so much time avoiding in Section 13.7.2 land right on top of us when the IF is at DC. Though their frequency is 0, their phases still go as multiples of $\Delta\phi$.

Like mixers, phase detectors are normally used with lots of LO drive to reduce their sensitivity to LO amplitude variations. For a pure phase application, such as FM detection, phase detectors work best with both inputs driven hard.

Gilbert cell multipliers (see Figure 14.16) have much the same characteristics as diode bridges. They are not as resistant to overdriving the RF input, are noisier, and require more ancillary circuitry. On the other hand, they exhibit gain, and if the LO is driven gently, they can have very low spurious products, owing to the accurately bilinear multiplication possible with bipolar differential pairs. Even driving the LO hard requires much less LO power than a diode bridge. Because of their complexity, they are less often used in discrete circuitry than diode bridges, but are a nearly universal choice in bipolar ICs.

Below 1 MHz or so, the CMOS transmission gate approach is best unless the dynamic range is very large. The main problems with these are duty cycle errors, which cause the nulls to be in the wrong places; switching glitches; and charge injection in the switches, which causes a level-dependent DC offset. The large-signal performance of CMOS switches is excellent, so the gradual compression of the output level that afflicts diode mixers is not present.




Figure 15.15. A discriminator is a frequency-dependent phase shift and a phase-sensitive detector.

15.5.11 FM Detectors

Good FM detectors work by taking two copies of the input signal, applying a frequency-dependent phase shift to one of them, and then running them into a phase-sensitive detector (Figure 15.15). Typical ways of getting the phase shift are delay lines (often just a chunk of coax) or resonators (e.g., *LC* tank circuits, quartz crystals, or SAW devices). The phase shift should be as linear as possible, and the total phase shift over the band of interest should be small enough to stay within the linear range of the phase detector (sometimes the nonlinearity of one can compensate that of the other, as in double-tuned quadrature detectors in FM radios). Linearity aside, there is an obvious trade-off between sensitivity and bandwidth for a given output voltage swing and noise level.

The other approach is to turn the FM into AM by sitting on the skirt of a filter, and detect the result with an AM detector. We're often stuck with that sort of thing in optics (e.g., the Pound–Drever technique for locking a diode laser to an etalon), but there's no reason for it in electronics.

15.5.12 Delay Discriminator

The simplest way of getting the phase shift is to use a chunk of coax or a serpentine microstrip trace, some odd multiple of $\lambda/4$ long. In that case, the phase is very linear with frequency, and the resulting signal is approximately

$$V_{\phi} \approx A \cos\left(\frac{2\pi f L}{v}\right). \quad (15.12)$$

Delay discriminators are very wideband, which is their best feature; if we can accept 5% nonlinearity at the phase extremes, a nominally $\lambda/4$ delay line works over a range of ± 0.55 radian, a bandwidth of more than an octave. The downside is that this makes them pretty insensitive. Delay discriminators are really useful in the lab, because you can wire one up with a Mini Circuits mixer, hybrid splitter, and two patch cords of different length, as in Figure 15.16.

An unequal-path interferometer is a delay discriminator too. If there are exactly two paths, and no spatial phase variations, (15.12) still holds; if as usual there are more, we have to put the right-hand side inside a double sum over all the pairs of paths. This isn't all that difficult, although for continuous distributions and multiple scatter (e.g., double Rayleigh scattering in fibers), it does become a bit subtle analytically. Usually, though, we have one path that is much stronger than the others, and the predominant effect is




Figure 15.16. Delay-line discriminator.

interference of stray paths with that one. In that instance, a single sum is enough, and (15.12) becomes

$$V_\phi \approx \sum_{i=1}^N A_0 A_i \cos\left(\frac{2\pi f L_i}{v}\right). \tag{15.13}$$

15.6 PHASE-LOCKED LOOPS

Since we can measure phase electrically and can control it with a voltage-controlled oscillator, we can make a feedback loop controlling the relative phases of two signals, a so-called phase-locked loop (PLL), as shown in Figure 15.17. We don't have too much space for them, but PLLs can do some very powerful things, including pulling good data out of a signal buried in noise, and doing detection in a bandwidth much narrower than the uncertainty in the carrier frequency. You can think of a PLL as a lock-in amplifier that makes its own reference signal by following the input.

Crudely understood, a PLL generates a VCO signal whose phase follows the input signal's within the PLL bandwidth f_0 , and ignores it outside. A narrowband PLL follows only the average frequency, and so the phase detector output can be used as a PM detector for modulation frequencies above f_0 . A wideband one follows the small dips and wheels of the instantaneous frequency, so the VCO control voltage can be used as an FM detector for sufficiently slow modulation.

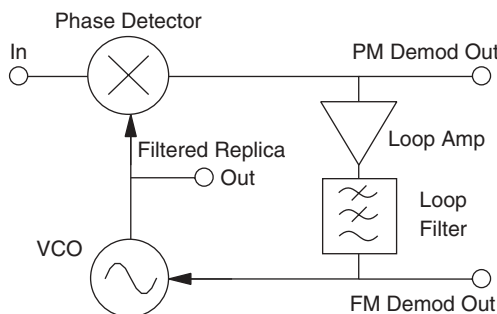


Figure 15.17. Phase-locked loop.

If you're going to build your own PLLs, read Gardner's *Phaselock Techniques*, and watch out for three things: acquisition of lock, which is difficult in wide-range loops with narrow loop bandwidths; severe reference frequency ripple due to the loop amplifier and filter having too little attenuation at f or $2f$; and, of course, frequency compensation in the face of gain variations in the VCO (K_{VCO} in rad/s/V) and the phase detector (k_ϕ in V/rad), which change with center frequency and input signal level.

15.6.1 Loop Design

In Sections 15.4.1 and 18.4.1 we talk about frequency compensation of amplifiers. PLLs are slightly more complicated because if the VCO's radian frequency sensitivity is K_O rad/s/V, its transfer function is

$$A_{VCO}(f) = \frac{K_O}{j\omega} = \frac{1}{jf} \frac{df}{dV}, \quad (15.14)$$

which is an ideal integrator. The reason is that we control f but measure $\Delta\phi$, and phase is the integral of frequency. (It's usually less confusing to stick to radians for PLL calculations.) The PD characteristic is roughly $K_\phi \sin \phi$, but all that matters is the slope near the operating point (i.e., 0 V), so its gain is assumed to be constant and frequency independent. If the loop filter gain is $A_F(\omega)$, the overall loop gain is

$$A_{VL}(\omega) = \frac{K_\phi K_O}{j\omega} A_F(\omega), \quad (15.15)$$

so all we have to do is choose A_F . Since we want a phase margin of at least 45° , and the integrator contributes 90° , we have 45° of phase shift to work with when $A_{VL} > 1$.

The phase detector produces a lot of ripple at $2f$, and usually some at f as well. Diode bridge phase detectors are generally best for high stability applications, due to their very low noise and predictable ripple. Phase-frequency detectors of the 74HC4046 type ideally produce no ripple at zero phase error, and don't need aided acquisition because the output is railed when the loop is out of lock, rather than producing a beat note like other PD types. Unfortunately, they exhibit a dead band right at the zero-ripple point, so you have to bias them a bit off center, which makes their claimed ripple advantage more apparent than real.

Generally, the best loop filters are op amp integrators with lead-lag networks, as shown in Figure 15.18 (assume $R_a = 0$ and $R_b = \infty$ for now), because their high DC gain makes the static phase error very small. For our purposes, the best simple method for compensating the loop is first to decide on the desired loop bandwidth f_L . For a reasonable phase margin, set the zero $1/(2\pi R_f C_f) \sim f_L/2$, and choose $R_1 + R_2$ and R_f so that the gain drops to 1 at f_L (remember that the phase detector has gain in V/rad as well). Then you can put in C_1 to roll off the ripple at f and $2f$. If necessary, you can use a notch filter to increase the ripple rejection, but that requires careful attention to phase shifts. Either way, you'll want to use a math program or SPICE to verify that your phase margin is at least 45° . Another source of phase shift is the phase detector itself; linear multipliers allow the output to react to an input phase shift immediately, but saturated or XOR PDs can't react until the next transition, which is a quarter-cycle delay, and a PFD can't react until the next cycle, which averages a half-cycle delay. Make sure you account for these effects if you're building a wideband PLL.

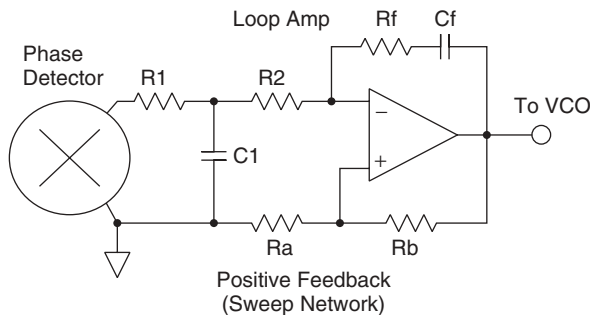


Figure 15.18. Narrowband PLL loop filter with automatic acquisition: when the loop is unlocked, the overall loop gain drops to 0 and the DC component of the phase detector output becomes very small. The positive feedback network R_a – R_b takes over, turning the loop amplifier into a slow triangle wave generator ($f \ll f_c^2$), which sweeps the VCO frequency until the loop acquires lock.

15.6.2 More Complicated PLLs

So far we've talked only about 1:1 PLLs, where the VCO frequency is the same as the input frequency. If you put a frequency divider between the VCO and the phase detector, you can make a frequency multiplier, and if you put dividers on both the input and the VCO, you have a *frequency synthesizer* that can in principle make any rational multiple of the input frequency. In practice, this hits limits due to noise multiplication [the phase noise goes up by a factor of $20 \log(f_{\text{out}}/f_{\text{ref}})$], and due to the requirement that $f_L \ll f_{\text{ref}}$ for ripple rejection. Offset loops, in which f_{VCO} is mixed down to a lower frequency f_{IF} before phase detection, are also frequently useful—especially when using acousto-optic deflectors, where they allow the detected signal to be mixed down to a fixed IF (see Section 15.2.1).

15.6.3 Noise

There are some pretty poor PLL chips out there. For high SNR applications, avoid anything with a digital-logic phase detector or an RC VCO. (Chips like the CD4046 are great for low performance applications.) Keep the resistor values low, as you would in a low noise amplifier, and use a high output diode bridge phase detector such as a Mini Circuits MPD-1. Those have nanovolt output noise densities and about 500 ohm output impedance, so use really quiet bipolar op amps such as the LT1028A. Use good quality VCOs, such as a VCXO for narrowband applications or a varactor-tuned LC VCO with automatic level control for wider ranges. The phase noise of the VCO is suppressed by the loop gain, so try to keep the bandwidth wide unless the oscillator is really quiet. High comparison frequencies are your friends, because they reduce noise multiplication. For more on PLL noise, see Gardner.

15.6.4 Lock Detection

It's obviously important to know when the loop is in lock. A phase–frequency detector makes this easy; PFDs usually come with a lock detection output, and (more usefully) they rail when they lose lock, so a window comparator on the VCO control voltage line

will tell you when the loop is unlocked or near the edges of its range. Lock detection is more of a puzzle when using diode or XOR phase detectors, because their DC output is near zero whether they're locked or unlocked. The usual approach is to use a second phase detector, driven 90° out of phase, and look for a large DC level there to tell us we're in lock. We're often interested in the AM information on our signal, so we'd likely be needing an in-phase output anyway (see Section 13.9.3).

15.6.5 Acquisition Aids

Slow PLLs take a very long time to acquire lock, and if they're sufficiently far off, they may never make it. Figure 15.18 shows a high gain PLL loop filter for use with a low noise diode bridge phase detector (e.g., a Mini Circuits MPD-1). This filter has a built-in acquisition aid—when the loop loses lock, it becomes a triangle wave oscillator that sweeps through its entire tuning range until it reacquires. Choose R_a/R_b to be a couple of times V_{OS}/V_O , where V_{OS} is the maximum offset voltage of the phase detector and V_O is the maximum output voltage of the amplifier. (Split supplies are assumed.) Alternatively, you can use a phase-frequency detector as an acquisition aid, but that's a bit more involved because PFDs work at 0° rather than the diode bridge's 90° (a divide-by-4 walking-ring counter is often useful). The only drawback of the triangle wave approach is that the hysteresis forces the phase detector to work slightly off null, which slightly degrades the AM rejection of the loop. If that's a problem, and you have a lock detector, you can use it to control an FET switch that grounds the + input.

15.7 CALIBRATION

A detection system is composed of a detector and a calibration; a barefoot detector isn't much use for quantitative data. Make sure you know how your detector is going to be calibrated before you go too far with the design. One particularly thorny issue you should wrestle with early is whether you can afford to do a two-dimensional calibration; I/Q measurements need to have the I and Q detectors calibrated over a wide range of amplitudes and phases—their output isn't just $(A \cos \phi, A \sin \phi)$. A clever strategy can accomplish this, but it doesn't happen by accident. The alternative is to use a phase-insensitive envelope detector plus a nulling phase detector, which needs only a 1D calibration.

Calibration is often made much easier through the process of linearization, which not only reduces the size of the corrections but often eliminates their cause; for example, putting a matched CMOS switch in the feedback loop of an op amp driven by an analog mux can correct for switch resistance, nonlinearity, and temperature drift all at once (see Section 15.10).

15.7.1 Calibrating Phase Detectors

Phase detectors are easy to calibrate, since phase is easy to turn into time and frequency. You can use two divide-by- N digital counters, driven by a crystal oscillator or synthesizer, to produce two phase-coherent signals for the signal and reference phase inputs of the phase detector (a mixer and filter arrangement can shift them up to frequencies too high for the counter outputs). Make one a pulse-swallowing counter, whose output phase

can be shifted by $2\pi/N$ under software control, and you can take N phase points at each amplitude level. Cubic spline interpolation will do an excellent job for a 12 bit nulling-type phase detector with $N = 100$. The major things to avoid, as in all time and frequency measurements, are leakage of one phase into another, multiple reflections, anything that will cause the phase steps to be unequally spaced. Use lots of shielded boxes, attenuators, and isolation amps in your calibrator, and test it by sitting on the pulse-swallowing button (to shift one channel's frequency slightly) while looking at the outputs on a spectrum analyzer. That'll show up spurious coupling pretty fast.

15.7.2 Calibrating Amplitude Detectors

It isn't that easy to calibrate an amplitude detector to the accuracy of your ADC, mainly due to the difficulty of getting accurate calibration levels. Probably the best way is to use a quartz crystal, excited by a long tone burst, and terminated in a very low, resistive impedance (Figure 15.19). When the burst cuts off, the crystal will ring at its natural frequency f_s (i.e., its series resonance), with exponentially decaying amplitude. Once the ring-down is started, samples taken at regular time intervals will have accurately known relative amplitudes and (even better) will produce a linear ramp from a logarithmic detector. Crystals have Q_s of 10^4 to 10^6 , which means the decay is nice and slow; a decent 80 MHz crystal will decay at around 1 dB/ms, which is very convenient.

A stable amplitude detector of limited dynamic range (e.g., a temperature compensated diode detector) can be used to measure the initial amplitude and the time constant, from which our *a priori* knowledge of the envelope's functional form allows a good calibration over a very wide range. The main opportunities for corruption here are nonlinearity in the buffer amplifier following the crystal, overdriving the crystal with the tone burst, or allowing wideband noise to dominate the signal at late stages. Keep the rms current in the crystal between $100\ \mu\text{A}$ and 1 mA during the tone burst.

One problem is that the tone burst has to be at $f_s \pm f_s/(2Q)$ or so, or the crystal won't ring very much. This is pretty easy to solve in the lab, because you can just tune your synthesizer until the output peaks, but is more awkward for an online calibration. In that case, you can make the crystal into an oscillator. It isn't as easy to get a good

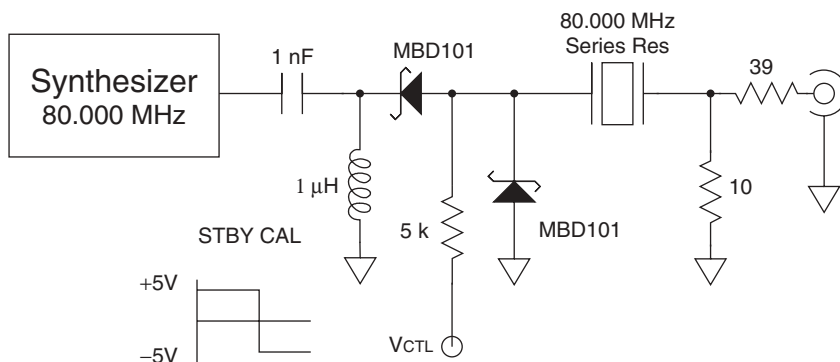


Figure 15.19. Crystal ring-down calibrator: For low oscillation amplitudes, the rate of decay of the oscillation is accurately exponential until it becomes comparable to the noise level in the measurement bandwidth. Since timing is easy to do accurately, this provides an excellent linearity calibration for amplitude detectors and limiters.

result, because you have to get the crystal to oscillate right at f_s , which requires external series inductance, and you have to use diode switches to ground the transistor end of the crystal to cut off the burst.

15.7.3 Calibrating a Limiter

We've talked about AM–PM conversion in limiters, which also requires calibration. The crystal ring-down method will produce the AM–PM characteristic in radians/dB as a function of time, which is very convenient. Putting a calibrated limiter ahead of your phase detector is another way to reduce the calibration to two 1D problems instead of one 2D problem.

15.8 FILTERS

In Section 13.8.5 we talked about filtering in general. Here we'll go into the gory details. A cascade filter is just a bunch of resonant sections strung together in a row. If we put buffers in between them so that they don't interact (as in an op amp active filter), we can calculate their individual transfer functions from the series and parallel impedance formulas, and then just multiply them together to get the filter transfer function, as we did building feedback loops in Section 15.4.3, but in the pure LC case it's a bit harder because the sections detune each other and so the design equations become complicated. Thus we usually start from some canned design, transforming it to fit roughly, then running a numerical optimizer and tweaking component values and tolerances to get the final design. *Filter etiquette note:* People sometimes talk about “selectivity” when they mean “stopband attenuation,” leading to embarrassing stares when they mention their “high-selectivity RC filter.”

15.8.1 LC Filters

An LC ladder filter is a cascade of series and shunt sections. (We'll specialize to the low-pass case for now, so it's series inductors and shunt capacitors, as shown in Figure 15.20.) Such filters have transfer functions that are rational functions (i.e., ratios of polynomials) in f . This is because the element impedances themselves are polynomials or rational functions of frequency, and putting them in series or parallel is equivalent to making rational functions of the impedances. Thus there is no way out of the class of rational functions for these filters' frequency responses. It is in principle possible to make an adequate approximation to any causal transfer function with enough LCR sections and enough delay, but in practice tuning highly complex filters is a serious problem.

Rational functions exhibit poles and zeros in the complex plane. These are nothing more than the roots of the denominator and numerator polynomials, respectively. The positions of these poles and zeros determine the form of the transfer function, so transfer functions are designed by moving poles and zeros around, as we saw in Sections 13.8.5 and 15.4.3.

Historically, people have cared more about amplitude than phase, so the usual procedure is to pick a convenient even-order rational function for $|H(f)|^2$, then factor the numerator and denominator polynomials to find the poles and zeros (picking the stable ones, of course). You'll probably want to look the values up in Zverev instead. (One comforting fact is that all ladder filters are minimum-phase networks.)

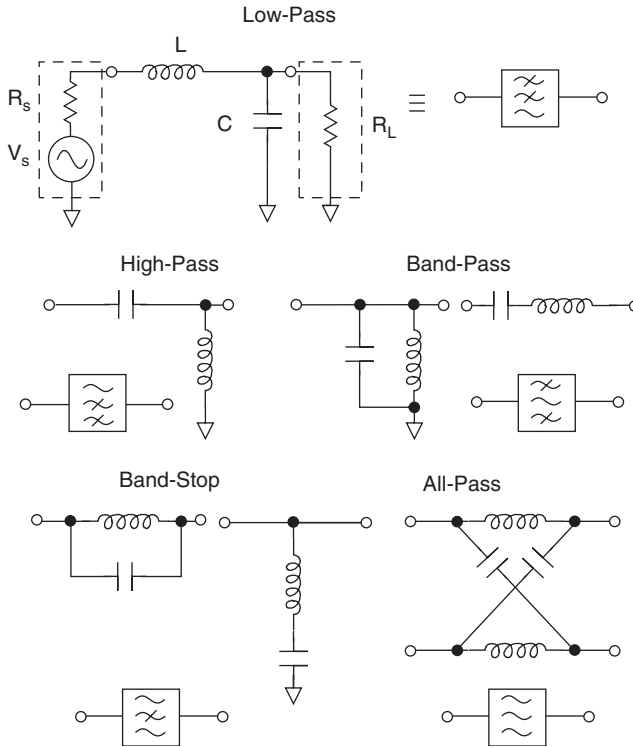


Figure 15.20. LCR filter sections; the all-pass is a balanced lattice section, the rest unbalanced ladder sections. Band-pass and band-stop filters have both series and shunt realizations.

15.8.2 Butterworth Filters

The simplest LC filter is the Butterworth, which has

$$|H(f)|^2 = \frac{1}{1 + (f/f_c)^{2n}}, \tag{15.16}$$

which is very flat near 0 (the vastly overrated *maximally flat* property), but rolls off continuously through the passband. Butterworths are easy to design and have only moderate group delay problems, but have few other redeeming features.

15.8.3 Chebyshev Filters

Chebyshev polynomials have uniform-sized ripples for $x \in [-1, 1]$, and steeply rising monotonic behavior outside that. A lowpass Chebyshev filter comes from picking

$$|H(f)|^2 = \frac{1 + \epsilon^2 T_n^2(0)}{1 + \epsilon^2 T_n^2(f/f_c)}. \tag{15.17}$$

There will be a small amount of passband ripple, approximately $\pm 10\epsilon^2/\ln(10) \approx 4.3\epsilon^2$ decibels. Obviously there is a trade-off between how fast the filter falls off and

how big the ripples are; in practice, we rarely use filters with more than 0.5 dB of ripple, and 0.1 or 0.2 is more common. Chebyshev filters distribute the passband error more uniformly and gain sharper skirts in the process, which is pretty worthwhile.

15.8.4 Filters with Good Group Delay

As we saw in Section 13.8.9, if we want good group delay we have to accept either very gradual skirts and some passband bowing (like a Butterworth but even worse), or use a group delay equalizer to even it out. Inside a feedback loop it is important to have minimum phase, so in that case look for Bessel, Gaussian, or equiripple phase filters.

15.8.5 Filters with Good Skirts

The filters we've considered so far are *all-poles* filters, that is, the numerator is a constant, so there is a nonzero response at all frequencies and a monotonic falloff in the stopband. By putting a higher order polynomial in the numerator, we can make the skirts fall off more rapidly, at the cost of poorer attenuation deep in the stopband. The *elliptic function* or *Cauer* filter is the classical design, based on elliptic function theory, but any filter with zeros at real frequencies is usually called an *elliptic filter* anyhow. An example is a lowpass filter with inductors in the series legs and series *LC*s in some of the shunt legs. These can be made with highly asymmetric skirts, which is great if you have some ugly close-in spur to get rid of. Look these up if you need to use one. The group delay of these things is not a sight for the faint-hearted.

15.8.6 Lowpass to Bandpass Transformation

Bandpass filters are generated from the lowpass prototype by a simple transformation, in three steps.

1. Choose a lowpass filter with the same bandwidth.
2. Find the component values for the prototype lowpass from tables (*e.g.* Zverev).
3. Resonate each element at $f_0 = (f_1 f_2)^{1/2}$, the design center frequency. For a bandpass filter, you want the impedances at f_0 to be correct for the center of the passband, that is, the equivalent of DC for the lowpass: series-resonate the series elements, and parallel-resonate the shunt elements. For a bandstop, you want f_0 to be the center of the stopband, so you do it the other way.

This transformation is an example of a *conformal map*, in the bandpass case $f' \rightarrow f - f_0^2/f$. The bandpass and bandstop filters designed this way have symmetrical responses when plotted against $\log(f)$. The imaging of the filter response from lowpass to bandpass results in the section reactances changing twice as fast with frequency, because the inductive reactances are going up and the capacitive ones are going down. The full bandwidth stays the same, because although each side is half as wide, we get an image on each side of f_0 . The settling time is doubled compared with the lowpass prototype. If you're trying anything the least bit fancy with this filter (*e.g.* constant group delay), you should stick a numerical optimizer on the design when you're done. Remember from Section 13.8.8 that the nonlinear frequency warp makes this transformation useless for linear-phase filters.

Aside: Component Limitations. The design procedure is general enough to let you design a 100 Hz wide Butterworth bandpass filter at 1 GHz, but unfortunately that would need 30 millihenry inductors and 30 attofarad capacitors, both with Q values of 10^8 , and you won't find those at Radio Shack. You can do the equivalent with a superhet, as we saw in Chapter 13. When your component values become extreme, you're reaching the limits of realizability of a filter.

15.8.7 Tuned Amplifiers

A tuned amplifier stage has a tuned circuit in its input, output, or both. A tank circuit in the collector of a common-emitter stage is a common example, and so is a photodiode amp with the diode capacitance resonated away with a series or parallel inductor, as in Section 18.5.1. Tuned amps are useful in reducing the number of filters you need, reducing the broadband noise, and giving you flexibility in the impedance levels you use. They have to be adjusted manually, so stick to low Q values (usually below 10).

15.8.8 Use Diplexers to Control Reflections and Instability

You can compensate a series RC by putting a series RL in parallel with it; the R s have to be equal, and the L is R^2/C . This actually gives constant, pure resistance equal to R at all frequencies. The nice thing about this is that the high frequency energy goes to one resistor and the low frequency energy to another; you can filter the output of some termination-sensitive device like a diode mixer without an impedance mismatch, sending a low frequency IF to the LC arm and the image and high frequency spurs to the RC arm. This is a simple example of a *diplexer*.

15.9 OTHER STUFF

15.9.1 Diode Laser Controllers

Diode laser power supplies need to have very low noise, preferably below the shot noise. This isn't hard, with care. The idea is to make it extremely quiet at AC, then wrap a slow feedback loop around it to make it stable at DC, while keeping very careful control of its transient response into weird loads. We saw in Sections 14.6.5 and 18.4.6 how to make a very quiet current source, and an op amp sensing the collector current and comparing it with a heavily filtered voltage reference finishes the job, as shown in Figure 15.21. Make sure you supply a mechanical relay contact across the outputs to protect the laser against electrostatic discharge when it's off.

15.9.2 Digitizing Other Stuff

So far, we've stuck with converting our signals to voltages and digitizing them afterwards. This is the usual way of doing things, but not always the best. For example, consider the phase digitizer of Figure 15.22. A state machine called a *successive approximation register* uses a DAC driving a voltage-controlled phase shifter to perform a binary search of the DAC's code space to find the null at the phase detector output. This keeps the phase detector right at 90° all the time, so the calibration only needs to cover errors due to amplitude, and not phase nonlinearities changing with signal level. The state machine

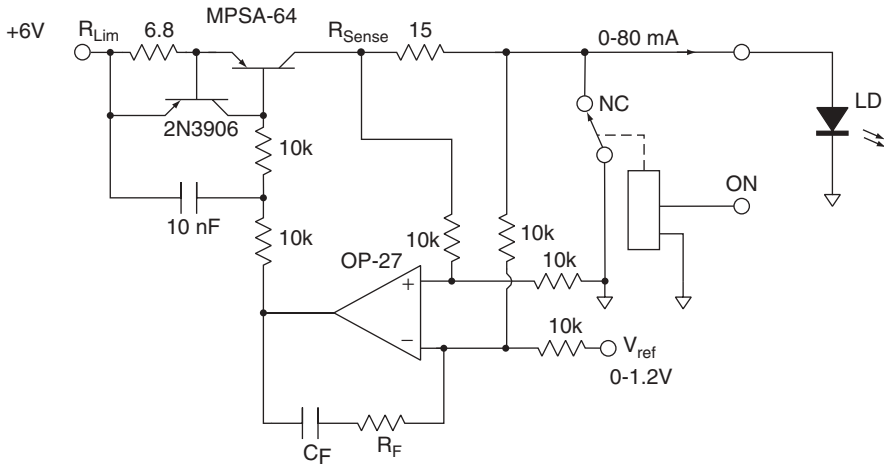


Figure 15.21. Diode laser controller design.

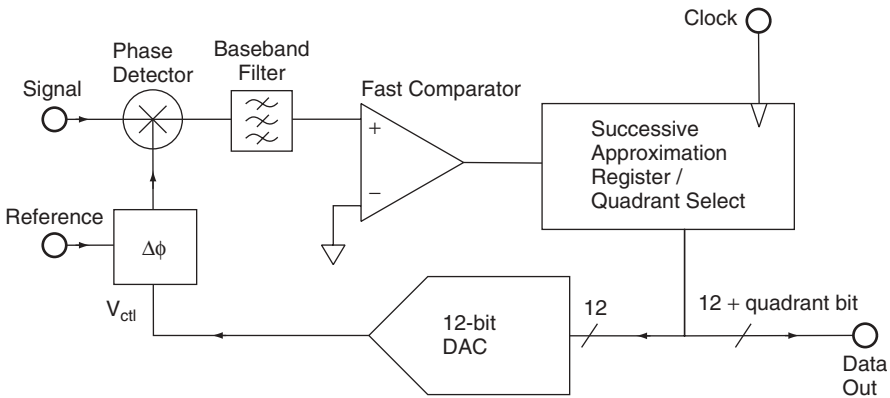


Figure 15.22. Successive approximation phase digitizer: the successive approximation logic finds the correct null by a modified binary search.

needs an extra flip-flop to make sure it's shooting for the right null—there are two, of course, and only one is stable. (Why?)

15.9.3 Use Sleazy Approximations and Circuit Hacks

This far into this book, it isn't necessary to repeat it too loudly, but there's usually a circuit hack to fix a seemingly intractable problem, as long as it's technical and not fundamental. For example, we can make really, really quiet amplifiers, but can't reduce the shot noise of starlight except by gathering more. It's worth having a good look at the temperature-compensated breakpoint amp (R. Widlar, National Semiconductor Application Note AN4) and the low phase error amplifier of Section 15.10 (Analog Devices Application Note AN-107). Some of the oscilloscope front end tricks in Jim Williams's books are really worth reading too. Chapter 18 has an extended example of

how to get through a difficult circuit problem (a shot noise limited front end for low light), and talks about the laser noise canceler, which is another neat circuit hack for doing ultrasensitive detection with noisy lasers.

15.9.4 Oscillators

We often talk about the characteristics of RF signals without much attention to where they come from, but unless you're listening to a pulsar with a crystal radio, all must come from some oscillator ultimately. An oscillator consists of an amplifier and a frequency-selective network, which feeds back the oscillator's output to its input. The network is often a resonator, but this is not always the case; the criterion for oscillation is that the total phase shift around the loop should be an exact multiple of 2π , while the gain equals 1. This condition sounds very delicate, but in fact is not; the frequency continuously adjusts itself to make the phase condition true, and if the gain is initially larger than 1, the amplitude of the oscillation will increase until circuit nonlinearity reduces it to an average value of exactly 1, the *self-limiting* condition. This intrinsic nonlinearity causes the low frequency noise of the amplifier to intermodulate with the oscillation, producing noise sidebands at low modulation frequency, which is why the best oscillators use ALC instead of self-limiting (see Sections 15.6.3 and 15.10.5). Oscillators have noise in both amplitude and phase, with the phase noise generally being much more objectionable, because it's hard to get rid of; amplitude noise can be eliminated with an external limiting stage, but reducing the phase noise requires redesigning the oscillator or phase locking it to some quiet reference source.

Oscillators are divided into categories, depending on the nature of their feedback networks and how the amplitude limiting is done. Stable and quiet oscillators result from quiet amplifiers, in feedback loops based on stable resonators whose phase shifts vary rapidly with frequency, and amplitude limiting schemes that minimize the nonlinear mixing of the amplifier noise with the oscillator signal.

They aren't that hard to build, but they need lengthy explanation, so for the purposes of this book, oscillators are something that you buy, or build from oscillator ICs such as the LMC555 astable or SA601 RF mixer/oscillator.

15.10 MORE ADVANCED FEEDBACK TECHNIQUES

15.10.1 Put the Nonlinearity in the Loop

We saw that changes in A_{VOL} get suppressed by a factor of A_{VL} . It isn't difficult to use this linearizing property to correct for other circuit elements. There are two ways to do this; we'll use a FET switch as an example. The first way is to put the switch in series with the output, with another one to keep the amp from railing when the first one opens, as in Figure 15.23a. This relies on the stability of the summing junction for its OFF isolation, and consequently isn't great at high frequency, but works on unmatched switches and gains other than -1 . You can also rely on monolithic matching and use a constantly closed switch in a unity gain inverter's feedback loop to provide a matched nonlinearity, as in Figure 15.23b (see also Section 15.12.5). It's possible to combine the two tricks ju-jitsu style, to make the op amp's input errors track themselves out, as in Figure 15.23c. This puts some stress on the amplifier's frequency compensation, so it's worth simulating it before using, just to make sure it does what you expect. (This is one place where even op amp macromodels might come in useful.)

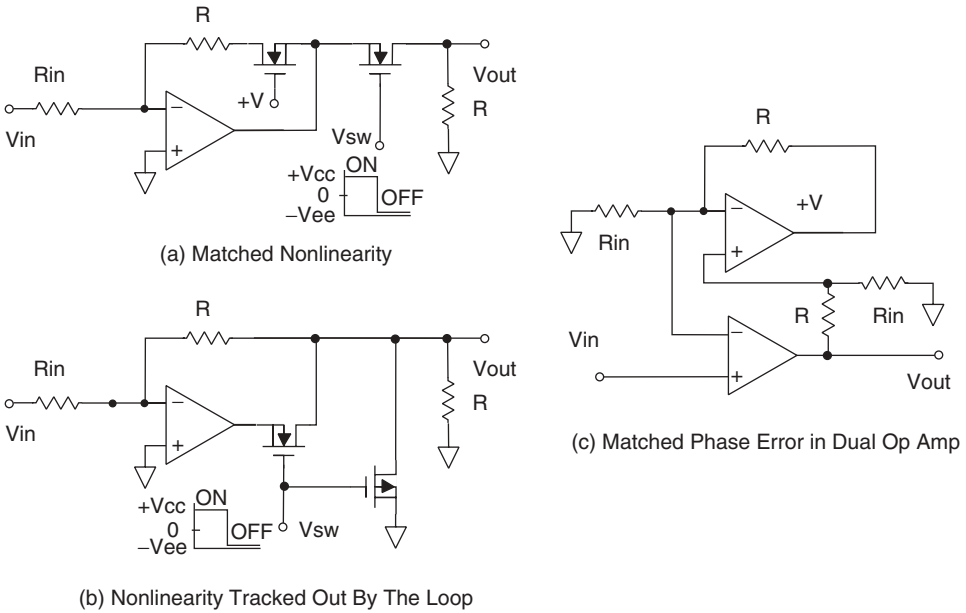


Figure 15.23. Eliminating nonlinearity: (a) matched on-resistance nonlinearities cancel; (b) nonlinearity tracked out by the loop (extra switch needed for off-state isolation; and (c) combining (a) and (b) phase error due to op amp frequency response can also be tracked out, although some amplitude peaking occurs (see Analog Devices Application Note AN-107).

15.10.2 Feedback Nulling

One interesting method of suppressing background is to use a servo loop that detects the background plus signal, and nulls out the whole works. Of course, this helps only when the desired signal is a transient, so that the loop doesn't suppress the signal as well as the background. Dithered systems using Bragg cells are especially suitable for this (see Section 10.8.3).

15.10.3 Auto-zeroing

A related technique is auto-zeroing, where we measure the background at a time when the signal is known to be 0, and subtract it out in analog. An example is correlated double sampling in CCDs (see Section 3.9.4), where the kTC voltage uncertainty is different for each pixel clock cycle, but we can measure it separately and so eliminate it completely. We can use it to subtract off the huge sloping baseline in current-tunable diode laser measurements, by adding in an adjustable amount of the scan signal, using a multiplying DAC. These examples may seem like sleight of hand, but it can help a lot; unless your system is taking data continuously at all times, the local servo loop will do a better job than postprocessing for this, and the demands on the system's dynamic range are enormously reduced. In the diode laser example, the laser power variation can easily be 50%, and if we're trying to detect absorptions in the 10^{-4} range, that's 74 unnecessary decibels.

Auto-zeroing at a rate of f_s suppresses noise at lower frequencies, and puts some funny ripples in the noise spectrum up to a few times f_s due to switching transients and redistribution of noise power. Signals at integer fractions of f_s are nulled perfectly, and intermediate frequencies go roughly as f/f_s . This is valuable in eliminating the horrible $1/f$ noise of CMOS amplifiers and is widely used in very low level DC amps.

15.10.4 Automatic Gain Control

It's tempting to extend this idea to automatically changing the gain as well as the offset, but that's not usually a good idea. AGC is used in radios with great success, but there the exact signal level is of no concern whatever, which is utterly unlike the situation in instruments. Linearity and absolute accuracy tend to get compromised pretty badly with AGC control, and fast AGC loops (faster than $\sim 1\%$ of the IF bandwidth) can produce parametric effects, such as asymmetric pulse decay, or phase modulation as a function of AGC voltage. It is often best to use a combination of range switching, redundant detectors, and digital scaling. One exception is in PLL phase demodulators, where optimum low SNR performance requires a combination of wideband and narrowband AGC (see Gardner).

15.10.5 Automatic Level Control

On the other hand, AGC's twin brother *automatic level control* (ALC) helps enormously in oscillators and frequency synthesizers, by allowing them to operate in a highly linear regime while having enough excess gain to start up reliably—see Section 15.9.4.

15.10.6 Feedback Loops Don't Have to Go to DC

It's a good idea to use DC coupling when you can, because it saves blunders. Things that are obvious when the DC is there (e.g., a photocurrent of 0 mA or a front end whose output is pegged) are easily missed in a purely AC-coupled system. On the other hand, feedback loops can sometimes do remarkable things when you don't DC couple them. This is usually because of the perfection of capacitors in subtracting off DC. One example is the capacitance multiplier of Example 14.1. Another is the high voltage supply quietener of Figure 15.24, which jiggles the bottom of a high voltage supply to keep the top still, as measured via a capacitor. Ten volts of 120 Hz ripple can be brought down to the 100 μV level this way for about \$5, which really helps the stability and noise of piezoelectric actuators and PMTs.

Aside: Motorboating. Instability usually results from too much phase lag at high frequency, but AC-coupled feedback loops can also exhibit instability at low frequency due to accumulated phase *lead*. This instability is called motorboating, an apt description of how it sounds over a speaker.

15.11 HINTS

Invert When Possible. If you can build your circuit inverting, do it. Op amps are not as happy when there's a big common-mode swing on their inputs, and their protest shows up in nonlinearity and offsets. Inverting operation keeps the inputs still, somewhere near

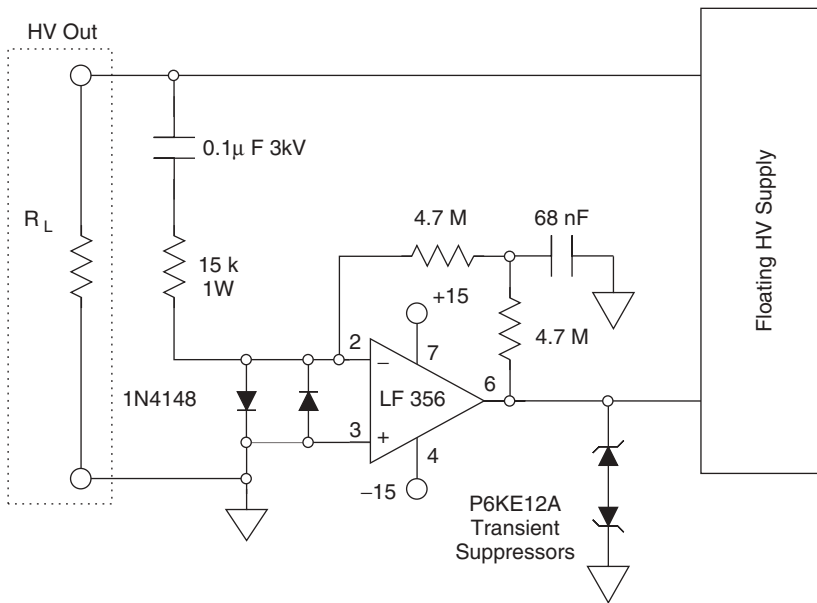


Figure 15.24. High voltage power supply quietener.

the middle of their range, and that helps a lot; you also never run into the problem of somebody's wide output swing exceeding the next guy's input common-mode range that way. In general, inverting stages with similar gains will be noisier than noninverting ones, because we often need to use higher value resistors in inverting stages, and because an amp with inverting gain $-A$ has a noninverting gain of $1 + A$, so you get more signal for the same noise by not inverting.

Watch for Startup Problems. There are lots of circuits that don't start up well. An LM7805 regulator, which is of course a byword for good behavior, will fail to start if you drag its output below ground while its input is coming up; its output remains at a high impedance, so it'll never pull itself back above ground. This can easily happen if the negative supply starts up quickly, but a rectifier connected from output to ground will prevent it. Circuits can also latch up, for example, old technology CMOS chips, which made quite reasonable SCRs if you took their inputs outside the supply range.

Besides failure to start, circuits can also be damaged by turn-on transients. For example, consider a 0–10 V multiple DAC driving a low breakdown voltage device such as a micropower Gilbert cell mixer used as a DAC-controlled gain source, or a ± 15 V op amp driving the base of a transistor with its emitter near ground, as in Section 18.6.3. These devices can be damaged by momentary overvoltages, and those are almost inevitable on startup unless you use a voltage divider, current limiting resistor, or protection diode. You can normally ignore this problem when using ICs running off the same power supplies, but not always. Note: For simulation devotees, SPICE often misses this sort of problem.

Subtract, Don't Divide. Since all the photocurrents in most instruments are exactly proportional to the source power, it is very tempting to measure the source power

independently and then divide it out, either with an analog divider chip or in software, yielding a beautifully normalized result. Except that it doesn't usually work very well.

Most of the time, we're measuring a small signal on a big background (the proverbial grass on top of the Empire State Building), so we have a little dynamic range problem. Let's say our light source has a $1/f$ noise δ of 0.2%/decade, and that our signal is a 5 mV rms ripple on a 5 V background. In 1–10 Hz, that's $5 \text{ V} \cdot \delta = 10 \text{ mV}$ of noise from the background, plus noise intermodulation of $5 \text{ mV} \cdot \delta = 10 \text{ } \mu\text{V}$ from the signal itself. An ideal divider fed a faithful replica of the source power will completely eliminate both. But let's go on carefully. Analog dividers with a $\pm 10 \text{ V}$ range have flatband noise of about $1\text{--}2 \text{ } \mu\text{V}/\text{Hz}^{1/2}$ with a full-scale denominator, and a $1/f$ corner of 10 Hz or so—about 60 dB worse than a good op amp. By the time we've thrown in that -60 dB signal-to-background ratio, we've given up 120 decibels of dynamic range before we even start. In our 1–10 Hz bandwidth, the noise is

$$\begin{aligned} e_N^2 &= (2 \mu\text{V}/\text{Hz}^{1/2})^2 \cdot \left(9 \text{ Hz} + \int_{1 \text{ Hz}}^{10 \text{ Hz}} \frac{10 \text{ Hz}}{f} df \right) \\ &= 2 \mu\text{V} \cdot \sqrt{9 + 10 \ln(10)} = 11.3 \mu\text{V rms}, \end{aligned} \quad (15.18)$$

which limits us to a 54 dB SNR. If the measurement intrinsically has a shot noise limited CNR of 140 dB in 10 Hz, then even with the background taking up 60 dB of that, we've lost 26 dB to the divider, which is just plain silly.

It's even worse if we digitize first, because even a noiseless, perfect 16 bit digitizer has only 107 dB; combining two measurements loses us 3 dB, and the 60 dB leaves us a maximum SNR of 44 dB, a 36 dB loss solely to the digitizer.

The right way to do it is to subtract instead. A op amp plus an 8 bit multiplying DAC, (with the source intensity feeding the reference) will get rid of the background down to the 10 mV level, including almost all of the additive noise; after that, digitizing the remainder and dividing by the source intensity in software will do the rest of the job.[†] Because of the high accuracy required, you have to do the subtraction right at the detector, and good background suppression requires great accuracy in the ratios of the two arms. For measurements in white light or with incoherent sources with lots of low frequency noise, this quasidigital method reduces the dynamic range to something your digitizer can probably handle. If you're doing a laser measurement, you can do a great deal better than this using a laser noise canceler (see Section 10.8.6).

15.12 LINEARIZING

In this day of high precision ADCs, fast computers, and cheap memory, why do we care if a measurement is linear? We're going to calibrate it anyway.

There are a number of reasons. Severely nonlinear measurements are a waste of good bits, because their resolution is widely variable over the input range. Nonlinearity inside a feedback loop changes the loop bandwidth with signal, which leads to weird parametric effects, poor settling, and possible oscillation. A measurement that is inherently linear will usually have a much more stable calibration curve, which helps a lot with real-world

[†]This of course assumes that the two are sampled simultaneously. If you're using a data acquisition card, they probably aren't, and it matters.

accuracy—the calibration can be simpler and less frequent. Harmonic and intermodulation distortion of the signal is much reduced by linearizing, which allows simple and extremely sensitive system tests using sine waves and filters.

The basic techniques for linearization are circuit balancing, matched errors, feedback, feedforward, and bootstrapping.

15.12.1 Balanced circuits

We've encountered lots of balanced circuits, for example, diode bridge mixers and transistor differential pairs. The symmetry of these devices suppresses even-order harmonics, because the positive and negative half-cycles are treated the same way. Since the coefficients of the distortion polynomial (Section 13.5) are usually steeply decreasing with order, this makes a big difference to the circuit linearity.

15.12.2 Off-stage Resonance

A varactor phase shifter can be linearized, and its range extended, by wrapping series and parallel inductances around it, as shown in Figure 15.25. A nearby resonance speeds up the variation of reactance with tuning voltage at the ends of the range, where the varactor tuning curve flattens out. The MMBV432 series-resonates with the 45 nH inductors just beyond the high capacitance (2.8 V) end of the range, and parallel-resonates with the sum of the 43 and 45 nH just beyond the low capacitance (10V) end. Proper placement of these resonances leads to a reactance range from near 0 to effectively ∞ , so we can get very nearly the theoretical range of π radians per section. Its linearity is much better than with just the varactor.

15.12.3 Waveform Control

Phase detectors have output characteristics that depend on the input waveform. If the waveforms both have flat tops (e.g., square waves or clipped sine waves), then until

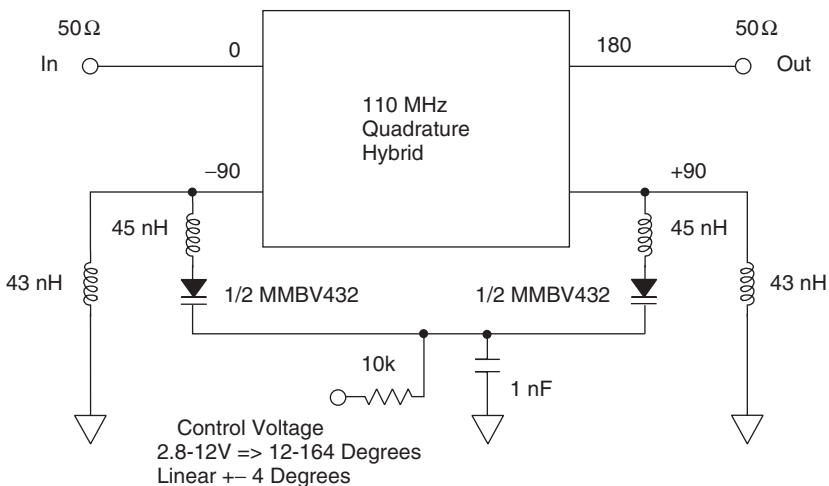


Figure 15.25. Linearized varactor phase shifter.

the phase shift gets large enough that the edges start to overlap, the phase detector's $V(\phi)$ curve will be linear. Only the ends of the curve will depend on the details of the waveform edges. Thus a Gilbert cell phase detector can be linearized by driving both the LO and RF ports hard enough to switch the transistors completely, and a diode bridge phase detector (which doesn't like being overdriven) can be driven from logic gates.

15.12.4 Breakpoint Amplifiers

A last resort is to make an amplifier whose distortion is the inverse of your circuit's. This is so tempting and so nearly worthless that we won't discuss it much. The origin of the nonlinearity of most circuits is some parameter that drifts with temperature and time, so that tweaking a breakpoint amplifier to fix it is a complete waste of time. One important exception is resistance heaters, whose output power is quadratic in the input power, so that the loop gain will vary all over the map with set point and ambient temperature; a breakpoint amp can limit the range, and since it's a low accuracy application and inside the feedback loop, it will work okay. You can stick diodes in the feedback loop of an op amp, which drifts like mad with temperature, or try Widlar's temperature-compensated, ultrasharp breakpoint amp idea from National Semiconductor's Application Note AN4, which is much more accurate but still slow. Good luck.

15.12.5 Feedback Using Matched Nonlinearities

You can also rely on monolithic matching and use a constantly closed switch in a unity gain inverter's feedback loop to provide a matched nonlinearity. If you have a nonlinearity that is inherently matched to your circuit's, then all is easy; you can just put the matched nonlinearity in the feedback loop of an op amp. A typical example of this is fixing the nonlinearity of a CMOS switch by putting another one, closed, in series with the feedback resistor of the op amp, as shown in Figure 15.23a. The same trick is often used with dual optocouplers, but since they are not monolithic devices it is somewhat less successful.

Still another variation uses one op amp's input error to null out another's: the low phase error amplifier of Figure 15.23 uses the input error of one-half of a dual op amp to null out the error of the other half, operating at exactly the same gain. Assuming perfect matching, this compensation makes A2's input errors quadratic in f rather than linear, so that the closed-loop phase shift is below 0.1° for $f < 0.11f_c$ instead of $0.002f_c$ as it would be otherwise. The price is a moderate amount of gain peaking (3 dB) near f_c , and having to work at a high enough closed-loop gain that the circuit doesn't oscillate. (It's not clear that production variations and layout strays will let you reach 0.1° in real life, of course.)

15.12.6 Inverting a Linear Control

One good way of getting linearity is to use a device known to be very linear, and putting it in a feedback loop controlling the nonlinear one. An example is wide range control of LED brightness by shining part of the LED output on a photodiode and using an op amp to force the photocurrent to equal an externally supplied reference current. This is useful with voltage-to-frequency converters, photodiodes, and heaters.

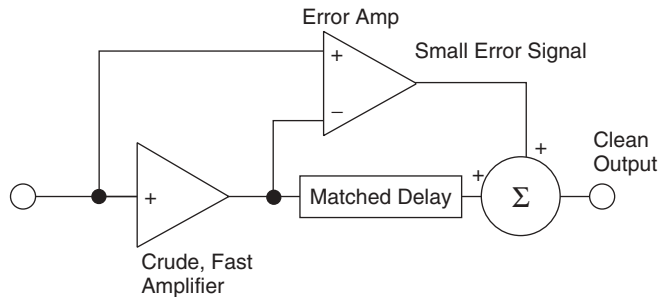


Figure 15.26. Feedforward system.

15.12.7 Feedforward

Feedback control is very useful, but it has its limitations, mainly caused by the unavoidable phase shifts and the attendant limitations on bandwidth, loop gain, and transient response. Open-loop control doesn't have these problems to the same degree, but it is inaccurate. The solution to this dilemma is *feedforward*, shown in Figure 15.26. The error signal is applied (in open-loop fashion) to the output of the amplifier instead of its input, with a suitable delay inserted to make the uncorrected signal and the correction sync up in time.

An example of this is the beam steering and focusing electronics in synchrotron accelerators, to control both systematic and random position and focus errors in the beam. The particle beam consists of individual bunches of particles going round in a circle. Bunch N , which has some position error, goes round the loop at almost the speed of light and can't be slowed down; however, the error signal goes across a chord, whose shorter path length compensates for the limited speed of the electronics and magnets—the correction arrives just in time to fix the error on bunch N , the same one that the measurement was done on. This is an excellent application of feedforward; a feedback design would have used a measurement on bunch N to control bunch $N + 1$, which in the face of random packet-to-packet errors is obviously inferior. This example is not pure feedforward, because the same bunch will come round to the sensor again, be measured again, and corrected again, which is a feedback process. If the system behavior drifts a bit, a supervisory program that occasionally estimates and tunes the feedforward parameters will usually do a good job.

15.12.8 Predistortion

Another feedforward-like technique is *predistortion*, in which the nonideal behavior of some amplifier or transmission medium is compensated in advance. A filter can boost high frequencies to compensate for dispersion and high frequency rolloff in high speed cable links, for example.

15.13 DIGITAL CONTROL AND COMMUNICATION

You'll spend a significant amount of design time just figuring out how to make all the data and control bits twiddle at the right times, so make it easy on yourself. The control and

communications tasks are not normally the places you should spend all your effort—they may influence the price of your instrument, but not the specifications. The general philosophy here is to use PCs (or single-board computers running DOS) to control instruments and network with the outside world, and build small amounts of special-purpose hardware to do the things the PC is bad at (e.g., sequencing). If your instrument is going to be built in tens of thousands, it starts making sense to use embedded microcontrollers, but in that case you'll have help (and need it, too, unless that's your field).

Real-time operating systems for embedded processors exist, but you don't want to open that can of worms unless you really have to.

For communication between the PC and instrument, choose (best to worst) a bidirectional printer port, an RS-232 port (if you can find either of them nowadays) a purchased bare-bones USB or TCP/IP communications module, a frame grabber, a data acquisition card, or a custom designed plug-in card.

15.13.1 Multiple Serial DACS

One of the pleasant consequences of Moore's law is a steady improvement in DACs and ADCs. One example is the octal serial DAC, a \$2 device that replaces eight pots and allows software control. There are also EEPots, which are multiplying DACs with EEPROMs remembering the setting, and are good for infrequently required adjustments. An octal DAC and a really good understanding of what your system does can make automatic tweaking of the operating parameters very attractive.

It is a simple matter to connect a serial DAC, or a string of them, to a PC parallel port, and twiddle bits to load them (you can make up your own data protocol with complete freedom on parallel ports, because every bit is under your control). Make sure you connect the data out pin back to the port, and verify that the correct data got loaded by reading it back.

15.13.2 Data Acquisition Cards

The author must make a confession: though he has bought the occasional data acquisition card over a span of 20 years or so, he really hates them. It isn't that there are no good ones, but the vast majority are really suitable only for undemanding low speed use. The main problems are lack of time coherence, short product lifetimes, lack of detailed schematics, and poor quality Windows software that hides bugs. The combination leads to flaky measurements and reasonable-looking wrong answers. There are honorable exceptions, but they are few, and the many bad ones have sowed confusion through two generations of experimentalists.

15.13.3 Nonsimultaneous Sampling

The sampling theorem states that we can reconstruct a band-limited function exactly from equally spaced samples. The key is "equally spaced." Most N -channel A/D cards have one actual sampling digitizer and an $N:1$ multiplexer. This leads to lots of flexibility and great-sounding specifications at low cost, but rarely to good measurements. Say you're sampling four inputs at 10 kHz. What is the time relationship between them? You don't know. Some cards take the samples as close to each other as they can, but others space the acquisitions out equally over a sampling period, regardless of what the period is,

so you can't improve the time coherence even if you slow to a crawl. If you use the manufacturer's software, you will get them plotted as though they were really coincident in time—in other words, the software constantly tells you lies about your data.

It's especially bad if there are high frequency components in the data you're sampling. You might have a measurement scheme such as subtraction or division that ought to be very stable regardless, but even small sampling time errors will function like a delay discriminator (see Section 15.5.12) and bring those out-of-band components right into your measurement where you don't want them.

Besides the multiplexing, how regular is the sampling, really? PC sound cards are quite good at this, and more modern A/D cards often have hardware timing as an option, but it isn't automatic.

15.13.4 Simultaneous Control and Acquisition

Most instruments don't just sit there and take data, like an old-fashioned chart recorder. They go somewhere, take data, move, take data, move, recalibrate, check the motor speed, take data, . . . , with control commands and data acquisition interleaved. They have to be interleaved correctly, too, observing all the timing restrictions, *especially* the requirements for perfectly periodic sampling and of simultaneous sampling of data that are to be combined later.

Doing this well, that is, with coherent timing, is usually hard with ready-made analog I/O cards, and finding out whether the settings you picked are actually doing what you think they are is even harder. If the card depends on the PC for any of its timing, interrupt latency and multitasking will completely disrupt the show by introducing delays of tens of milliseconds at unpredictable intervals. This makes many kinds of measurement impossible. Pretty-face GUI programs—even very expensive measurement suites—make this even worse, because there's no way to find out what is really going on under the covers.

To be sure of good performance, it's usually better to wire up something simple yourself, so you know exactly what it's doing and why. If you're building a commercial product for volume production, choose a major microcontroller family and write the code yourself. You can sit the processor in a housekeeping loop, and use one of the internal timers to interrupt the processor at each data point. This can be done surprisingly fast, and the timing will be very good (with a few flip-flops sprinkled around for resynchronizing, it can be very good indeed).

If you're building lab stuff, relax. Just use a couple of PLDs and maybe a FIFO buffer such as a 74ACT7804 to make a finite state machine. It'll effortlessly go 500 times faster than a PC-based solution, have very low jitter, and not bring your PC to its knees. Alternatively, consider using a small MCU module such as those from Parallax, perhaps with a bit of hardware assist, or a single-board computer running DOS. Either way, once you get good at it and have a drawer full of parts, it only takes a day or so to reconfigure the hardware. That's a good deal if your experiment is going to last awhile.

Put the complicated stuff like detailed housekeeping, communication, and data reduction on the PC side, where you can write it using a well-upholstered C or C++ compiler and debugger. Another approach is to use a flash EPROM with a counter generating addresses sequentially; you basically just store the whole timing diagram of the system in the flash, and the counter steps you through it. You may need to resynchronize the

output of the flash memory with a register (e.g., a 74HC374) so that it only changes state synchronously with the clock edge. An in-circuit programmable flash will allow you to control the sequence on the fly from your PC. This is a poor man's arbitrary waveform generator. For more complicated things, you can use a small single-board computer running Linux, but that involves a bit longer learning curve for most people. Still, for \$300 or so you get a pretty whizzy set of tools—ports, flash, Ethernet, video, and lots of memory and CPU power, all with open-source tools.

A bidirectional printer port (just about all printer ports are bidirectional, if you can still find one) can be used for all manner of control and data acquisition tasks. All the control and data lines are software programmable, so you can make up your own rules as you go along (you need a small library to allow your code to twiddle the hardware directly, but you can download several). This is really a sweet solution to the state machine problem. Of course there are also FPGAs if you're good at that.

15.14 MISCELLANEOUS TRICKS

15.14.1 Avalanche Transistors

There are lots of situations where you have to make something switch quickly with minimal jitter. If the voltages involved are small, say, less than 2 V, you can generally do it with logic parts such as fast CMOS or picosecond ECL. When higher voltages are involved, e.g. switching longitudinal-mode Pockels cells or driving a high speed magnet, and ordinary power MOSFET switches aren't making it, there are three ways of proceeding: thyratrons, spark gaps, and avalanche transistors. Of the three, spark gaps are simple and fast, and can stand high voltages, but require a lot of maintenance and have poor jitter, because the initial few ionization events are stochastic; thyratrons are powerful but slower; and avalanche transistors are lower powered, but very fast and much more repeatable than the others. Avalanche-rated transistors are available from Zetex, along with a lot of good application notes. Their devices tend to be quite a bit slower to avalanche than the old standby 2N2369, which can easily do 300 ps with picosecond jitter (see Figure 15.27). Interestingly, it's the old, slow, diffused-junction parts that avalanche fast. You can also get few-kilovolt pulses as short as 100 ps from the reverse recovery of rectifier diodes, if you hit them hard enough.[†]

15.15 BULLETPROOFING

You have to expect failures. It's critically important to design your system so that it can fail without melting down. For example, make sure your power supplies have well-defined current limits, that there's a thermal cutout on all power circuits and a limit switch on all mechanical motions. Most of all, make sure as far as possible that no hardware damage or human injury can occur as a result of software failures, turn-on transients, or any conceivable abuse of the inputs and outputs. It is proverbial that you can't make something idiot-proof, but you can at least make it idiot-resistant.

[†]See the publication of I. V. Grekhov et al., High-power subnanosecond switch. *Electron. Lett.* **17**, 422–423 (1981).




Figure 15.27. Avalanche transistor flat-top pulse generator. When Q1 avalanches, the upper transmission line discharges a constant current into the lower one. With a sufficiently light load, the open-circuit reflection doubles the voltage at the far end of the lower line. You can also play other transmission-line games, such as using a balun to get higher output voltage, or increasing the emitter resistor to get higher peak voltage at the expense of worse reflections.




Figure 15.28. This current driver will blow up the diode laser when the wiper of R1 opens momentarily.

Example 15.2: How to Kill a Diode Laser with a Pot. At the risk of getting somewhat ahead of ourselves, consider Figure 15.28. The LM317LZ is a cheap but good voltage regulator. It keeps a constant 1.25 V drop between its output and adjustment terminals. Putting a voltage divider between the output and ground, with the adjustment pin connected at the middle, makes a positive voltage regulator adjustable from 1.25 V up. It works quite well for this, since the adjustment pin sources a small and reasonably constant 50 μ A. Here it is used as a not-too-great constant current source (far above shot




Figure 15.29. Provided that R_2 is chosen correctly, this diode will survive failure of R_1 . R_3 is there to handle most of the current, rendering R_1 more reliable as well.

noise); it keeps a constant 1.25 V across R_{1a} , thus keeping the diode current constant at $I_{\text{diode}} = 1.25\text{V}/R_{1a}$. (This circuit may have problems on startup, but we'll ignore them for now.) The circuit works fine unless the wiper of R_1 opens momentarily (100 μs is long enough). If that happens, the adjustment pin can't source its 50 μA , and the output rises to $V_{\text{in}} - 2\text{V}$ or so, destroying the diode laser instantly.

A better way to do this is shown in Figure 15.29. Here if R_1 's wiper opens, the diode current rises only to a maximum value set by R_2 . R_3 is there to take most of the current, which otherwise would come through R_1 's wiper. Another advantage of this scheme is that you can't blow up the diode by turning R_1 too far. Remember: pots do get noisy eventually, and you don't want your instrument to die on the first blip.

15.15.1 Hot Plugging

The most common cause of death for prototypes is losing ground or supply connections. Lethal scenarios depend on how many supplies of each polarity you have, and whether there are any loads connected across the supplies (e.g., heaters or motors). You have to be systematic about what can happen, by checking each possible order in which the supply and ground connections can open. The actual fatal fault is always either transient overvoltage or supply reversal. If the person doing the plugging is ham-handed about it, it is also possible to short a supply pin to somewhere and blow up the on-card regulator or the somewhere.

It is rare and usually stupid to connect a load between supplies of the same polarity, for example, a load running on the 7 V between +5 and +12. If the load is a heavy one, then if whatever uses most of the +5 gets unplugged, the +5 V line will get dragged away upwards and probably blow something up. It is less rare and more intelligent to connect a load (e.g., a motor or heater) between positive and negative supplies, and this is also easier to protect against. When one supply is lost, the other one will try to drag it through ground and beyond, which is bound to kill something. Fix this with Schottky rectifiers between each supply lead and ground (reverse biased, of course). That way they can't be dragged more than about 0.4 V in the wrong polarity.

If ground is lost first, the usual failure mode is that one supply draws much more current than the other (positive, usually), and so the "ground" lead gets dragged rapidly toward the positive rail as the positive supply's bypasses get drained down. This puts the whole supply range (30 V for a $\pm 15\text{V}$ system) across the negative rail, which is

often enough to blow something up. A transient-suppressing Zener diode across the lower current supply fixes this. The Zener should be able to take the full supply current of the opposite supply, plus discharge all its bypass capacitors. ON Semiconductor sells good ones cheaply—check out the P6KExxxx series.

Otherwise, the main failure modes are that the loss of ground before the signal leads causes big voltage excursions on the signal connections, sometimes blowing up logic parts. Low impedance signal connections (i.e., those without $>1\text{ k}\Omega$ resistors in series with them) should have some sort of protection, such as diodes to the power supplies.

Plugging in has some dangers as well. Most of them are the same as for unplugging, but there are one or two unique ones. Boards are often plugged in slightly cocked, so that one end of the connector mates a little before the other, leading to the supplies and signals coming up in different orders and at different times. Some voltage regulators don't start under weird loads, especially foldback limiters and 7800 series devices, which refuse to start up if their inputs are brought below ground. When the supply contact is suddenly made, all the bypass caps charge up very rapidly, requiring a lot of current. If the supply contact doesn't get made soon enough, this current may wind up coming through the signal leads. A few $100\ \Omega$ resistor arrays, wired in series with your signal inputs and outputs, can save headaches here. This also protects inputs and outputs from damage when connected but unpowered, which happens a lot in prototypes and is especially nasty because they're so precious and usually harder to repair. If you're using card edge connectors, it's possible to arrange the pads so that the ground connections are made first, which helps a lot. On a multipin connector (e.g., DIN Eurocard), you can put grounds at both ends of the connector instead. This isn't as good, because small amounts of crud on the pins can cause delayed contact, but it helps.

15.15.2 It Works Once, How Do I Make It Work Many Times?

This is a really generic problem, and a major source of difficulty in putting a design into production. There are two basic causes: first, it is a corollary of Murphy's law that all prototypes turn out to contain the normally unobtainable perfect component (at least one, sometimes several). Lenses with no decentering, transistors with infinite beta, op amps with no offset current, AR coatings that really have zero reflectance at your wavelength, the list goes on. Prototypes invariably seem to depend critically on these perfect parts, leading to pain in replication. Second, most R&D people (including the author) have a tendency to work on the irritating problems until they're no longer irritating, then stop. This is fine for graduate study, but a disaster in a technology development or transfer situation. The problem is that faults which do not irritate the designer may render the instrument unusable by others. This leads to difficulty handling the system off, either to customers or to a product division, and so to heartburn on all sides. Early customer involvement is a good way to avoid this problem. Complex systems that really work are invariably found to have developed out of a close interaction between one or a few gifted designers and a small, highly involved community of users.

The users involved should include a few querulous curmudgeons, too, but it is best to pick ones who are unlikely to be loud and public in their criticism. If the users are given work to do, their level of commitment increases, because they gain a feeling of ownership. That way, they'll be more likely to stick by you when the project gets in serious trouble for a while (most successful projects do this at least twice).

15.15.3 Center Your Design

The perfect component problem is best handled by simulation if possible. The author is very sceptical of the uncritical use of simulations, but one place they are frequently indispensable is in making sure that the design is *centered*—that its components are chosen in such a way that normal component variations are allowable on both sides of the nominal design. It is common to design a circuit or instrument in such a way that (for example) a 5% shorter focal length is a disaster, leading to an inability to focus at infinity, while a 5% longer one is no problem, leading only to a barely noticeable difference in the closest focus distance. This will hurt when your lens supplier ships a bunch of slightly short lenses, so if the specified tolerance is $\pm 5\%$, the nominal focal length should be made slightly longer to reduce the danger. Simulation can help in avoiding this sort of trouble.

In a mixed-technology system, getting simulations right is considerably more difficult than it is in lens design or digital circuit design. Because of the labor involved, it should be used judiciously.

Electronic Construction Techniques

Experience is a hard teacher because she gives the test first, the lesson after.

—Vernon Law[†]

16.1 INTRODUCTION

Advanced mixed-signal electronic hardware can be designed on a computer and taken straight to multilayer surface-mount circuit cards without prototyping. Then—unless you're a Certified Wizard—it will fail to work and have to be cut apart with a Dremel grinder, laboriously debugged, partly redesigned, and laid out again.

This chapter is mainly about electronic prototyping and the reduction of a working prototype to a working PC board, with specific application to high frequency, high dynamic range, mixed-signal circuitry. We'll talk about laying circuits out intelligently, calculating the circuit board strays, assembling breadboards, and tuning *LC* filters. By the end, you should be equipped to build what you've designed, without suffering from avoidable layout and assembly bugs such as ground loops. You'll need to be comfortable with test equipment, know how to solder (preferably blindfolded), and have basic prototyping skills such as stuffing boards, wiring connectors, and drilling boxes.

16.2 CIRCUIT STRAYS

Simulations often don't match up with experimental reality, due to modeling errors and circuit strays. Strays arise from Maxwell's equations and the known properties of well-behaved conductors and dielectrics. We may not know the exact boundary conditions well enough to produce a perfect model of the self- and mutual inductance and capacitance of everything in the circuit, but we can usually take a good first cut at it; since we hope that their effect is not dominant, that's often enough. It doesn't replace prototyping, but it helps a lot in reducing the prototype to a circuit board, where the strays are usually much worse.[‡]

[†]*This Week*, August 14, 1960.

[‡]It is surprising to many people that a small circuit card can have worse strays than some big ugly prototype with wires in mid-air, but it usually has.

As a general philosophy, if you're using circuit simulators, for any critical analog functions such as front end amplifiers, you should check the simulation against the prototype, and adjust both until they match and you know why they match, and you can explain it to a sceptical colleague. Routine stuff can be designed, simulated, and laid out without always needing a prototyping step, especially with the rise of rapid-turn PC board websites, which often have their own CAD software you can use.

16.2.1 Circuit Boards

Most PC boards are made of epoxy-impregnated glass cloth, type G-10 or FR-4 (a flame-retardant G-10 variant). It's not a terribly well-controlled material— ϵ_r varies between 4.3 and 4.6, but it's cheap, strong, and reasonably stable. In North America, the copper thickness is conventionally specified in ounces per square foot and comes in 0.5, 1, and 2 oz weights (thicknesses of 0.017 mm, 0.035 mm, 0.07 mm, respectively).

If you're wise, you'll make a habit of building production circuits on ground plane PC boards exclusively; the extra fabrication cost is more than repaid by stable, predictable performance. Mixed-signal boards (those with both analog and digital circuits) must always be built on ground planes, and if you don't use a ground plane with a board containing ADCs and DACs, it will quite definitely fail. We'll specialize our discussion of circuit cards to the ground plane case, where all traces are really microstrip transmission lines.

16.2.2 Microstrip Line

We've already encountered microstrip line, and since we know how to calculate the impedance of a piece of mismatched transmission line, we have a completely general solution that can be used with traces of any reasonable length.

As shown in Figure 16.1, the impedance of microstrip line of width W and negligible thickness, with dielectric of thickness h and dielectric constant ϵ , is approximately[†]

$$Z_0 \approx \frac{377}{\sqrt{\epsilon}} \frac{h}{W} \frac{1}{1 + 1.735\epsilon^{-0.0724}(W/h)^{-0.836}}, \quad (16.1)$$




Figure 16.1. Microstrip line, that is, a metal trace over a ground plane.

[†]Ralph S. Carson, *High Frequency Amplifiers*. Wiley, Hoboken, NJ, 1982, pp. 78–79.

and the propagation speed $c' = c\lambda/\lambda_0$ is

$$c' = c \frac{\lambda}{\lambda_0} \approx \frac{1}{\sqrt{1 + 0.63(\epsilon - 1)(W/h)^{0.1255}}}, \quad \frac{W}{h} > 0.6$$

$$\approx \frac{1}{\sqrt{1 + 0.6(\epsilon - 1)(W/h)^{0.0297}}}, \quad \frac{W}{h} < 0.6. \quad (16.2)$$

If the trace thickness t is nonzero, we can adjust the width to take this into account, replacing W with W_{eff} , where

$$W_{\text{eff}} = W + \frac{t}{\pi} \left[\ln \frac{2h}{t} + 1 \right]. \quad (16.3)$$

These formulas look ugly, but they're pretty accurate; anyway, since you've probably only got three different widths of trace on your board, calculating Z_0 and c' isn't onerous. You can ask your board vendor the width of a 50- Ω line in your geometry, and what their guaranteed impedance tolerance is.

Because of the asymmetry of microstrip line, its impedance and propagation speed are both frequency dependent, even in the absence of material dispersion. The low frequency limit resembles a TEM mode with $\epsilon_{\text{eff}} \approx (\epsilon_r + 1)/2$, whereas in the high frequency limit almost all the field is inside the dielectric, and $\epsilon_{\text{eff}} \approx \epsilon_r$.

Aside: Hook. Ordinary FR-4 or G-10 board is a reasonable high frequency dielectric; however, its dielectric properties are too dispersive for building really good wideband amplifiers. This effect is known as *hook*, due to the shape of the ϵ versus f curve. If you need really wide bandwidth, say, DC-1 GHz with good time-domain response, choose a different material such as alumina or Teflon-glass. Microstrip is slightly dispersive above 1 GHz anyway, but is quite acceptable on a good substrate. Boards are available with ϵ from 2.2 to 10.

16.2.3 Inductance and Capacitance of Traces

To calculate the inductance and capacitance of traces, we use the equation for the impedance of a mismatched transmission line. For lines whose electrical length $\theta \ll 2\pi$, we can take the low frequency limit of (14.13) using a shorted line to get the inductance and an open-circuited one for the capacitance. We find that a length x of line of propagation velocity c' and characteristic impedance Z_0 has low frequency inductance and capacitance of

$$C = \frac{x}{Z_0 c'} \quad \text{and} \quad L = \frac{x Z_0}{c'}. \quad (16.4)$$

Note that this assumes that x is much longer than W and h ; otherwise the strays at the ends of the line become dominant. Thus we can't use this simple idea for pads, only traces. Also, the temperature coefficient of the dielectric constant of FR-4 board is almost 1000 ppm/K, which is really inconvenient sometimes.

Example 16.1: Inductance and Capacitance of PC Traces. Consider the PC board trace of Figure 16.2, which is built on 0.33 mm thick FR-4 board with a ground plane. The line




Figure 16.2. PC trace.




Figure 16.3. Circuit model.

is a 7 mil (0.18 mm) wide trace of 1 oz copper, which has $Z_0 = 91.4\Omega$ and $c' = 0.57c$. Its total length is 3 cm, and it has a pad measuring 0.8×1.3 mm at each end and another one in the middle. The parallel-plate capacitor formula isn't adequate for the pads, so we have to do a few measurements of the actual capacitance, and we find that it's 0.2 pF per pad. By (16.4), the two bits of trace look like 8.0 nH of inductance and 0.95 pF of capacitance. This model is only valid far below resonance, so to leading order it doesn't matter where we put them in the model; since the circuit is symmetrical, we'll add them symmetrically, so our final low frequency model is as shown in Figure 16.3.

16.2.4 Stray Inductance

We saw in Section 14.3.7 that a straight b centimeter length of round wire of diameter d centimeters in free space has an inductance of about

$$L = (2 \text{ nH/cm}) \cdot b \ln \left(\frac{2b}{d} \right). \quad (16.5)$$

This is a special problem in low impedance circuits (e.g., 1 inch of #22 wire (22 nH) will have a reactance of 50Ω at 360 MHz) and in power supply wiring, where the supply lead inductance often resonates with the bypass capacitors, making a mockery of our supposedly well-decoupled supply. Inductance is reduced by shortening the wire and by running it closer to the ground plane.

A single-turn circular ring of diameter D made of the same wire has an inductance of[†]

$$L = (6.3 \text{ nH/cm}) \cdot D \left(\ln \frac{D}{d} + 0.08 \right). \quad (16.6)$$

Thus a 1 meter diameter loop made of the shield of RG-58A/U coax (5 mm diameter) has an inductance of only $3.3 \mu\text{H}$, an interesting fact that we'll come back to in Section 16.5.2 in connection with ground loops.

[†]Bureau of Standards Circular No. C74, US Government Printing Office, Washington, DC 20402 (1924).

16.2.5 Stray Capacitance

Stray capacitance comes from two places: the self-capacitance of objects and the mutual capacitance between the object and other objects, mostly ground. For small objects, the self-capacitance becomes important, because it goes as the radius whereas mutual capacitance goes as the area.

A large pad over a ground plane with a thin layer of dielectric constant ϵ has roughly the parallel-plate capacitance,

$$C_{pp} = (1.12 \text{ pF/cm}) \frac{\epsilon_r A}{4\pi d}. \quad (16.7)$$

This formula considers only the region between the plates, and neglects fringing fields. It thus gives an underestimate, which becomes progressively more serious as the pad gets smaller; for a 1 mm square pad and 0.3 mm G-10 dielectric (typical of four-layer surface mount boards), it's too low by a factor of 2, which is somewhat surprising initially.

The explanation is the image effect from undergraduate E&M: a charge q above a conducting plane effectively generates an image charge $-q$ the same distance below the ground plane (Figure 16.4). Thus that 1 mm pad and ground plane have the same field configuration as two 1 mm pads, separated by 0.6 mm (now the pad looks a lot less infinite), and a major fraction of the fields are outside the pad area. The fringing capacitance will tend to go as the perimeter of the pad, whereas the parallel-plate capacitance goes as the area.

As a corollary, the area of ground plane contributing significantly to the pad capacitance extends about half a dielectric thickness past the edge of the pad, so that if you need to chop a hole in the ground plane to accommodate a high impedance, high frequency node, you need to make the hole a bit bigger than the pad.

You can make a guess that, on a surface mount board, the stray capacitance is about 0.15 pF times the number of components attached to that node. It's somewhat worse in through-hole boards. The other thing to remember is the self-capacitance of large objects; a 1 cm radius sphere has a self-capacitance of 1.12 pF, besides mutual capacitance with everything else. Fast, high impedance circuits have to be kept physically very small.




Figure 16.4. Electrostatic field lines for a pad over a ground plane.

16.2.6 Measuring Capacitance

It is possible to measure the capacitance of a 1 mm pad, but it isn't too easy. A better method is to make a 100× scale model, whose capacitance will be 100× larger, and measure that instead. At 10 cm on a side, it's easy to handle, and its capacitance will be a convenient 20 pF or so, rather than 0.2 pF. You can build it out of G-10 board or heavy-gauge aluminum foil and cardboard. Use air for the dielectric and multiply the measured capacitance by ϵ , or for more accurate results, dunk the model in oil of known ϵ to represent the dielectric.

This method is accurate and very general—you can use it to find the capacitance between some horrible stack of pads and vias and the ground plane, which is difficult with any other method. It takes a bit of work, but you only have to do it once for each pad stack design, and usually a judicious use of scissors can turn one model into another pretty fast. Alternatively, there are free programs available on the Web to automate this, but it's much more work than a quick measurement. Whichever way you do it, you'll rapidly discover that the parallel-plate capacitance is a wild underestimate if the separation is anything like the pad dimensions.

16.3 STRAY COUPLING

16.3.1 Capacitive Coupling

As we saw in Section 14.4.5, the fields of a balanced line fall off more rapidly laterally than those of a single wire, and the scale of the falloff is the distance between the conductors. Because of the image effect, wires near a ground plane have the same lateral confinement of the fields, and bringing them closer to the ground plane increases the confinement.

Thus two traces above a ground plane have much less mutual capacitance than the same traces alone. (That's one of the ways that ground planes help.)

16.3.2 Inductive Coupling

The mutual inductance of two identical microstrip lines whose self-inductance is L , separated by a distance s , is roughly[†]

$$M = L \left(\frac{1}{1 + s^2/h^2} \right). \quad (16.8)$$

16.3.3 Transmission Line Coupling

Two transmission lines in close proximity will couple together due to the overlap of their field patterns, in just the way we saw in Section 8.7.1. Transmission line coupling is only distinguishable from the other two when you've got a long run of line, but if you have, terminating the line on each end with Z_0 will send the forward and reverse wave coupling to the front and back end termination (this is how directional couplers work).

[†]See, for example, Johnson and Graham.

16.3.4 Telling Them Apart

Inductive coupling depends on the direction of the current flow, whereas capacitive coupling depends only on the voltage swing. Weak stray capacitive coupling is high- Z and moves the whole trace up and down together; weak inductive is low- Z and moves one end of the trace with respect to the other. Thus one way of telling whether your coupling is capacitive or inductive is to open-circuit the far end of the lead; inductive coupling will go away, and capacitive coupling won't. Similarly, short-circuiting one end to ground will get rid of the capacitive coupling but preserve the inductive.

If you have a combination of inductive and capacitive coupling, you can sort them out by changing the direction of the source current; the inductive coupling will change sign whereas the capacitive won't.

16.4 GROUND PLANE CONSTRUCTION

16.4.1 Ground Currents

Kirchhoff's current law (aka charge conservation) says that the net current flowing into any node of the circuit is always 0, so currents flow in loops; for every current flowing in a wire, there's an equal and opposite current flowing back again somewhere else. Though elementary, this is often overlooked.

Most of the currents we think about flow in wires, so that their routes are predefined. Where there are multiple routes (e.g., paralleled wires or a ground plane), however, we need to be a bit more careful. Whenever there is more than one path, the current divides itself according to the admittance of each path; at DC, this means that the current through each path is proportional to its DC conductance G . At AC, the resistance of the path rises slowly with frequency (as \sqrt{f}) due to skin effect, but the reactance rises linearly, since $X_L = j\omega L$. Accordingly, inductive effects become dominant at higher frequencies.

The inductance goes up as the loop gets bigger, so as f rises, the current tends more and more to return directly underneath the trace. In building circuits over ground planes, we must be careful not to obstruct these ground currents by cutting holes in the ground plane for other conductors. Even a narrow gap forces the current to go around, which adds significant inductance to the ground path, as well as leading to EMC vulnerability (see below).

I once turned a beautiful 2 GHz amplifier into a 400 MHz bookend by deleting the ground plane and thereby effectively placing large inductors in the circuit.

—Steve Roach[†]

16.4.2 Ground Planes

A ground plane provides very low impedance paths from everywhere to everywhere else, so that different parts of the circuit agree pretty well on just what ground is, and (through field confinement) provides isolation between different parts of the circuit.

Its main drawback is increased stray capacitance to ground. In most places in a circuit, this isn't too horrible if you plan for it.

[†]Steve Roach, Signal conditioning in oscilloscopes and the spirit of invention, in Jim Williams, ed., *The Art and Science of Analog Circuit Design*. Butterworth-Heinemann, Boston, 1995.




Figure 16.5. Minimizing capacitive loading by minor circuit changes (bandwidth-limiting meshes shown in bold lines).

Example 16.2: Stray Capacitance in an Emitter Load. The common-emitter amplifier of Figure 16.5 has a series RC circuit from emitter to ground, to set the AC gain of the stage. With the capacitor on top, in parallel with the bias resistor, there are five pads and three traces to provide capacitance from E to ground; with 2.5 mm of 7 mil width per trace (0.33 mm dielectric) that’s about 0.16 pF per trace, for a total of 1.5 pF. If we split the bias resistor with the bypass instead, most of this appears harmlessly in parallel with C , so there are only two pads and one very short trace, about 0.5 pF. The common-emitter stage will exhibit a rising gain with a 3 dB point at 140 MHz in the first case, but 420 MHz in the second.

16.4.3 Relieving the Ground Plane

Sometimes, though, you can’t stand that extra 0.5 pF (e.g., at the collector of a cascaded photodiode amplifier; see Section 18.4.4), and you have to chop small holes in the ground plane under a few pads. This is okay as long as you don’t make ground currents detour even slightly.[†] You have to make the hole about one dielectric thickness bigger than the pad because of fringing. If the ground plane is near one side of the board, you can often do additional relieving with a milling machine or Dremel tool; failing that, lift the pads and run the offending circuitry in mid-air, dead bug fashion.

16.4.4 Skin Depth

Down at DC, current in a long piece of ordinary imperfect conductor flows uniformly throughout its cross section. At AC, the changing current produces an AC B field, which

[†]In extreme cases, you can bridge the hole with jumpers, for example, 0 ohm resistors.

by Faraday's equation,[†]

$$\nabla \times \mathbf{E} = -\frac{1}{c} \frac{\partial \mathbf{B}}{\partial t}, \quad (16.9)$$

produces an AC E field inside the conductor. Since it *is* inside a conductor, this E gets mostly shorted out, preventing B from changing much, which in turn suppresses the AC current in the bulk of the conductor. The result of all this is that the current flow is progressively confined to the surface of the conductor as f increases, and distributed so that $B = 0$ deep inside (for a round wire, that means that the current density J is independent of angle). The finite conductivity of real wire means that these effects occur over a finite thickness of the conductor; at lower frequencies, $\partial B/\partial t$ is smaller, so the induced E is smaller. In circuit terms, this means that the source impedance of the induced E is smaller at low frequencies, and so it takes more conductor thickness to short it out effectively. Since these effects are linear and homogeneous, we expect some sort of exponential falloff of J with depth, and that's in fact what we get.

An infinite half-space ($z > 0$) of imperfect conductor has a current density J that falls off with depth as $\exp(-z/\delta)$, where δ is the skin depth,

$$\delta = \sqrt{\frac{\rho}{4 \times 10^{-9} \pi^2 f \mu}} \quad (16.10)$$

(in centimeters) and ρ and μ are as usual the electrical resistivity ($\Omega \cdot \text{cm}$) and relative magnetic permeability. This applies both to conduction currents from applied voltages and to eddy currents caused by ambient magnetic fields.

"Sufficiently highly conducting," in the section just above, means that the skin depth in the ground plane is much less than the thickness of the ground plane or of the dielectric. That points to where ground plane construction breaks down—low frequency magnetic coupling. For copper ($\mu = 1$, $\rho = 1.7 \mu\Omega \cdot \text{cm}$), δ is $7 \mu\text{m}$ at 100 MHz but 0.85 cm at 60 Hz.

The presence of a shield will strongly suppress the B field nearby, but even neglecting that, a shield 7δ thick will give >60 dB isolation, and 10δ will give >85 dB. Thus a double layer of 2 oz copperclad board (0.14 mm total copper) will give >60 dB isolation anywhere above 10 MHz, but progressively less below there; that isn't too much of a problem, though, since below 10 MHz we're usually using magnetically shielded coils anyway, to get high enough inductance without massive copper losses, and that together with the lower frequency reduces $\partial \mathbf{B}/\partial t$ substantially in the surrounding circuitry. As a practical matter, you'll probably never see an inductive coupling problem in a circuit above 1 MHz if you use tall, continuous copper-clad shields; if you're worried about it, use ordinary sheet steel instead (note the factor of μ in (16.10)).

Aside: Split Grounds and Electromagnetic Compatibility. If you physically split your ground plane into "analog" and "digital" regions, you've made what is called a "patch antenna." It's really amazing how well structures like that couple electromagnetic interference (EMI) in and out. Commercial products have to pass stringent electromagnetic compatibility (EMC) tests, and split ground planes are one excellent way to make yours flunk. Do the layout carefully on a single featureless ground plane—vias and a few very small cutouts to relieve a pad are OK, but no big holes and *no slots*.

[†]Okay, we do need one Maxwell equation in this book, but we got pretty far without it.

16.5 TECHNICAL NOISE AND INTERFERENCE

16.5.1 What *Is* Ground, Anyway?

The first *koan* of prototyping is: When is ground not ground? (Appropriate Zen answer: always and never.) Ground is a very useful approximation, but that's all. All conductors have resistance and inductance, so as soon as your circuit gets turned on and current starts to flow, all your grounds will be at different voltages. This isn't necessarily obnoxious, but sometimes it is; you have to watch, or those voltage drops in your grounds will bite you. Typical problems include ground loops at low frequency and ground pigtailed at high frequency.

16.5.2 Ground Loops

Ambient magnetic fields produce voltage swings in conductors, and multiply-connected grounds (*ground loops*) will allow them to produce large ground currents. Remember that 1 m loop of RG-58A/U shield ($3.3 \mu\text{H}$) in Section 16.2.4 that we were going to come back to? With a 1 milligauss field at 60 Hz, it produces $30 \mu\text{V}$ of signal with a reactance of 0.0012Ω ; the shield has a resistance of around 0.02Ω . This impedance is lower than almost anything—it's the impedance of a big car battery. If the coax had been RG-214 (half-inch coax with two layers of braid), it would be a dozen car batteries in parallel.

Low voltage plus extremely low source impedance make ground loops do unintuitive things, and this is what makes them resemble gremlins—we don't usually expect to encounter a millivolt-level signal with a circulating current of hundreds of milliamps. Once you realize that the source impedance of the ground loop can be smaller than the contact resistance of the coax jack, it becomes obvious why wiggling the cable sometimes causes the hum to change by 10 dB.

This low impedance can easily cause a voltage drop across a chassis or ground plane, which is one source of the trouble. Even more important, magnetic coupling will induce the same voltage in the shield and the center conductor; shorting out the induced voltage in the shield while leaving the center conductor alone turns the common-mode swing into a differential swing that the amplifier sees as signal. Fortunately, it isn't hard to fix once you know how it occurs; just putting a 1 ohm resistor in series with the shield will attenuate ground currents by 30–60 dB, and adding a differential receiver will help even more.

Still, we shouldn't create ground loops where we don't have to. Plug all your instruments into the same wall circuit where possible, and use tape to insulate coaxial connectors that touch each other or the optical table (the author has seen cases where a file card slid between a patch cord and the optical table made a 60 dB improvement). Grounds in house wiring and water pipes can easily have 10 m loops.

Example 16.3: Ground Loop in an AC Amplifier. Consider the circuit of Figure 16.6, a 15 cm long by 5 cm wide, 1-ounce dead bug prototype with a 1 m loop of RG-58A/U attached. A weak stray field of 1 mG at 60 Hz will induce a voltage of $30 \mu\text{V}$, with a $20 + j1.2 \text{ m}\Omega$ source impedance. The ground plane has an end-to-end resistance of $3.0 \text{ m}\Omega$, so 13% of the open-circuit voltage is dropped across the length of the ground plane. The center conductor will have the same open-circuit induced voltage, of course, but will not be shorted out by the ground plane and will not exhibit the large gradient




Figure 16.6. Ground loop example.

across the circuit. If the AC amplifier has a gain of 10^4 , and its ground connection is 3 cm away from the input connector, it will see a differential signal of $30\mu\text{V} \cdot (1 - 0.13 \cdot 12\text{cm}/15\text{cm}) = 27\mu\text{V}$, and produce 270 mV of hum at its output.

16.5.3 Floating Transducers

It's a really bad idea to put transducers (e.g., photodiodes) on cables, but sometimes there's no alternative. If one end is going to be grounded, all is easy; connect that end to the shield, and ground it at the receiver end of the cable. Sometimes, though, both ends have to be floating, and that's more difficult. The key trick in this case is to use shielded twisted pair and ground the cable shield *at the transducer*.

The reason is that, if the shield is bouncing up and down with respect to the transducer's idea of ground, capacitive currents will flow between the shield and the wires. The cable won't be exactly perfect, so they won't be quite balanced, and a differential mode signal will result. If you ground the shield at the transducer end, there is no capacitive current and hence no interference. This remains true if the transducer has a differential line driver attached to it instead of really being a floating device.

16.5.4 Mixed-Signal Boards

Mixed-signal and RF boards require special treatment, especially when high speed, high resolution A/D converters are involved. Keep the analog and digital traces far apart; fast logic slews at 1 V/ns or faster, which will couple spikes into everything, especially when several outputs are switching at once. Metal shields, continuously connected to the ground plane by solder or finger stock, can be a big help. (See Figure 16.7.)

It's often tempting to split the ground plane into analog and digital sections, the way many ADC application notes suggest. While this is one way to keep digital junk out of your analog section, it's such an EMC nightmare that you're way better off not doing it. Ground currents drop off fairly steeply from the lowest-inductance return path, so just keeping the two sections physically separate is the better way to go. If you have something really really horrible to deal with, such as the A/D clock, which can hardly be kept entirely out of the signal processing path, you can run that in coax with a separate ground.

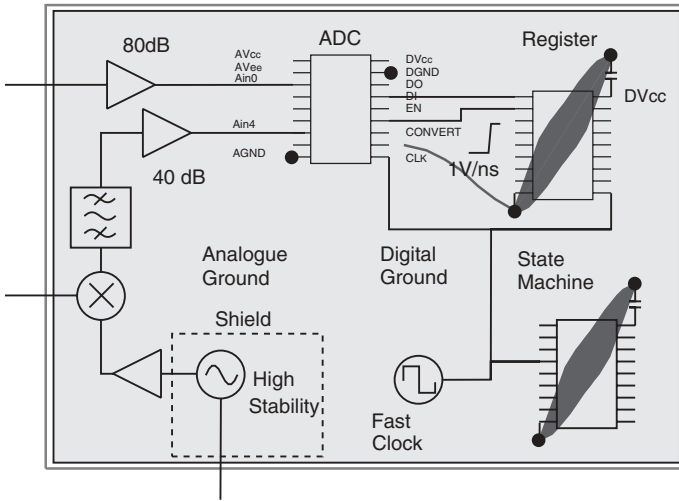


Figure 16.7. Connecting analog and digital grounds: use a single plane and separate the analog and digital sections so their ground currents don't interact.

Digital systems usually have separate power and ground planes, with densely sprinkled bypasses and the interplane capacitance to short them together at AC. Don't do this with your analog sections, unless the circuitry can be limited to the side nearest the ground plane; the supply plane always bounces a bit, and coupling between it and the traces running over it will cause you headaches (see Example 19.1). You can split the power plane, but it's usually best to have only ground planes in analog circuitry and run the power in traces.

A few things that can help quite a lot if you have analog and digital sections on the same card: put them on opposite sides of the board with two ground planes (analog and digital) in between and a shield over top, and don't use an external memory bus on your microcontroller. Serial I/O peripherals work very well for most things, and you don't have 30 wires jiggling up and down on every instruction cycle. This need not be a serious constraint, either; you can even get 64 megabit flash memories with serial interfaces.

16.5.5 High Impedance Nodes and Layout

Getting the layout of high-Z circuitry right is extremely important. Pickup is so insidious that you'll spend a long time finding it if you're not very thoughtful about layout. Life has gotten easier in one way, because SMT components pack more tightly and hence have less trace length and loop area to pick up junk, but at the same time, more circuitry is packed into a smaller space, so that the sources of pickup are that much closer. SMT isn't always a win; see Example 19.1.

16.5.6 Connecting Coaxial Cables

The ground loop example shows that connecting cables isn't as simple as it looks. One good way to eliminate ground loop problems is shown in Figure 16.8: ground the shield through a 1 Ω resistor, and use a differential amplifier or transformer to restore a local ground reference.




Figure 16.8. Isolating ground loops.

A common-mode choke, made by wrapping the coax many times around a low frequency toroid core (even a link of heavy steel chain works), is another possibility in bench setups, but expensive and inconvenient in commercial instruments. You need at least 5 mH or so to help very much.

16.5.7 Bypassing and Ground/Supply Inductance

Power supplies have to be considered as parallel resonant circuits.[†] The resonances are of several kinds, but the ones that cause problems are the parallel ones, because the supply impedance becomes high at resonance. These have the common property that what causes trouble is overkill—too much bypass capacitance and not enough resistance for effective damping of the resonance, resulting in a high- Q circuit and excessive ringing. At high frequency, supply resistance is your friend, not your enemy.

Example 16.4: Supply Bus. Consider a 150 mm long supply bus made up of a 6 mm wide copper trace, 1.5 mm above the ground plane, with 0.1 μF capacitors to ground every 25 mm. The impedance is about 30 Ω , and the velocity factor is 0.52. Figure 16.9 is a SPICE simulation of the voltage at the middle and one end of the trace, for unit AC current at the midpoint. The bypassing rapidly gets worse as ESR drops below 0.3 Ω —that is, as the capacitors get better.

The moral of the story is to be content with good enough bypassing, and to remember that killing the Q of a resonance requires putting in some loss: a series resistor, a lossy ferrite bead, or a lossy capacitor such as an aluminum electrolytic. Use a 0.1 μF monolithic ceramic capacitor on each supply pin of each IC, and put in a 1 μF solid tantalum on each supply for every 8 or 10 chips using it. When the layout comes back from the draughtsman, *then* you can eliminate the unnecessary ones. If you have ground and supply planes, the board itself will provide some nanofarads of very quick capacitance, so the need for individual bypasses is reduced.

We often need to use the same supplies for analog and digital circuits, or for all sections of a high gain or high isolation circuit. Use series resistors and ferrite beads in the supply traces between sections for high frequency isolation, and capacitance multipliers for low frequency (Example 14.1). Use a low frequency and high frequency bead together (Ferroxcube 3E2A and 4B ferrites work well), and be ready to use two or three sections if necessary.

[†]Paul Brokaw, *An IC Amplifier User's Guide to Decoupling, Grounding and Making Things Go Right for a Change*, Application Note AN-202, Analog Devices.




Figure 16.9. Resonances in a well-bypassed supply bus. Vertical scale is impedance in ohms.

16.5.8 Bypass Capacitor Self-Resonances

Bypass capacitors have a series-type self-resonance, with f_{SR} somewhere in the VHF or UHF range. From what we've seen of supply resonances, that sounds bad, but (surprisingly) it is entirely benign unless you make the lead inductance excessive. The reason for this is that a well-bypassed supply is one with very low source impedance at all frequencies, but just what the phase angle of that impedance is matters very little. The ESL and wiring inductance will eventually dominate the impedance of the bypass, but remember that parallel and series LC impedances are symmetrical functions of $\log(f)$ —if your bypass is effective at $f_{SR}/10$, its impedance will have just the same magnitude at $10 f_{SR}$, so it will be just as good a bypass there. It's parallel-type resonances you have to worry about, and you can fix those by spoiling the Q with a series resistor or lossy capacitor.

16.6 PRODUCT CONSTRUCTION

We'll only sketch the most important things to think about in constructing products, because by the time your prototype instrument works, you're over the hump.

16.6.1 Cost Versus Performance

The cost of many purely electronic systems is dominated by boards, cables, and connectors. A careful trade-off is required between flexibility, reusability, manufacturing yield, and per-piece cost.

16.6.2 Chassis Grounds

Ground the power supply common lead to the chassis only at the backplane or system board. When grounding something to a chassis, case, or safety ground lead, always use a star lock washer or a lug with a lock washer pattern. Assuming the plating of the washer and the chassis are compatible, this makes a dozen or so gas-tight contacts between the screw, chassis, and lug. This doesn't work as well with PC boards, because the board flows slightly at elevated temperatures (the glass transition temperature of G-10 board is about 125 °C), which loosens the connections. A combination of a star washer with a spring washer such as a Belleville may be best for this, if your boards will be used or stored at high temperatures.

16.6.3 Magnetic Shielding

We saw that copper shields become ineffective at low frequencies, so that we'd need several centimeters of copper to shield 60 Hz really effectively. Equation (16.10) for the skin depth contains a factor of μ in the denominator, so a magnetic conductor is a better shield than a nonmagnetic one; from the skin depth alone, a material such as super-permalloy, with μ up near 10^6 , reduces the skin depth by a factor of 10^3 for the same conductivity, and mu metal ($\mu \approx 4 \times 10^4$) can manage a factor of 200.

High- μ materials can shield magnetic fields just the way a Faraday cage shields electrostatic ones. If the ambient fields are above an oersted or two, your high- μ shield is liable to saturate and become transparent; a mu-metal enclosure in a high field environment will need a steel shield around it. This is particularly true when the shield is held together with fasteners; the field concentrates in the regions of good contact and causes early saturation. High- μ materials are vulnerable to shock and bending, which pin the magnetic domain boundaries; although mu-metal comes in sheets, as soon as you bend it into a box and spot-weld the edges, it isn't mu metal anymore and has to be carefully annealed once again.

All in all, low frequency magnetic shielding is a lot harder than just soldering together some chunks of FR-4 board. Most of the time, you're much better off moving the offending transformer or switcher away from the sensitive circuitry and reducing its strays, either by using a wall wart power supply or by going to toroidal power transformers, using twisted pair wiring (which cancels the magnetically induced voltage on every twist), breaking ground loops, and doing everything you can to reduce the area of the loops collecting the stray flux. Small loop area, differential amplifiers, and wide separation of source and receiver are your best defenses, with steel and mu metal as the strategic reserve.

16.6.4 PC Boards

Circuit boards have to be debuggable, so you have to be able to get at the wiring. In ground plane boards, running signals on the top and bottom facilitates ECs[†]—you just cut the trace with a scalpel, instead of having to drill through it. One thing to remember is that the placement of the ground plane matters. Sensitive circuits can easily be driven nuts by a single pad lying over a noisy supply plane, for instance (see Example 19.1),

[†]Engineering changes (*i.e.* hacks), usually involving adding *roach wires*.




Figure 16.10. Cross section of a $\frac{1}{16}$ inch (1.57 mm) four-layer board.

whereas ground planes are really quiet. Figure 16.10 shows a typical four-layer board; you can get almost 1 mm thickness if you put signal wiring on the top, nothing on layer 2, ground on layer 3, and noncritical traces on layer 4. PC board houses can make the layers asymmetric in thickness, which is worth considering in some cases. If you have more than one ground plane, stitch them together with lots of vias. (There is some difference of opinion about this; if your signals are all at high frequency, the capacitance between the planes may be enough that you don't need so many vias.)

There are a number of quick-turnaround PC board houses with decent Web interfaces and free design tools, so you can lay your board out, hit the button, and have boards in a few days for a very reasonable price.

16.6.5 Design for Test

Before your board goes to layout, use the knowledge you gained in debugging the prototype to add test points and places to break the circuit (e.g., jumpers on power supply leads). They should be pads with via holes where possible, so that you can solder a pin to grab with your oscilloscope probe. Vias will stand a lot more soldering and messing about than an ordinary pad, which matters for test points, and your scope probe will put a lot of torque on the pin. Each test point should have a grounded pad and via within half an inch or so, because that probe needs a good ground if you want it to tell you the truth.

16.6.6 Connectors and Switches

Connectors, cables, pots, and switches are the main sources of flakiness in electronic systems. Choose them carefully; they're expensive to buy but even more expensive when they fail in the field. Don't run signal lines through panel mounted switches; to avoid pickup, use diodes, relays, or analog switches to switch signals, and run only control currents through the switches. Since most switches are especially unreliable at very low currents, this helps circuit reliability too. As for a panel pot, at least one of its leads

will have a high impedance to ground, and that invites capacitive pickup. In both cases, control lines can be heavily bypassed, but signal lines can't.

Use rugged BNC and SMA connectors for panel mount applications, but stick to the cheaper and smaller SMB or SMC connectors inside the instrument.

16.6.7 Multicard Systems

Multiboard systems have the advantage of flexibility, but are at an increasingly severe cost disadvantage, and are nowadays less necessary since circuit density has increased, logic parts are now programmable, and most of the display and control tasks are done on a computer. A 50,000 gate FPGA chip costs the same as one 200-position back-plane connector. Consider using a single-board design, perhaps with a daughter card if necessary, or *population options*: parts of the circuit that can be omitted by simply not stuffing that part of the board. If you really need multiple cards, check out the Eurocard system. It's a good all-round way to assemble a system using multiple boards, because it combines widely sourced hardware with a wide choice of single contacts, high current points, and coax connectors. Card edge connectors are now nearly obsolete, except in computer plug-in cards such as PCI Express, because they lack this flexibility.

16.6.8 Computer Plug-in Cards

Be very careful if you're building hardware to plug into a PC. The external interfaces (serial and parallel ports, USB, Ethernet, GPIB, and so on) are pretty robust, but if you plug a wrongly wired board into a PCI slot, you can blow up your motherboard. Be very sure that your circuit is correctly wired, and that the bidirectional bus is working properly, before you plug it in, and don't test it on your development machine. If you can possibly avoid a proprietary plug-in card, avoid it.

16.7 GETTING READY

16.7.1 Buy a Stock of Parts

You should have a stock of at least 50 pieces each of 1% metal film resistors, in values spaced approximately 10% apart. Get the $\frac{1}{4}$ watt type with leads. The mil-spec goes only up to 301 k Ω , but most manufacturers make them up to 1 M Ω in the commercial grade, and you can get 10 M Ω in the half-watt size. Nowadays, your production boards will be surface mount, but leaded components are indispensable in prototyping. Make sure you get some values such that $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{4}$, $\frac{1}{5}$, $\frac{1}{9}$, and $\frac{1}{10}$ of them are in the set too, for example 10.0k, 4.99k, 3.32k, 2.49k, 2.00k, 1.11k, and 1.00k, because that helps a lot in making amplifiers of integer-valued gains. Murata and other manufacturers make very nice collections of surface mount resistors, capacitors, inductors, and trimmers (both pots and capacitors). For RF work, and for stuffing your first few engineering boards, you should have a set of these kits too. It isn't a huge investment, and it'll make you fast.

16.7.2 Get the Right Equipment

You need at least a good fast analog scope, access to a digital one, plus some good way of measuring noise. The absolute minimum acceptable bandwidth is three times your maximum operating frequency; since a scope is a time-domain instrument, you need to

be able to see the shape of your waveform, and that means reproducing the fundamental and enough of its harmonics. The first, second, and third harmonic will allow you to see asymmetry and peaking or flattening of the tops of the waveform. You need at least the first 10 harmonics to be really sure you know what your signal looks like, for example, a 1 GHz vertical bandwidth for a 100 MHz amplifier.

The best way to measure noise is with a really well-calibrated spectrum analyzer[†] and low noise input amp, but setups like that don't grow on trees; for work up to 10 MHz, a filter whose noise bandwidth is accurately known[‡] plus a HP 3400A RMS voltmeter is a good substitute. (You can also use peak-reading meters such as the HP 400EL, but remember to add 1 dB to the RMS reading to get the noise power accurately—see Section 13.6.5.) For high frequency work, you also need access to a spectrum analyzer capable of detecting spurious oscillations up to twice the maximum f_T of your fastest component. Everything in the world resonates between 1 and 5 GHz, so if your circuit might have gain that far out, look for oscillations there, and look *hard*.

16.7.3 Soldering

Most soldering problems are caused by overcooking or undercooking. Surprisingly, both are caused by an iron that's too cold or too dirty. It may be a bit counterintuitive that a hot iron keeps components from overheating, but consider soldering with an iron at $T_{MP} + 1^\circ$: the parts would stay very hot for a long time before the solder melted. The amount of heat dumped into the package through the lead goes as the time integral of the temperature gradient; somewhat hotter irons used for a much shorter time win.

Accordingly, get a decent 60 W, temperature-controlled iron. The right temperature is about 750 °F (400 °C). Much cooler and you cook parts; much hotter and you can't keep the tip clean. Use a fairly fat conical tip (about 2 mm diameter, e.g., Weller PTA-7) for general purposes, and a skinnier one (1 mm, Weller PTK-7) for SMT prototyping. Stick with the conical ones and avoid the ones that neck down to a narrow cylinder; they look convenient but actually cook parts because they can't heat the leads fast enough. For RF boards, a large (100–150 W) soldering iron is very helpful for attaching the shields to the ground plane. You'll need a soldering station with a wet sponge and an iron holder. Get a good quality solder sucker (e.g., a Soldapullt). (Don't worry about the conductive plastic kind, they're not nearly as good as the plain Teflon tips.) Learn to use it by touch and with your less-skillful hand (cock it on the bench or on your knee). Get some solder wick for surface mount parts, and if you're going to be doing a lot of work with SMT ICs, get a SMT desoldering station, which will help keep the pads from being torn or cooked off. For those with bigger budgets, metals are very nice.

16.8 PROTOTYPING

16.8.1 Dead Bug Method

Building prototypes of high performance circuits needs a low capacitance, low inductance, flexible prototyping system: the *dead bug method*. You fit a piece of copperclad board into the lid of a suitable die-cast aluminum box, as shown in Figure 16.11, and use

[†]Emphasis on *well calibrated*. Be suspicious.

[‡]For example, from swept measurements or a first-principles calibration, as in Section 2.5.4




Figure 16.11. Dead bug prototype of a dual channel, linear phase bandpass filter and detector for modulation-generated carrier interferometry. A surface mount prototype appears in Figure 16.12. (Photo by Frank Giordano; a full-resolution color version is available at <http://electrooptical.net/www/beos2e/photos.html>.)

cyanoacrylate glue to attach ICs and transistors to the ground plane. These work like terminal strips, so you can hang your passive components right on the pins.

It's often ugly, but it's fast, it's easily modified, and it works up to 500 MHz or higher if you're careful.

Die-cast boxes are easy to drill and make good shields when you finish the prototype and attach the bottom. BNC bulkhead mount connectors hold the board into the box, and with the aluminum, they provide a very low impedance ground to minimize ground loop voltage drops (see Section 16.5.2).

Some people put the ICs upside down, some bend the leads and mount them right side up, depending on whether they have more trouble with breaking leads or forgetting which IC is which. If you go upside down, mark the package on the notched end, so you know which is pin 1. Try to keep all the chips going the same way, so you don't get confused.

Dead bug prototyping is easiest by far with DIP packages, and you should use those when possible. Unfortunately, they're starting to disappear, so we have to get good at prototyping with SMTs. SMT transistors are easy enough to get wires on, but their leads are so delicate that you have to be careful. As for SOICs, you can do the 0.050 inch ones by staggering the leads like saw teeth, but the 0.025 inch ones you can just forget. Buy or make an adapter to spread the leads to some useful spacing.[†]

It is possible to use SMT chip resistors and capacitors in dead bug prototypes, but it's fiddly, so do it only where really necessary; leaded components are so much easier to use that you'll be a factor of 10 faster with them.[‡]

16.8.2 Laying Out the Prototype

A common error is to run out of room. If this happens at one edge of the board, it isn't too serious; bolt another piece of board on that side, run some copper tape all along the seam, and solder it along both sides so that the ground plane is continuous.

[†]One good source of adapters is the Solder-Mount System, from Wainwright, Andechs, Germany.

[‡]There's not enough space here to go into all the details of dead bug prototyping high speed circuits. For more, have a look at the classic (and very large) Linear Technology Application Note AN47 by Jim Williams.

If it happens in the middle, where you perhaps were a bit optimistic about how much gain you could get out of one stage and now need two, it's much more awkward. So the best way to proceed is to build each functional block (an IF strip, say) on a small piece of board, and stick them together on a much larger sheet with screws and copper tape. Alternatively, you can just leave extra space in places where you're not positive it will work perfectly. Building prototypes that can be modified easily is vital—all prototypes eventually get to the point where they're too flaky to use, but good ones don't get there until you're finished with them.

If you're building something big that needs to be one piece (e.g., an RF back end), start with a 6×10 inch piece of copperclad board for a ground plane. You can attach smaller chunks of board to it as your prototype grows. Don't make the mistake of grounding the wrong side of the board, or you'll make very leaky shields.

Prototype your signal flow starting at one end and going to the other in a line. This keeps the input and output from coupling together and making the whole string oscillate. If your VHF circuit needs good isolation (≥ 60 dB) between any two places, keep them apart and use copperclad shield walls in between them. Use these shields to separate VHF or UHF gain stages into 20 dB chunks, and run the signal and power leads through small notches cut in the bottoms of the shields. Minimize loop sizes; a big loop will pick up RF on one side of the shield and reradiate it on the other side.

Think about ground currents; run especially noisy grounds right to the appropriate supply bypass cap if you can. If necessary, run two parallel wires a couple of diameters apart, which is nearly as good as foil because the **B** field between them tends to cancel. Remember that the inductance of a loop depends on its area; a big U made of wide foil is not necessarily lower in inductance than a shorter run of fine wire. Bridge any big loops.

16.8.3 Adding Components

The major skill in dead bug prototyping is making connections incrementally, in such a way that they don't come apart again when you remelt the solder to add the next part. You must avoid bending stresses in the component leads to prevent this.

To add a component (say, a resistor), gauge by eye how long the leads need to be. Cut and tin one lead. Using the other lead as a handle (bend it if that's more convenient), solder it to the desired node, keeping it as far as possible in its final position, to avoid bending stresses in the leads. Solder the other side. Now—and here's the vital step—*go back and reflow the joints, in the same order*. This releases the spring tension in the leads, and prevents their springing apart when you add another component. Now wiggle both ends, to make sure that you have everything properly connected, rather than two groups very near each other. Whatever you do, *don't* make the mistake of grouping all the leads to be joined, and then soldering them all at once. This will make a haywire prototype that will be hard to fix and hard to get working.

Dressing the leads short and close to the ground plane improves performance, reduces coupling, avoids shorts (since the component bodies wedge the leads apart), and generally makes it easy to build and modify. Most leaded components are amazingly strong—you can really crank on their leads if you need to, but really, really don't leave residual spring tension in them.

Wire the power supplies first, then the bypass caps on the ICs, then their supply connections. This puts the supply wiring on the bottom, which is good since it rarely needs changing, and gives the bypasses the shortest available path.

16.8.4 Hookup Wire

Most of the time, you'll use the clipped-off component leads littering the bench for hookup wire, and strips of copperclad board glued to the edges of the card as power supply bus bars. When you do need insulation, use #26 stranded, PVC jacketed hookup wire,[†] and use a consistent color code. Get the stuff with the really skinny jacket (OD is about 33.5 mils or 0.85 mm).

Strip only about one jacket diameter back from the end, and always twist the wire and tin it before bending the stripped part—otherwise strands will get loose. This is especially important if you like stripping the wire in several places to run daisy chain power connections. Always twist the wire at each stripped place and tin it.

Most people try using wire wrap wire for solder prototyping—usually only once. The trouble is that it's so fragile and the insulation melts too easily. The insulation is very useful by itself, though, for insulating things like thermocouples and thermistor leads.

16.8.5 Wire It Correctly and Check It

This is another wrinkle on the build-what-you-designed principle. Your gizmo has never worked yet, so there may be design errors in it, but wiring it wrongly will have you chasing your tail trying to figure out what's wrong. Orient the components so that you can read their markings, and always check the value or type number of each part as you install it; errors of factors of 10, or simple misfilings, are common, and lead to mystifying problems that take a long time to solve. SMT components are marked poorly or not at all, so pay special attention to getting them right, for example, by installing all the 100 k Ω resistors first, then all the 10 k Ω ones, and so on. Capacitors are usually unmarked.

Pay special attention to the wiring of power, ground, chip enables, and resets. On each chip, make sure you account for every last pin; if you wire only the ones you've thought about, there's sure to be a chip enable or latch input dangling somewhere. Use a yellow highlighter on a copy of the schematic to mark off the parts you've wired, so that if you're interrupted you'll know where to begin.

Building the whole prototype before turning it on is a mistake. You'll have to do the smoke test eventually,[‡] but test the circuit in small bits first: the voltage regulators, the first stage, the power output stage, and so on. Testing each one individually is very helpful in localizing faults—it's probably in the last thing you did. Keeping the subunits small limits the damage if something's been wired wrongly (and it probably has). As a corollary, always wire the power first, and test it thoroughly before connecting it up to the other bits.

16.8.6 Cobbling Copperclad Board

Ordinary G-10 or FR-4 copperclad board is a wonderful construction material, both for small structures and for boxes and shields. Use the $\frac{1}{16}$ inch, 2 ounce, two-sided kind. You cut it using either a big sheet metal shear if you have one, or by scoring and breaking it, and then cleaning it up with hand-held compound-action sheet metal shears. (Note: fiberglass board dulls shears quickly.)

[†]Stranded wire has a bad rep for prototyping, on account of shorts due to stray strands. It is much easier to work with, however, and that is more important.

[‡]That is, turning it on and seeing if smoke comes out.

Copperclad also makes good boxes and shields to separate gain stages and filter sections. Single isolated walls are good enough most of the time. Put sides on the sections when you get above 100 MHz, and bad cases such as power amplifiers and high- Q resonators will probably need a lid too.

Cut the board to size, tack-solder the sides in place, then use a high wattage (at least 60 W, preferably 100 W) iron with a broad screwdriver tip to run a continuous solder bead along all the seams. Remember that the inside and outside foils are insulated from each other, so make sure that at least one of them is continuous all the way around.

Sometimes people solder the box together using the inside foil, then attach it to the ground plane by the outside foil, and wonder why the shield is ineffective. Ideally, you should solder both, because it makes the whole thing stronger. Use copper tape, sold for PC board repair and for leaded glass hobbyists, to bridge the outside edges, and solder it down both sides (the adhesive is insulating).

If you do need a lid, you can make the hinge from copper tape, soldered all along both edges, and just tack-solder the lid here and there to the other sides of the box when you flip it down. Solder copper tape along all the edges of the lid to short the top and bottom foils together firmly.

16.8.7 Perforated Board

You can use plain 0.100 inch grid perf board (without pads) for small prototypes using DIPs and discretes, but make sure you use copper tape for power and ground; the ICs need some sort of solid mechanical connection in order for the prototype to be solid and reliable. There also exists copper-plated perf board, where you cut out pads yourself with a pad cutting tool, which is a bit like a miniature hole saw with a guide pin in the middle. The problem is the shreds of metal they leave behind, and the great big holes in the ground plane that result from cutting pads on adjacent holes (e.g., for an IC). If you put the ground plane on the component side, you can touch each hole with a drill to avoid shorts, which avoids pad cutters but isn't as mechanically robust.

16.8.8 Perf Board with Pads

Vero makes ground plane boards with pads and supply busses, which are good for more complicated digital and mixed-signal stuff where you want to use sockets. The ground plane is relieved around each hole, and solder mask is applied, which avoid shorts but make the ground plane more of a mesh (which works fine). When using this kind of board, make sure that you run the wires (longer than 0.3 inch) on top of the board, away from the solder pads. That way you can avoid solder bridges by running a small scraper (the edge of a thin metal ruler or the back of a broken X-Acto blade works well) between the pads before applying power. This is a sensitive tactile test for solder bridges if you pay attention. Using something sharper runs the risk of slicing off copper shavings and causing shorts instead of curing them. Make sure the ground plane on the top is stitched to the ground traces on the bottom at least every 2 cm, and more often if you have a high current, high frequency (or sharp edged) waveform coming from some chip.

The cheaper breadboards with the interdigitated power and ground areas, intended for digital ICs, are fine for slowish logic if you use enough bypass capacitors, but are not good enough for mixed-signal or RF work.

16.8.9 White Solderless Breadboards

These superficially attractive chunks of white nylon allow you to shove component leads right into an array of socket pins, which are interconnected in rows. A simple circuit can be wired up without using any solder at all. That's the good news. The bad news is leakage due to the low grade hygroscopic plastic, flakiness due to poor contacts and metal debris scraped off the leads, high stray capacitance and inductance, and poor bypassing due to ground inductance. Nobody can do good work on one of these. If you're a hopeless case, at least resist using white protoboards above 50 kHz, 100 mA, 50 V, or when a randomly sprinkled 0.1 Ω or 100 nH in series or 10 pF and 10 M Ω in shunt will screw you up.

16.8.10 Prototype Printed Circuit Boards

Your circuit will almost certainly be built on a PC board eventually, so why not do one for the prototype? You can if you like—there are quick-turn PCB houses such as PCB Express that will turn around small boards in a few days inexpensively. The bad news is that you then have to use CAD for your initial design, which will probably slow you down. The one exception to this rule is microwave circuits or very high impedance RF circuits, where to a large degree the strays are the circuit, and you expect to have a couple of board iterations anyway. Such boards are usually small—if they aren't, you've put too much on them, so consider putting the special stuff on its own card.

16.8.11 Blowing Up Prototypes

Prototypes and development hardware get a lot of abuse, and you don't want to spend all your time fixing things that you broke; thus you should spend a little on bulletproofing. For example, supply leads are always falling off, which can result in reverse-biasing some of your components, and—confidentially, now—sometimes we all put the banana plugs in the wrong binding posts on the power supply, or drop a screw in the wrong place. Protect yourself against power problems by using big beefy Schottky diodes in inverse-parallel with your supplies, adjusting the voltage and current limits on the supply so they will contain the damage, and turning power off at the supply each time you pick up a soldering iron.

16.9 SURFACE MOUNT PROTOTYPES

You can do dead bug prototypes with SMT devices, but it requires a delicate touch; the best bet is to use a mixture of leaded components and surface mount, which isn't as hard. The only thing to worry about with SMTs and dead bug is that the SMTs won't take much torque, so you can't just solder one end of an RN55C resistor to an SOT-23 transistor and then twist it to where you want it; you have to take the strain with long nose pliers instead, or the transistor lead will break off. Chip resistors and capacitors are worse still. The 1206 size is best, with the 0805 size[†] being barely tolerable, and anything smaller, you can forget.

[†]Short for footprints of 0.120 in. \times 0.060 in. (3.05 \times 1.52 mm) or 0.080 in. \times 0.050 in. (2.03 \times 1.27 mm).

16.9.1 Stuffing Surface Mount PC Boards

Stuffing surface mount boards by hand is a fiddly job, because of the small size of the components, their close spacings (inviting solder bridge shorts), and especially their unforgiving character. It takes about 5 or 10 times longer than when you're using leaded parts, on account of their small size and their lack of leads to allow sequential assembly and to use for hookup wire. You have to get rid of solder bridges with solder wick, because the usual trick of touching it with a clean iron usually doesn't work. You can't usually get a solder sucker in there, either. Solder wick is very hard to use in mid-air, because you have to grind it in with the soldering iron to get it to work. It is very important to rest your hand on a stable surface while working on SMT boards, so get a good selection of high quality, sharp-pointed, curved-jaw tweezers, which make that a lot easier. You have to get every lead in good contact with its corresponding pad, and well aligned, before applying any solder. You can use small amounts of glue on IC packages, but that makes them hard to remove, which is bad when you're debugging. A better way is to use a rare-earth magnet under the board plus a nut on top of the IC to keep it flat, so that all the leads match up and stay in contact with the board. You can hold capacitors and resistors with tweezers.

Wear a 3× eye loupe or some drugstore reading glasses,[†] and angle the soldering iron so that the handle tips toward the IC package, which helps the solder jump across the gap. Make sure you lean your elbow on the table, and the heel of your hand on the vice—it's steadier, you can keep it up longer, and by making you wrap your arm around, it keeps the soldering iron away from your cheek.

Alternatively, but more slowly, you can do it under a low power stereomicroscope (see Section 9.3.8). Using a single loupe lets you switch magnification by switching eyes, which saves time but will give you a headache if you do it all day—don't stuff a whole big board at once this way. Use solder paste and an SMT oven instead.

A chair of easily adjustable height is very important for this, because otherwise SMT prototyping is an ergonomic nightmare—if you use one of those tall lab stools, run and don't walk to get a pneumatically adjustable chair. You'll get a kink in your neck if the microscope isn't set up well, and severe eye strain if your magnifiers are of poor quality. Go with the good stuff—with an aching head and a sore neck, you're not at peak debugging efficiency. The author has a 1960s-vintage Zeiss surgical stereomicroscope, which is perfect.

16.9.2 Debugging SMT Boards

Debugging surface mount boards is a true pain. Make sure the signal leads are on the outside! One of the main troubles is that it is so hard to stuff an SMT board incrementally (see above). Another way, if you have an SMT oven for production, is to use solder paste and the oven to stuff a couple of boards completely, and hack them up. They don't survive very much hacking, so you'll need two or three just for test purposes.

Once you've soldered wires to SMT pins on the board a few times, the connections between the pins and the board become unreliable, and flakiness sets in. Put in a few test pads, ideally with largish via holes, to allow attaching wires without damaging the board. Vias are pretty robust even on SMT boards, but plain pads will get cooked off

[†]Drugstore glasses are made with an interpupillary distance of 66mm. If that's right for you, you're in good shape, but otherwise it will lead to severe eye strain. You can mail-order custom glasses for \$10 or thereabouts.

after only a couple of solder/desolder cycles. Pads with vias also allow you to bodge in some leaded components dead bug style while debugging, or to tack on a 450 ohm resistor in series with an SMB jack to connect to a 50 Ω scope input—the best kind of $\times 10$ probe for RF signals.

It is worth using a few 0 ohm resistors (SMT ones are available) that can be removed to isolate sections of the circuit for debugging purposes; in low power circuits, we often use 10 Ω resistors in series with the supply lead to each section for decoupling, which is a natural place to break the circuit. Put vias on those pads, too.

16.9.3 Probe Stations

One of the main problems with SMTs is that the pins are either inaccessible or too wimpy to clip a probe to. The solution is a poor man's probe station: a corner of an optical table with a few magnetic probe holders, and the board held firmly in a vice. That way you can make firm connections to the test point without needing to grab onto it (and risk tearing it off or shorting it out).

16.9.4 Hacking SMTs

Sometimes you won't have the right physical size resistor or capacitor—you'll need an 0805 but only have the larger 1206 ones. One good way to handle this (and the problem of needing to change values generally) is to mount your chip components *edgewise*—so that they stand up like miniature billboards. That way you can get your iron and tweezers in there much more easily, and can tilt the component to avoid shorts. Figure 16.12 shows an SMT board that was used for debugging. (It's the analog back end of a 12 Mpixels/s, hand-held, micropower, interferometric 3D surface scanner.)

16.9.5 Board Leakage

High impedance circuits such as pyroelectric detector front ends live or die by the quality of their insulation. Novolac epoxy, as used in IC packages, is a very good insulator, Teflon is better, and air is better still. Clean glass-epoxy circuit cards are pretty good, fingerprints are bad, solder flux is lousy, and damp solder flux is atrocious. A board that hasn't been cleaned is a failure waiting for the next rainy day. Cleaning surface mount




Figure 16.12. Hacked-up SMT board: this is a follow-on to the dead bug prototype of Figure 16.11, with only one channel plus some additional function. (Photo by Frank Giordano; higher resolution version at <http://electrooptical.net/www/beos2e/photos.html>.)

boards is especially difficult, because the cleaning solvent is held under the components by capillary action. The best method for getting rid of flux seems to be an ethanol or trichloroethylene ultrasonic rinse followed by a trip through an ordinary domestic dishwasher with ordinary dishwasher detergent. A good conformal coating applied in a low humidity environment will prevent condensation from causing problems. (The equilibrium thickness of the adsorbed water layer on a clean surface reaches one monolayer at about 40% relative humidity, and increases rapidly, to infinity at 100%—the board will leak even when it isn't dripping wet.) Guard rings around sensitive nodes are a good idea, but not a cure-all; if the leakage resistance gets low enough, you may run into funny noise peaks and instability caused by leakage from the guard ring becoming an important part of the feedback network.

At high temperatures, glass–epoxy board starts to leak through the bulk, so that wider guard rings (a few times the board thickness) and guarding on internal layers may become necessary.

16.10 PROTOTYPING FILTERS

Complicated filters are not that easy to prototype, even if you have network analyzers and such. Finding components is no problem; variable capacitors and inductors are easy to get. For small inductors (5–500 nH) you can easily wind your own, using (14.4), and squash or stretch them to adjust their values (that's why the coils in Figure 16.11 are irregular). The problem is tuning, that is, how to go about adjusting them to give the desired response. Filters with only a couple of lowpass sections aren't hard to tune, but complicated bandpass filters are more difficult; your starting point has to be close to the desired filter.

The best way to ensure this is to make sure the component values are very close to correct before you start, and that means using an accurate standard inductor or capacitor to resonate each one somewhere near the operating frequency, and then calculating the value. *LCR* meters aren't a good solution at high frequencies.

Bandpass networks are much more sensitive to the resonant frequencies of their individual sections than to the exact component values, so for phase-insensitive applications you may be able to get away with just measuring the resonant frequency of the combination (e.g., with a grid dip meter). See Section 16.11 for a more organized way to tune high-*Q* filters.

16.10.1 Standard Capacitors

Highly accurate picofarad capacitors are scarce, and so are accurate inductors of any sort. So how do we do this? A good way at HF and VHF (10–300 MHz) is to use shorted transmission line sections. You can measure the propagation velocity of your transmission line very accurately by making a shorted stub of a known length, and then finding where it is exactly a half-wave long (i.e., where it looks like a short again). Use the setup of Figure 16.13, and adjust the frequency until the signal goes away (verify that it's $\lambda/2$ you've found and not λ). If you want to use a fixed generator, you can adjust the stub's length using a thumbtack as a movable short, then measure the length of the cable. Because you don't know the exact impedance of the generator, scope, and tee connector (which will change the effective length of the stub), find the first *two* nulls, which are exactly $\lambda/2$ apart, and measure the distance between them.




Figure 16.13. Calibrating test capacitors: generator, oscilloscope, open-circuited cable, and thumbtack.

Once you know the propagation speed, measure Z_0 with a long length of cable; sweep the frequency while looking at the generator end with a high impedance probe (i.e., a 2 k Ω chip resistor and the 50 Ω input setting of the scope, *not* a $\times 10$ probe) and adjust the resistance at the far end until the ripples in the envelope go away completely. Use a very small pot, or (ideally) solder a 56 Ω chip resistor right across the coax at the end (no pigtail) and then scribble on the resistor body with a soft pencil to adjust its resistance to Z_0 (most 50 Ω coax is really 51 or 53 Ω , so a 49.9 Ω 1% resistor won't do). Alternatively, use an emery board to adjust the value of a 49.9 Ω metal film resistor. Either way, measure the resistance with an ohmmeter, and you know Z_0 . Note: Don't skip this step, because there's a factor of Z_0 in all the formulas for backing out LC values from coax reflections.

16.10.2 Calibrating Inductors and Capacitors: A Hack

A quick and quite accurate calibration method uses the setup of Figure 16.13 minus the thumbtack. Cut a coax patch cord to a length that has its first open-circuit resonance ($\lambda/4$ electrical) in the frequency band of interest, and just short the inductor or capacitor across the open end (no pigtails—stick it right on the featureless chopped-off end). Then if the first OC resonance is at f_0 , and the first resonance with the shunt element is f_1 , the inductance is

$$L = -\frac{Z_0}{2\pi f_1} \tan\left(\frac{\pi f_1}{2f_0}\right). \quad (16.11)$$

Similarly for a capacitor,

$$C = \frac{1}{2\pi f_1 Z_0} \cot\left(\frac{\pi f_1}{2f_0}\right). \quad (16.12)$$

This is OK as long as (1) the frequency shift is much larger than f_0/Q , (2) you get the right null, and (3) you really know Z_0 accurately (see above).

We can use this to get a set of beautifully calibrated 1206-series COG chip capacitors epoxied to popsicle sticks, that we can use for a long time to come for measuring inductors and so on. We can do the same with the inductors, but the caps are more

temperature stable and less susceptible to outside influences, for example, proximity to nearby magnets or conductors, which can change their inductance.

Aside: Variable Standards. It is sometimes possible to make *a priori* standards that don't need extra calibration. For example, consider using a trombone line as a low frequency capacitor. Provided that the shaft is long enough for end effects to have died away, adding a length Δl doesn't change the end contributions; it just lengthens the radial-field region in between. Thus the added capacitance $\Delta C/\Delta l$ is that of a section of an infinitely long coaxial line, which (in MKS units) is

$$\frac{\Delta C}{\Delta l} = \frac{2\pi\epsilon_0\epsilon_r}{\ln(b/a)}. \quad (16.13)$$

Finding the offset C_0 is easily done by putting an inductor in parallel, measuring the resonant frequency at various Δl values, and fitting the curve to extract L and C_0 . Similar approaches include the cutback method for waveguide attenuation measurements, the crystal ring-down calibrator of Section 15.7.2, and the turbidity meter of Example 13.11. It's worth looking for this sort of idea anywhere your calibration requirements start looking awkward.

16.10.3 Filter Layout

There are two other difficulties with filter prototyping: the inductance of the leads and stray coupling between sections. Stray coupling is mostly a problem when you need really good rejection in the stopband. Orienting the coils at 90° to one another helps, and in bad cases, copper shields between sections make a big difference. Make sure you don't screw up your inductors by mounting them right next to a metal surface nearly perpendicular to the coil's axes: that looks like a shorted turn, which reduces the inductance and Q by as much as 25%.

16.10.4 Watch for Inductive Coupling

The main drawback of large air-core coils is that they tend to couple to other coils and even to component leads. If you have a mental picture of what the field lines of a solenoid look like, you can orient the coils to reduce the coupling fairly easily: tip them so that the *field lines* due to the two coils are roughly orthogonal inside the coils. The commonly heard advice to put the axes at right angles is wrong unless they intersect at the center of one of the coils. This approach is good for about a 20 dB reduction in coupling. If that isn't enough, use shielded coils or shield the air-core coils from each other with copperclad board.

16.11 TUNING, OR, YOU CAN'T OPTIMIZE WHAT YOU CAN'T SEE

Tuning a network is a knack, like aligning optical systems on the bench. You get good at it after a while, but it's slow starting out. The best way to learn is to sit at the elbow of an expert—as a 21-year-old junior designer, the author was taught it by a crusty old technician of about 26. You spend some time getting a good starting point, then go round the circuit, adjusting one thing after another until it looks the way the simulation says it

ought to. Like any iterative optimization process, the steps start out big and get smaller as you get closer. If there are only a couple of components (e.g., an L -network matching section between two RF amplifiers), tuning is easy. Much the same is true if the sections don't interact much, for example, a multistage tuned IF strip or an active filter using several op amps in cascade—it can be broken down into several easy problems.

On the other hand, tuning multiple-section LC filters with carefully optimized responses can be tough if you're new to it, because all the controls interact. What's more, the reason we've spent so much design effort on that filter is usually that its characteristics govern the performance of the whole instrument. A typical example is the last IF filter in an RF signal processor; its bandwidth sets the SNR, and the sidelobes of its impulse response determine how fast we can take data. That one just *has* to be right, and we often can't buy it off the shelf. (Custom filters can be had in a week for a hundred dollars or so from TTE.)

In numerical optimizers, each trial step can adjust all the L and C parameters, and every step produces a numerical value of the penalty function that can be compared accurately with all previous values. Real-world tuning isn't like that—we have to adjust the components individually, and we're susceptible to noisy readouts and our own mental inconsistency. Thus if we find ourselves down some long narrow valley in parameter space, fairly far from the optimum, we might never tune our way out of it. It is critically important to have stable, detailed, real-time indicators of network performance, and to start out near the optimum in the first place. After that, a few pointers will get you well on your way.

1. *Tune passive networks for return loss.* Figure 16.14 shows the right way to tune filters when we care only about amplitude response (e.g., wideband RF stage filters for image rejection) or when the desired signal looks like noise and not pulses.

You tune for return loss because the peaks of the transmission response correspond to $\Gamma = 0$; you're looking for easily spotted nulls rather than the very subtle peaks of the insertion loss. It really isn't possible to tune a high performance filter looking only at the insertion loss, since your cables, measurement setup, and generator variations with tuning will disguise small insertion loss changes. Do check the insertion loss when you're done, just to verify that you're not matching beautifully into the copper losses in your coil, and that the skirt selectivity is what you expect—the skirts are as invisible in reflection as the peaks are in transmission, and for the same reason.

If you have a fair amount of microwave work to do, it's really worth getting familiar with a magic labor-saving device: the Smith chart. It's beyond our scope here, but you can find it in microwave books and online tutorials: trust me, it's worth learning—it can make you 100 times faster at microwave design.

2. *Use Dishal's method for high- Q filters.* Sharp filters are made of high- Q sections (reactances large compared to Z_L), so when they're detuned, series sections look like open circuits and parallel sections look like shorts. This allows us to use a shortcut called *Dishal's method*.[†] Starting with the input section, use the return loss test setup to tune its resonance to $1/(2\pi\sqrt{L_1C_1})$. Go down the ladder and tune each section to its calculated resonance in turn, and when you're done, your filter is pretty nearly perfectly tuned. (For a lower- Q filter, you can do the opening and shorting with a soldering iron.) The reason

[†]M. Dishal, Alignment and adjustment of synchronously tuned, multiple-resonant-circuit filters. *Electrical Commun.*, pp. 154–164 (1952)

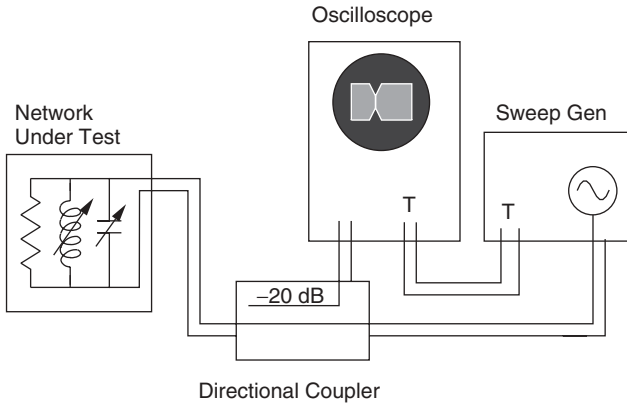


Figure 16.14. Filter-tuning setup with directional coupler.

this works is that (in the high- Q limit) series resonances don't move around when you put other reactances in parallel with them, and parallel resonances don't move when you put reactances in series with them. For low- Q filters this isn't terribly accurate, but it's usually close enough to get all the bumps showing (see item 4).

3. *Tune pulse filters for impulse response.* In pulse filters, the things that matter are the impulse response and the noise bandwidth, but what are usually specified are the amplitude response and phase nonlinearity limits, which is a bit indirect.

Whichever you choose, you have to have both things on the screen simultaneously if you want to optimize both. A vector network analyzer will do this, and is a must in a production environment, but there are other ways for the development lab. The easiest is to do a preliminary tune using the return loss method, and then use the setup of Figure 16.15 for the final tweak. Make sure the pulses are clean, and are at most 10% of the width of the desired impulse response (or 10% of one cycle of f_0 for BP filters), and that the duty cycle is 1% or so; since the pulse has a flat spectrum out well beyond the

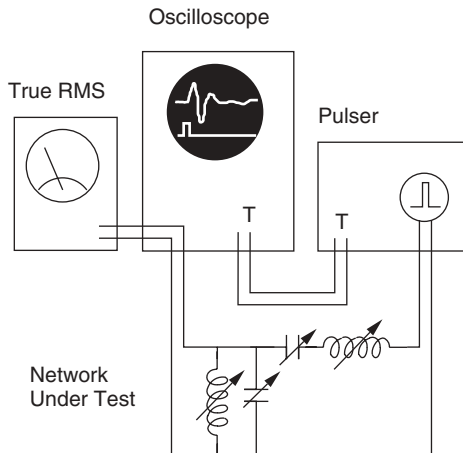


Figure 16.15. Setup for tuning pulse filters.

filter cutoff, the total AC power output represents the noise bandwidth, and the impulse response comes out directly. You'll need a true RMS power meter with really good crest factor performance, of course; a thermal detector is one good way to do it, or a calibrated square-law diode detector, or even a sufficiently fast digital scope plus some software. In any event, make sure you have a real-time, analog indication of the total RF power. That means a scope trace or a good quality electromechanical analog meter—the little bar graphs on DVMs are beneath contempt.

4. *Get the right number of bumps first, then make them look nice.* This advice applies even for filters that aren't supposed to have bumps (e.g., Butterworths); a (mildly mistuned) N -pole lowpass or highpass ladder filter should have N bumps in its transfer function, counting both positive and negative going ones, and bandpass ones should have $2N$. That way you're sure to have all the resonances in the right neighborhood.

5. *Get the center frequency and bandwidth right early.* In filter design, for a given number of sections, you have to trade off performance pretty rapidly to gain bandwidth. In tuning, it is tempting to get the filter looking nice at a slightly too-narrow bandwidth, and then expect to tweak it to get the bandwidth right. Don't fall for it—there are so many local minima when you do this that you might as well start over. The desired filter design may not even be nearby when you've done it this way. Get the right number of bumps, spread them over the right bandwidth, then optimize.

The reason for this is interesting: when we tune a capacitor or inductor, we're changing not only the resonance of the section, but its impedance as well; in moving the section resonances slowly around (as $(LC)^{-1/2}$), we change the coupling coefficient between sections rapidly (as C/L) in the process. If we had an adjustment method that changed f_0 and k separately,[†] this advice would not apply. The key is to make the adjustment ranges no wider than necessary to accommodate component tolerances; a circuit with tunable coils and trimmer capacitors as the only reactive elements is a bad circuit.

6. *Tune front ends for minimum noise figure.* RF front ends may exhibit minimum noise figure somewhere other than at best matching. Look at the SNR of a small modulation on the carrier to get the best tuning point.

7. *Don't be fooled by running out of range.* The capacitance of a trimmer has a value that depends on the area of overlap of the stator and rotor plates. Minimum and maximum capacitance are half a turn apart, and you can go either direction to get there. A network tuned with variable capacitors should thus exhibit *two* optima per turn, if the real optimum is inside its tuning range. If there's only one optimum per revolution, all that's happened is that you've run out of tuning range. This is a frequent problem.

Tunable inductors are a little less common, but will also exhibit multiple optima in general; the magnetic effects of the core are greatest when it is centered in the winding, so positions to one side are equivalent to those on the other. This is not true if there is coupling to other circuitry involved, however, since the two positions will exhibit different fringing fields and so different coupling. In any case, it's pretty easy to tell if you've run out of range on a coil—the slug falls out. If you have reason to know that the true optimum is somewhere in the adjustment range, don't worry about it; just keep tuning the other elements until the railed one comes back in range. Otherwise, try sticking in a fixed element to move the center of the range.

[†]We're talking about k the coefficient of coupling here, not k the relative permittivity, k the propagation constant, k the imaginary part of the complex refractive index, or k for Boltzmann's constant. Isn't interdisciplinary work fun?

Example 16.5: Tuning by Finite Differences. It is usually best to tune complicated networks by finite differences, especially when the readout is noisy. The author and his colleagues were once faced with a motorized π -network match box (one motorized tapped inductor and two motorized vacuum variable capacitors), coupling a 5 kW RF generator into a plasma etch chamber. Solid state RF power amps are much more vulnerable to mismatch than tubes; the big standing waves in V and I will blow up the rather delicate output devices. This one automatically reduced its gain to keep the reflected power below 100 W or so, producing nasty discontinuities in the reflected power with tuning. What's more, the RF impedance of a plasma chamber is wildly nonlinear, so we couldn't just turn the power down, and the onset of power limiting detuned the network severely, so that it had to be coaxed out of power limiting. Even worse, the forward and reverse power meters were digital, and the plasma was unstable, so that the digits kept bobbling up and down even at a fixed setting.

After a change in the cathode design, the network wouldn't tune at all; capacitor C_{fast} was very tweaky, and C_{slow} did almost nothing. Every adjustment of C_{fast} would send the amplifier into power limiting, so the usual iterative method of slowly trading one control off against the other while moving the system toward match failed—we were in one of those long narrow valleys, between cliffs, with lots of loose rocks. The solution was to tune by finite difference: tune the network for best VSWR with C_{fast} at one setting of C_{slow} , then move C_{slow} a fairly long way and tune it again using C_{fast} . If the minimum VSWR went up, that was the wrong way to move C_{slow} . By taking big enough leaps that the VSWR change was unambiguous, we got enough directional information to tune the network correctly. The final fix was to put another capacitor in series with C_{fast} to reduce its tuning sensitivity and confine it to the right neighborhood, but we couldn't tell what value to use until we'd tuned the network.

Digital Postprocessing

“Son,” the old guy says, “no matter how far you travel, or how smart you get, always remember this: Some day, somewhere,” he says, “a guy is going to come to you and show you a nice brand-new deck of cards on which the seal is never broken, and this guy is going to offer to bet you that the jack of spades will jump out of this deck and squirt cider in your ear. But, son,” the old guy says, “do not bet him, for as sure as you do you are going to get an ear full of cider.”

—Damon Runyon, *The Idyll of Miss Sarah Brown*

17.1 INTRODUCTION

Everybody needs to know something about postprocessing data. There are two benefits of this: to be able to do it, of course, but as importantly, to gain good technical taste for what should be done in digital postprocessing and what must be done in the front end and the analog signal processing. Discerning that is a high and useful skill and should be cultivated.

In this chapter, we start with the raw numbers at the output of the digitizer and see how to produce real experimental results at the end. The focus is on strategies and algorithms, rather than on code. We’ll explicitly ignore image processing, partly because it is a field unto itself, but more particularly because it is mainly good for pretty pictures and heavy processing such as feature extraction and machine vision, which are well beyond the scope of this book.

The basic thesis here is that postprocessing is not a substitute for good data. The best postprocessing strategy is one that makes few assumptions about the data and does no violence to them either; by and large that limits us to linear operations. Nonlinear processing can make nice pictures but you must use it very sparingly if quantitative data matter to you—enough processing can make a given picture look like anything at all. You might as well just multiply by 0 and add what you think it should look like.

Be thoughtful too about what order to do nonlinear processing in; for example, if you are measuring the phase of an RF signal directly, with a tracking phase detector, and then postprocessing the phase numbers as your signal, you should probably unwrap the phase before filtering. On the other hand, if the measurement is in rectangular coordinates (e.g., with an I and Q detector; see Section 13.8.7), you might want to filter and then unwrap, especially if you have lots of additive noise to contend with. If your intuition runs into

a wall, do some numerical experiments on canned data—and use real data if it’s in any way possible.

Choice of domain matters a lot. Certain operations are very much easier in one domain than other ones: filtering is easy in the frequency domain; image compression and some matrix operations in the wavelet domain; echo cancellation in the cepstral domain. We’ll begin with the simplest sorts of postprocessing (e.g., gain and offset correction), then state the basic theorems and use them to develop digital filtering, quantization effects, and the recovery of signals from noise. The emphasis is on interrelating these topics and showing how the lore is grounded in the theory, but if you’re new to the topic you’ll need a couple of good books, for example, Oppenheim and Shafer for signal processing theory and *Numerical Recipes in C* for algorithms and coding advice (see the Appendix).

17.2 ELEMENTARY POSTPROCESSING

Even instruments that don’t need a lot of filtering, deconvolution, and so on usually need to massage the raw digitizer output into a properly calibrated measurement.

17.2.1 Gain and Offset

The most basic postprocessing operation is to figure out where 0 and full scale are in measurement terms, and to apply a numerical gain and offset to transform the raw data into measurement units. For example, if a 12 bit digitizer was used with the turbidity meter of Example 13.11, and calibration data indicate that 0 signal corresponds to 8.3 ± 0.15 ADU and that a turbid extinction of $0.1\%/cm$ ($\pm 1\%$ of value) gives a reading of 3037.5 ± 2.2 ADU, the turbid extinction is given by

$$E = 0.1\%/cm \times \frac{(R - 8.3)}{(3037.5 - 8.3)}. \quad (17.1)$$

The standard error is found by differentiating with respect to all uncertain quantities, multiplying each partial derivative by the RMS error in the quantity, and taking the RMS sum of the result, just as we did in Section 10.6.2. Since all subsequent stages will depend on this calibration, you need to be a bit paranoid about it; in the turbidity meter example, clean water and $0.1\%/cm$ standard turbid water are probably not enough as calibration points; it needs some intermediate values. This improves accuracy as well as allowing sanity checks—if you have several independent measurements, one improperly mixed calibration solution won’t cause you to ship a miscalibrated instrument. Keep a close watch on how you could be fooled: for instance, in Section 13.11.5 we talked about making sure that there’s enough of an offset that 0 signal is guaranteed to give an on-scale reading. If you can’t possibly do that, for example, if you’re using preexisting equipment, then take several data points, use curve fitting to get the offset, and watch its drift vigilantly.

17.2.2 Background Correction and Calibration

A strategy involving more than a single measurement location (e.g., spectroscopy or imaging) has a somewhat more complicated calibration requirement, as the system gain

and offset may change with scan position. Of course, in designing instruments we try to make any such variation as small and as stable as possible, with stability being much the more important of the two. The simplest case is where the gain is close enough to constant, but the offset varies, for example, a mechanically scanned system where the background light depends on scan position; all we have to do then is subtract the background signal.

A more general and more common case is when the gain changes with scan position as well, as in a tunable-laser spectrometer, where the laser power depends on the tuning and there are etalon fringes in the sample cell. If the gain and offset are stable with time, we can usually produce calibration tables containing gains and offsets for each scan point, as we did in Section 3.9.19. The usual way is to make several calibration runs (from a minimum of about four to many thousands, depending on how expensive the calibration runs are and how much real data you intend to take). Remember that any residual noise from the calibration contributes fixed-pattern noise that will not go away when you signal-average your calibrated measurement data, so make sure the calibration noise is small compared to the residual random noise after signal averaging.

The other key to good calibrations is to calibrate often enough that background shifts do not corrupt the measurement. These two requirements are of course in direct conflict and must be traded off depending on the situation. A bit of ingenuity will usually produce some heuristics and sanity checks that you can use to determine when to calibrate and whether the current calibration is good. An instrument might be programmed to recalibrate itself any time the temperature changes by more than 0.87°C from last time, the system has recently been turned on and has now stabilized, 10 hours have elapsed, or a single-scan sanity check falls outside the 0.1% confidence limit.

As always in instrument design, make sure you calculate your signal-to-noise ratio carefully at each step; if you get stuck somewhere, that shows you where your unexamined assumptions are going to trip you up. Remember that we live and die by the SNR, so it is well worth your time to become clear, or to get some consulting help.

17.2.3 Frame Subtraction

There is a poor man's version of this careful correction procedure, one that may be unwise to even mention, as it sometimes constitutes a temptation to cutting corners unnecessarily. It is to make two runs where the background is the same but the signal is different, and simply subtract them. A legitimate example is a CCD with slow shuttering, for example, a frame transfer device, used in bright light. In that case, the bright background effectively gets shifted across the whole array during readout, making a big bright smear across the image. At the price of more or less serious dynamic range reduction, you can eliminate the smear caused by the slow shifting of the image into the storage array by taking two exposures, say, 2 s and 500 ms, and subtracting them. The smeared part of the image will be the same in both images (and so will cancel), but 60% of the unsmeared photons will survive. Providing that there's no bloom or nonlinearity to worry about, and that the image contrast is not too high,[†] summing several images taken in this way will recover the lost SNR.

[†]Why does the contrast matter?




Figure 17.1. Baseline restoration.

17.2.4 Baseline Restoration

When the background fluctuates more rapidly than a canned calibration can deal with, we have to try to use internal cues instead, leading to the idea of baseline restoration (Figure 17.1). The prototype is the DC restore function of a TV signal. The video signal's instantaneous voltage controls the local picture intensity (black is most positive). However, the video detector output is AC-coupled, which might seem to make this difficult. At certain times in each scan line, reference levels are transmitted, which are sampled and used as references to adjust the offset line by line, establishing a stable zero point. The same idea is useful in many situations where the background fluctuates widely, on a time scale that is long compared to the line rate, but too fast for a static calibration to help.

Baseline restoration must be applied carefully, or it will introduce artifacts by subtracting the wrong DC value; make sure the baseline sampling is done at a time when the signal is known to be 0, for example, when the light source is cut off, the shutter is closed, or the sample is not present, and—crucially—allow enough settling time on both ends when switching between signal and baseline acquisition. Remember that for baseline restoration to be applicable at all, the rapidly drifting part of the background has to be independent of scan position, which is not always the case.

Baseline restoration can also be applied point-by-point in sufficiently slow measurements, in which case it turns into a digital version of auto-zeroing (Section 15.10.3).

Example 17.1: Baseline Correction and Deconvolution for a Pyroelectric Imager. In Sections 3.11.16 and 18.7.2 we talk about pyroelectric sensor design; here we'll look




Figure 17.2. The transfer function of an unusual pyroelectric imager: the top curve shows the pixel charge response, which after sampling becomes the bottom curve. The two-stage inverse filter (dashed) is implemented as a nine-sample FIR filter for high frequency boost (solid black) followed by a running sum, and produces the well-behaved final response (solid grey). The ill-conditioning of the integration step is removed by adding a slow negative drift to the data and applying a positivity constraint: people are warmer than the floor.

at the DSP aspects. Pyroelectric detectors convert a temperature change into a charge, rather than a current as quantum detectors and bolometers do. That means that pyros are differentiating; the current is proportional to dT_{sensor}/dt , and when the charge is sampled periodically for measurement, a constant temperature corresponds to zero current. While this effect is fast, what we actually want to measure is scene temperature, not sensor temperature; due to the weak coupling and the thermal mass of the sensor, its response to scene temperature is a slow one-pole rolloff at $f_c = G/(2\pi M_{\text{th}})$, where G is total thermal conductance and M_{th} is the thermal mass of the detector element. Above f_c this rolloff cancels the differentiator slope, so that the response of a pyroelectric becomes flat for a bit, until thermal transient behavior becomes important.[†]

Thermal conduction acts exactly like a load resistor on a photodiode: it reduces the signal available and it adds thermal noise. Thus we would appear to win by insulating the sensor very well, so that its response is dominated by radiation, as Figure 17.2 shows. Unfortunately, this also slows the sensor down. What's more, the differentiating action means that forming a decent image is a bit of a puzzle; the inverse of a differentiator is an integrator, which applies infinite gain at zero frequency, where the sensor has an SNR of zero (not 0 dB, *zero*), which makes it hard to keep a stable background.

Each pixel has its own gain and offset, so we have to flat-field the sensors as in Section 3.9.19, but all pixels have the same thermal response. After applying the inverse filter, we still have to fix the DC problem, which we do by subtracting a slow linear ramp (equivalent to about -10 mK/s) from the data, and then forcing all negative ΔT values

[†]See Chapter 20, "Thermal Control," available at <http://electrooptical.net/www/beos2e/thermal2.pdf>.

to zero. This *positivity constraint* and ramp force all the quiescent pixels to stay near $\Delta T = 0$. Interestingly, even the kTC noise is reduced by this approach; the noise charge has nowhere to go except the pixel capacitance, so it gets recycled into the following sample, and the running sum tends to make it cancel out in the final data.

17.2.5 Two-Channel Correction

If the background fluctuates too fast even for baseline restoration, you'll have to move your measurement to a frequency where the background is quiet (Section 10.9), or use a second detector that sees the background and not the signal. This is done in Section 18.6.3 in analog, but you can do something similar (though far less effective) by using a second digitizer channel to measure the background. It is crucial that both detectors see the same background, which may require some rejiggering of the apparatus, for example, a noisy arc lamp might have to be put in an integrating sphere. Unless your background measurement is somehow quieter than your signal measurement, two-channel correction will cost you at least 3 dB in ultimate SNR.

17.2.6 Plane Subtraction and Drift

Sometimes the error in adjacent scan lines is highly correlated, for example, a scanning tunneling microscope (STM) picture where there is a secular thermal drift in the z actuator, which controls tip-to-sample separation. The feedback loop will keep the tip-sample separation constant, so the z drift will instead show up in the control voltage, which is the measurement data. If the raster is regular and unidirectional, with a line period T_{row} and a pixel period T_{col} , a constant drift rate v will produce a topographic image that differs from the true topography by a linear function of the row and column indices i and j ,

$$Z_{\text{meas}}(x, y) = Z_{\text{true}}(x, y) + ai + bj + c, \quad (17.2)$$

where $a = v \cdot T_{\text{row}}$ and $b = v \cdot T_{\text{col}}$. The error terms are an analytical description of a plane, and so correcting for them is known as *plane subtraction*. Plane subtraction is less vulnerable to artifacts than line-by-line background restoration, since only one rate and one offset need be estimated, which takes a great deal less calibration data. It also has only two adjustable parameters, which increases our confidence in the measurement since you can't sweep much dirt under a rug that small.

17.2.7 More Aggressive Drift Correction

If plane subtraction is good for controlling drift and is pretty safe, how about using a higher order model? Sometimes you can measure the drift independently, for example, when you purposely build in a signal-free calibration time in each scan line to facilitate baseline restoration, and then more complicated background subtractions may be OK; however, in cases where the fitted background is coming from the data itself, it's a mistake. There are just too many unexpected ways of fooling yourself and others. Whatever operation you use to get the coefficients for the fancy fit, it will soon disappear from your consciousness, and may be silently generating reasonable-looking wrong answers long afterwards. There's no escaping it: you'll have to get better data.

17.3 DEAD TIME CORRECTION

We saw in Sections 3.6.2 and 3.6.4 that to get stable and accurate photon counting measurements, you have to add a bit of *dead time* τ after each photocount, which makes the measurement nonlinear, because some photons will be absorbed during the dead time and so be missed. The simplest way of looking at this is to say that at a photon detection rate λ' , the proportion of dead time is $\tau\lambda'$, so the true photon arrival rate λ is

$$\lambda = \lambda' / (1 - \tau\lambda'). \quad (17.3)$$

This approach gives excellent results as long as the photon arrival is a Poisson process with a constant rate (λ is constant during time τ) and $\tau\lambda \lesssim 0.2$; otherwise more thought is required.

17.4 FOURIER DOMAIN TECHNIQUES

This section really isn't an introduction to digital signal processing; it's an attempt to connect the digital realm to what we already know about time, frequency, and SNR, emphasizing rules you can hold on to and some nasty potholes to steer clear of. Readers unfamiliar with the basic material should consult Bracewell for a firm and intuitive grounding.

17.4.1 Discrete Function Spaces

The data set from our digitizer consists of a finite number of samples d_n , taken at known discrete times t_n . It can be looked at as an ordinary column vector, of the sort we're familiar with from linear algebra. We know that such a vector can be written as a linear combination of basis vectors, and that we can use any basis set we like as long as the basis vectors are linearly independent. Transforming from one basis set to another requires a matrix multiplication, and the results can be seriously corrupted by roundoff error if either the new or the old basis set is nearly linearly dependent. The best case is when both bases are orthonormal; that is, the basis vectors within each set are mutually orthogonal and have length 1.

Our data can similarly be represented as weighted sums of other functions, which is the idea of a function space. Continuous functions of infinite length give rise to complicated function space problems, but we don't have to worry about that here.

The data values consist of a mixture of signal, noise, and spurious products due to nonlinearity and the digitizing process itself; it turns out to be possible to use the functional decomposition of our data vector to separate the signal from the crud pretty well in most cases, unless we've allowed the spurs to land on top of the data.

By far the most useful basis set is the complex exponentials, $\phi_n = \exp(i2\pi nft)$, which lead to modeling signals as trigonometric polynomials (weighted sums of a fundamental and its harmonics). They are not only linearly independent, but orthogonal, both in the continuous- and discrete-time cases (where they are also called geometric sequences). Trig polynomials are intuitive and mathematically convenient, and have a close relationship to concepts we're familiar with from continuous-time analysis, such as bandwidth and SNR. A slightly subtler way of saying this is that they are the eigenfunctions of

linear, time-invariant discrete-time networks, which are expressible as systems of linear difference equations with constant coefficients. This “geometric in, geometric out” behavior is exactly analogous to the “exponential in, exponential out” property of linear circuits. (Note that the sampling rate has to be kept constant for the network to be time invariant.) A very useful property is that sines and cosines are the correct orthogonal functions for equal weighting on both the finite continuous interval *and* on an equally spaced grid of points, in any number of dimensions, so we can move easily between integral transforms and finite Fourier series.

Sines and cosines don’t exhaust the class of complex exponentials, of course. Do keep in mind that there are lots of functions that are not particularly easily represented this way, especially those with asymptotes, discontinuities, or rapid changes of behavior, and that other representations are better for some uses.

Sampled data can be represented as vectors, and linear operations are matrix multiplications (although they’re not implemented that way). However, what we care about is faithful representation and reproduction of the continuous-time input waveform, and furthermore, a finite length of sampled waveform has a continuous Fourier transform, so that we’re not in a pure linear algebra situation.

17.4.2 Finite Length Data

A function $g(t)$ that is nonzero only on the domain $[t_0, t_0 + P]$ is said to have *compact support*. (We can choose $t_0 = 0$ without loss of generality.) It will have a transform that is a continuous function of frequency, so that an uncountably infinite number of Fourier components are needed to represent it. This is not an attractive property, so we naturally seek ways around it. The most common way is to make it periodic, so that it is representable by a Fourier series, which is only countably infinite, and thus ever so much more practical. A function with compact support can be made periodic in many ways. The easiest is to join many copies together, defining $\hat{g}(t) \equiv g(t \bmod P)$. In the special case of compact support, this is equivalent to

$$\hat{g}(t) \equiv \sum_{n=-\infty}^{\infty} g(t + nP) = g * \text{III} \left(\frac{t}{P} \right), \quad (17.4)$$

which is a fruitful observation, as we’ll see.

Such a periodic function has a Fourier series representation,

$$\hat{g}(t) = \sum_{n=-\infty}^{\infty} a_n e^{i2n\pi t/P}, \quad a_n = \frac{1}{P} \int_0^P \hat{g}(t) e^{-i2n\pi t/P} dt. \quad (17.5)$$

Aside: Other Fourier-Type Series. The complex Fourier series is not the only possibility. For example, if we construct \hat{g} by alternating normal and time-reversed copies of g , we make an even function of period $2P$; the sine components of (17.5) go away, and we have the Fourier cosine series,

$$\hat{g}(t) = \sum_{n=0}^{\infty} b_n \cos \frac{n\pi t}{P}, \quad b_0 = \frac{1}{P} \int_0^P \hat{g}(t) dt, \quad b_n = \frac{2}{P} \int_0^P \hat{g}(t) \cos \frac{n\pi t}{P} dt. \quad (17.6)$$

If we multiply the time-reversed copies by -1 , \hat{g} becomes an odd function of t , producing the Fourier sine series,

$$\hat{g}(t) = \sum_{n=1}^{\infty} c_n \sin \frac{n\pi t}{P}, \quad c_n = \frac{2}{P} \int_0^P \hat{g}(t) \sin \frac{n\pi t}{P} dt. \quad (17.7)$$

The sine and cosine series have real coefficients at half the frequency spacing of the complex series. We aren't doing this for our health, but to look for a way to represent arbitrary functions using finite arithmetic on a computer; therefore we want a series representation that will give us accurate reconstruction of the data with the fewest coefficients (we assume that the Fourier coefficients of our signal will eventually die off with increasing frequency). If g has zero derivatives at the ends of the interval, the cosine series is probably best; if it goes to 0 linearly there, the sine series is worth a try; otherwise, stick with the complex series. Measurement data are seldom well enough behaved to make the sine or cosine series worth worrying about in practice.

17.4.3 Sampled Data Systems

The next question that comes to mind is: What did I do to my spectrum just now? The convolution relation (17.5) is the key. In Chapter 13, we saw that convolving two functions means multiplying their Fourier transforms; thus the spectrum of the periodic continuous function \hat{g} is the product of G and the transform of $\text{III}(t/P)$, which is $(1/P)\text{III}(fP)$. Thus as we see in Figure 17.3, the spectrum we wind up with is a sampled version of the spectrum G , with each sample multiplied by an appropriately placed δ -function to make the integral of \hat{g} come out right.

Functions known analytically can have their Fourier series computed as far as necessary, but we're building instruments—if we knew in advance what the data were going to look like, we could just sleep late instead. What we have to work with is a finite number of samples of a function of finite (but possibly great) extent in time. From elementary linear algebra, we know that you can't get more independent coefficients out of a basis transformation than you put in, so we accept that our knowledge even of the Fourier series will be limited. In fact, what we mean by *spectrum* is a bit problematical, since we saw that the way we go about making the function periodic affects the spectral information we get. We do have the intuition that if we sample a continuous function densely enough, we should be able to reproduce it accurately.[†]

Changing over to Fourier series also leaves us in a bit of a conceptual bind when we think about 1 Hz SNR and bandwidth generally. Both presuppose that frequency can vary continuously, whereas the Fourier series has infinite PSD at discrete frequencies and zero elsewhere. At the risk of getting slightly ahead of ourselves, we see from Figure 17.4 that all is not lost. If we take these spikes and spread them out over the areas in between, then increasing the sampling rate does seem to produce better spectral estimates. Figure 17.4 plots the N -point discrete and true transform of $g(x) = \text{sech}(\pi x)$, which is its own Fourier transform, over the range $-5 < x < 5$. Because the length of the interval in x is constant, the frequency spacing is too; the extra data points from the

[†]Errors resulting from truncating an infinite series, or sampling a continuous function, are generically called *truncation error*, to differentiate them from roundoff error, which is merely a consequence of imperfect arithmetic rather than of the algorithm itself.




Figure 17.3. Sampled data in the Fourier domain.




Figure 17.4. Transform of the function $g(x) = \text{sech}(\pi x)$, and the N -point DFTs of samples on the domain $-5 < x < 5$, for $N = 8, 16, 32$. Note how the spectral estimates improve with finer sampling.

finer sampling go into increasing the frequency range. We need to firm up these ideas, which is what we'll do next.

17.4.4 The Sampling Theorem and Aliasing

In the time domain, intuition tells us that if we take a sufficiently slowly varying continuous function, and sample it densely enough, we ought to be able to reconstruct it very accurately from the samples alone. Our experience with polynomial interpolation (e.g., Taylor's theorem with remainder) suggests that the error we commit might go as some high power of the spacing between samples, but in fact we're luckier than that. Nyquist's sampling theorem, which is the mathematical statement of this intuition, sets out conditions under which the reconstruction is perfect.

Sampling a function $g(t)$ at a rate $\nu = 1/\tau$ can be represented in the continuous-time domain as multiplication of the signal by an appropriately normalized *sha* function $\tau \text{III}(\nu t)$, which produces the multiple copies of the transform, as we just saw. This time the samples are in the time domain and the multiple copies in the frequency domain:

$$g(t)\tau\text{III}(\nu t) \supset \sum_{n=-\infty}^{\infty} G(f + n\nu). \quad (17.8)$$

In general, these copies of G will overlap, a phenomenon called *aliasing* that causes irreparable loss of information. We had a closely related problem with positive and negative frequency lobes overlapping when we discussed modulation—but that was just two spectral lobes; here there are ∞ . This effect is what gives rise to the inaccuracies in Figure 17.4—you can see the computed spectrum flattening out at the ends, where the overlap is occurring.

A component at $f = f_0 + n\nu$ will produce exactly the same set of samples as if it were at f_0 , making the two indistinguishable. On the other hand, if $G(f + n\nu) \equiv 0$ for $n \neq 0$, then no information is lost, and the true transform G is trivially recoverable from the sampled one. Since for real data, $G(-f) = G^*(f)$, we must center the allowable band at 0, so that $G \equiv 0$ for $|f| > \nu/2$. (Such a function is said to be band limited, that is, to have compact support in the frequency domain.) This is the Nyquist condition: the sampling rate must be more than twice the highest frequency present in the data. (Energy exactly at $\nu/2$ is a singular case: the cosine term is sampled correctly, but the sine term disappears.) Note that since we can recover the true Fourier transform from the sampled data, our continuous-time notions of bandwidth and 1 Hz SNR have survived sampling absolutely intact, odd as it may seem.

Really band-limited functions are extinct in the wild, and in fact finite bandwidth is incompatible with compact support in the time domain. We can do very well, though, with analog filters ahead of our sampler, and the expression (17.8) furnishes us with an error estimate in practical cases. The ADC contributes noise of its own, after any filtering has been accomplished. This noise of course aliases down into the fundamental interval $[-\nu/2, \nu/2)$.

In order to make the analog filter easier to build, tune, and test, it's worth sampling a fair bit faster than the Nyquist rate (Figure 17.5). The CD audio specification samples a 20 kHz bandwidth at 44.1 kHz, which is only 10% higher than Nyquist, and would need a really impressive analog filter if the sampling were really done that way (it's actually




Figure 17.5. It’s easier to filter gently, sample quickly, and then decimate. Here the decimation is only $2\times$, but $10\times$ is not uncommon.

done by decimation, see Section 17.8.1). Group delay isn’t a big problem in audio,[†] but elsewhere it usually is, which makes the filter just that much harder to build. If we oversample by a factor of 2 or so, the analog filter becomes much easier, and $4\times$ or greater oversampling permits a really simple filter with loose tolerances. The explosion of nearly redundant data this causes can be controlled by the use of decimation, which we’ll get to.

17.4.5 Discrete Convolution

We saw in Section 1.3.8 that filtering in the continuous-time case is a convolution operation between the input signal and the impulse response of the filter. The same is true in the discrete-time case, with suitable definitions of convolution; we can connect the two domains via the idea of a sampled function being equivalent to a continuous function times a III function. Convoluting two infinite trains of δ -functions is not pellucid to the

[†]A good deal of money has been made by appealing to people’s vanity by telling them that true audiophiles can hear small phase differences. You may or may not have moral scruples about this.

intuition, but a diagram plus a simple limiting argument[†] can be used to show that the convolution of two sampled functions is

$$g(n) * h(n) = \sum_{m=-\infty}^{\infty} g(m)h(n - m) = \sum_{m=-\infty}^{\infty} g(n - m)h(m). \quad (17.9)$$

17.4.6 Fourier Series Theorems

In Section 1.3.8, we stated some of the basic theorems of Fourier transforms. Most of them can be carried along unchanged into the sampled data case, if we use the III function representation of our sampled data. The shifting theorem is identical; the power theorem becomes

$$\sum_{n=-\infty}^{\infty} y(n)x^*(n) = v \int_{-v/2}^{v/2} Y(f)X^*(f)df \quad (17.10)$$

and so does Rayleigh’s theorem, which becomes the Parseval relation

$$\sum_{n=-\infty}^{\infty} y(n)y^*(n) = v \int_{-v/2}^{v/2} Y(f)Y^*(f)df. \quad (17.11)$$

Scaling is a bit more fraught, since the Nyquist criterion must be satisfied both before and after scaling[‡]; as long as we remain within that bound, and within the fundamental interval, it is the same. The derivative theorem is replaced by one for finite differences, which follows directly from the shift theorem:

$$\Delta g(n) \equiv g(n) - g(n - 1) \supset (1 - e^{-i2\pi f\tau})G(f). \quad (17.12)$$

Aside: Trend Removal. It is of course possible to force the data to 0 at the ends, and even force one or more of its derivatives to be 0 as well, by subtracting a suitably chosen function (e.g., a polynomial). This approach can reduce the number of coefficients required to represent the function to a given accuracy. It is blameless in itself but has an unfortunate tendency to be forgotten—especially when we simply remove a trend line, to force the function to go to 0 at the ends. This is equivalent to subtracting the transform of a great big sawtooth wave from the frequency samples, which is a nontrivial modification. If we don’t take account of the effects of the trend line removal on subsequent processing steps, we will be misled, as in the example of Figure 17.6, where removing the trend apparently doubles the frequency of the major component of the data.

If you do trend removal, don’t just throw the linear part away—make sure to carry its coefficients along in subsequent steps. This isn’t too hard, since we can usually apply the subsequent steps to it analytically, and so keep the true measurement data without being forced to use an unnecessarily large number of coefficients. Just be sure you keep

[†]Construct the two III function as limits of two trains of rectangles of unit area whose widths go to 0, and let one set go to 0 before the other one.

[‡]In case you’re wondering why aliasing didn’t destroy Rayleigh’s theorem, when squaring the function doubles the width of its transform, it’s because we’re evaluating GG^* at 0, where the aliased copies don’t quite reach.

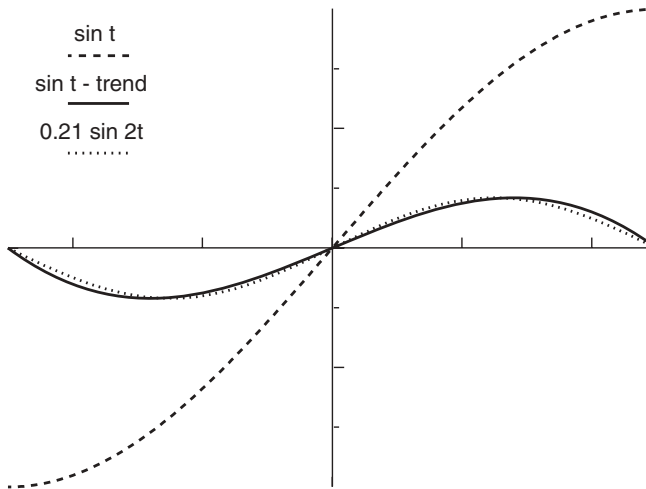


Figure 17.6. Trend line removal can create serious artifacts; here $\sin t$ is transformed nearly into $\sin 2t$.

your normalization straight, and check by comparing the data reconstructed from the big long Fourier series without trend removal to that from the compact form. When we get to spectral analysis with DFTs, having Fourier coefficients that fall off sufficiently rapidly will become important to us, and trend removal can help.

17.4.7 The Discrete Fourier Transform

So far we've used the frequency representation only implicitly. Sometimes we need to know the frequency spectrum of the data, and even more often, filtering is more efficient in the Fourier domain. We've seen that the Fourier coefficients a_n in (17.5) are correct samples of the true continuous transform G of the compactly supported function g , providing we manage to avoid aliasing. Two problems remain. The first one is purely practical: computing a_n appears to require a numerical integral for each sample point, which we probably can't afford; there's a clever algorithm that gets us round this. There is also a conceptual problem that is more worrisome: a finite run of data cannot have a compactly supported transform, so we cannot even in principle get exact samples of the continuous transform G from a finite length of data; the infinity of copies of G will leak into one another at least a bit. Our job is to keep the leakage small and bounded.

The real damage done here is not the approximate nature of our results—we're used to that—but the weakening of our intuition as to what these mysterious coefficients a_n really mean, so that we are less quick to spot nonsensical results. We'll work on keeping that intuition strong.

Once again, we have to change the problem definition to a special case. If we have a periodic, band-limited function $r(t)$, its true, continuous-time Fourier transform can be computed from samples making up exactly one period; the Nyquist criterion requires the number of samples to be at least equal to twice the number of harmonics present (including DC and the fundamental, of course). The result is the discrete Fourier transform, or

DFT:

$$D_n \equiv D(f_n) = DFT(d) = \sum_{m=0}^{N-1} d\left(\frac{mP}{N}\right) \exp\left(\frac{-i2\pi mn}{N}\right),$$

$$d_n \equiv d(t_n) = IDFT(D) = \frac{1}{N} \sum_{m=0}^{N-1} D\left(\frac{m}{NP}\right) \exp\left(\frac{i2\pi mn}{N}\right),$$
(17.13)

where $t_n = n\tau$ and IDFT is the inverse DFT. The IDFT is easier to motivate, in fact, since it's just the Fourier series of (17.5), truncated to reflect the band-limited nature of $d(t)$, and normalized slightly differently. By plugging the DFT definition into the IDFT definition, you can easily verify that the two are really inverses of each other. The normalization is chosen to preserve the useful property that the convolution of two functions is the transform of the product of their transforms.

There is a class of fast algorithms, generically called fast Fourier transforms (FFTs), which can compute the DFT and IDFT in $\alpha N \log_2 N$ operations, where α is around 1, but depends on the algorithm and on whether you count additions as well as multiplications. These algorithms are the most important nontrivial numerical algorithm in current use—the workhorses of DSP. They are fast only for highly composite N , and the simpler and more common algorithms assume that N is a power of 2.

Aside: FFT Routines. Don't write your own FFT routine if you can help it. Lots of smart people have worked on this problem and have spent a lot more time at it than you can spare. Use third party code for this work if possible, such as the open-source FFTW package, or one of several others available on the Internet. These codes all have a bit of a learning curve.[†] Alternatively, the lower performance but much more self-contained *Numerical Recipes in C* code can be a good choice when the FFTs are a small fraction of the execution time of your program. (Note that these routines are *not* in the public domain, but some others of the same sort are.)

Sharp-eyed readers may have noticed that the exponentials used in (17.13) do not obey the Nyquist criterion. They don't stop at $N/2$, but continue up to $N - 1$. What are the higher frequency coefficients for? In fact, they are the negative frequency components and alias down to frequencies between $-v/2$ and 0. Thus in order to identify DFT coefficients with frequency, we have to fold the DFT, that is, split the frequency samples at $N/2$, slide the upper $N/2$ coefficients down in frequency without disturbing them, and attach them to the left of 0 (i.e., sample $N - 1$ is really frequency sample -1 , and sample $N/2$ is really $-N/2$).

There's no deep magic here, it's just that the natural range of the index n is $[-N/2, N/2 - 1]$, whereas we parochially insist on using $[0, N - 1]$ (or even $[1, N]$ if we're crusty old Fortran IV bigots). The odd rearrangement of the data is the price we pay for this. If it weren't for the FFT algorithms' requirement that N be a power of 2, we might want to use an odd value of N to make the positive and negative frequency components more obviously symmetric about 0, and to avoid having to worry about the peculiar point at $N/2$, which is exactly at the Nyquist frequency.

The point $n = 0$ in the time domain corresponds to time $t = 0$. Because of this, if you have a function centered at $t = 0$, you have to do the same rearrangement in the time

[†]Watch out for license terms when using open-source code!

domain, if you want an even function of time to produce a purely real transform; before taking the DFT, split the data at $N/2$, and interchange the segments $[0, N/2 - 1]$ and $[N/2, N - 1]$. This is just a shift of time origin in the periodically extended function, so it isn't mysterious either.

The DFT has one or two important theorems associated with it, especially the DFT version of the Parseval theorem:

$$\sum_{n=0}^{N-1} x_n x_n^* = \frac{1}{N} \sum_{n=0}^{N-1} X_n X_n^*. \quad (17.14)$$

Aside: DFT Blunders. DFTs are very blunder-prone in practice. We get the folding wrong, load the data one sample off, or forget that the indices go from 0 to $N - 1$ with 0 counting as a positive frequency sample. Fortunately, it's easy to check that you've got the rearrangement right: try taking the DFT of a real, even function with no flat places, and make sure the transform is real and symmetric as well. If you've got 0 in the wrong place, the transform will have symmetric magnitude but will have a phase ramp on it, and if you've flopped the negative times or frequencies incorrectly, it won't be real and symmetric. Really bang on it to make sure you've got it right, and whatever you do, don't just poke things until it seems to work. Rearranging for even-order DFTs is easier—you just do a circular shift of $N/2$, that is, sample n in normal order is sample $(n + N/2) \bmod N$ in DFT order, and the operation is symmetrical—to shift back, just apply the same shift again. For odd-order DFTs, which are sometimes useful, it's a bit more subtle: forward and backward shifts can't be the same since they have to add up to an odd number. When the DFT returns, DC is at sample 0, whereas in normal order DC is at sample $(N - 1)/2$. Thus sample n in DFT order is sample $(n + (N - 1)/2) \bmod N$ in normal order, and sample n in normal order is sample $(n + (N + 1)/2) \bmod N$ in DFT order. (See Figure 17.7)

17.4.8 Does the DFT Give the Right Answer?

No, not exactly, unless you really live in the special case above (if you do, send a postcard); but as usual we can fix it up so that it does, in principle, to any accuracy we like. The main problem is spectral leakage, where energy that belongs in one bin appears in another. Leakage arises in two ways: overlap of the true Fourier transforms of the compactly supported data we actually have (aliasing), and frequency components that don't go through an integral number of cycles in the time P .

The aliasing problem is easily spotted in practice; generally speaking, if your computed DFT coefficients go to 0 faster than linearly as you approach $N/2$ from both sides, overlap leakage is not serious (taking more samples will help somewhat if it is). Serious aliasing usually leads to the a_n flattening out to some nonzero magnitude near $N/2$, although phase cancellations will often lead to peaks and nulls rather than a smooth approach.

Components at frequencies between frequency samples is usually a worse problem. It leads us to the idea of data windowing, which will solve both problems more or less completely.

17.4.9 Leakage and Data Windowing

Any given data set must contain a finite number of samples, drawn from a finite length of time, which violates the assumptions of Fourier analysis. Truncating the real measurement

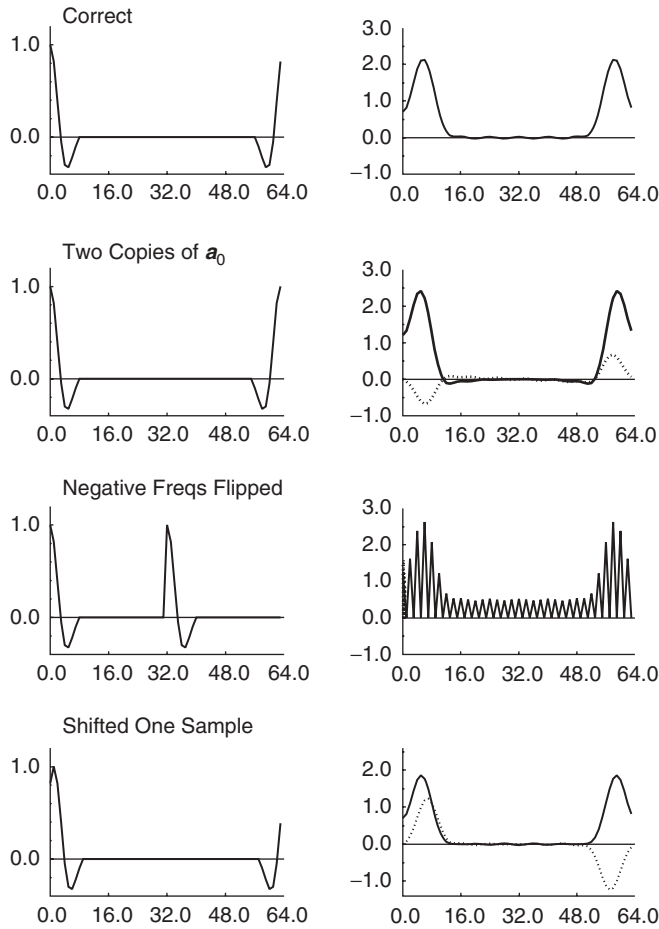


Figure 17.7. DFT blunders are easily spotted experimentally, using even and odd test functions and insisting that the theorems hold. Even big mistakes like these ones can easily go unnoticed with arbitrary data, which will turn your measurement into nonsense.

data in time this way is mathematically equivalent to multiplying it by a rectangle function; thus by the time we get our hands on it, the true continuous transform G has been convolved by the sinc function transform of the truncation window as well as by the shah function from the sampling. The sinc function convolution leads to nasty ripples and other artifacts in the transform, which are familiar in mathematical physics as the Gibbs phenomenon and in physical diffraction as Airy rings (see Section 5.6).

Convolution is commutative and associative, because it is multiplication in disguise; thus the transform we actually get our hands on out of the DFT is

$$DFT(g_n) = G * (Psinc(fP) * \text{III}(f\tau)). \tag{17.15}$$

The III function represents the wraparound, and the sinc function the leakage due to truncation in time. Thus each sample of the DFT is equivalent to the output of one

channel of a bank of identical filters, each with a sinc characteristic in frequency. If we slide the DFT—that is, recompute it every time a new time sample arrives (letting sample N fall off the end each time), we get N filter channels, each with an identically shaped sinc function filter characteristic, and each of which has the same data rate as the input.

The sinc function dies away only slowly as we move out in frequency and, in fact, is itself a source of significant wraparound as well as leakage into nearby channels. Figure 17.8 shows the problem; the sinc function centered on sample m is zero at all the other samples, so there is exactly 0 leakage provided that the function really contains only harmonics of $1/P$. However, we don't live in that case really, and as an estimator of the real continuous spectrum we care about, it stinks. There is a workaround for this problem, which is called *data windowing*.

17.4.10 Data Windowing

We saw that the output of the DFT is samples of the true spectrum G after corruption by a repeating sinc function, and that the sinc function arose from multiplying the true time series by $\text{rect}(t/P)$. If we don't like the sinc function, we'll have to change the rect function, and weight the data samples unequally, as shown in Figure 17.13a,b. A visual metaphor for what we're doing to the data is to imagine samples trooping by outside our office window, abruptly appearing at one edge and disappearing at the other; we need to make it appear more gradually by tinting the edges. Since this effectively narrows the sampling window somewhat, we can expect the FWHM frequency resolution to deteriorate accordingly, but that's a good trade if we can avoid wholesale smearing by the sinc function. As shown in Figures 7.9 and 17.10, there are a whole bunch of different windows you can use, which have different trade-offs of main lobe width (i.e., frequency resolution) versus close-in and distant sidelobe levels. Which is best depends on your intended use. A good general purpose window, and one that is easy to remember,




Figure 17.8. Each sample of a DFT is equivalent to the output of a filter whose transfer function is $\text{sinc}(f - n\tau/N)$. This function is 0 at all the other sampling points, but not in between them.




Figure 17.9. Window functions (128 points wide).

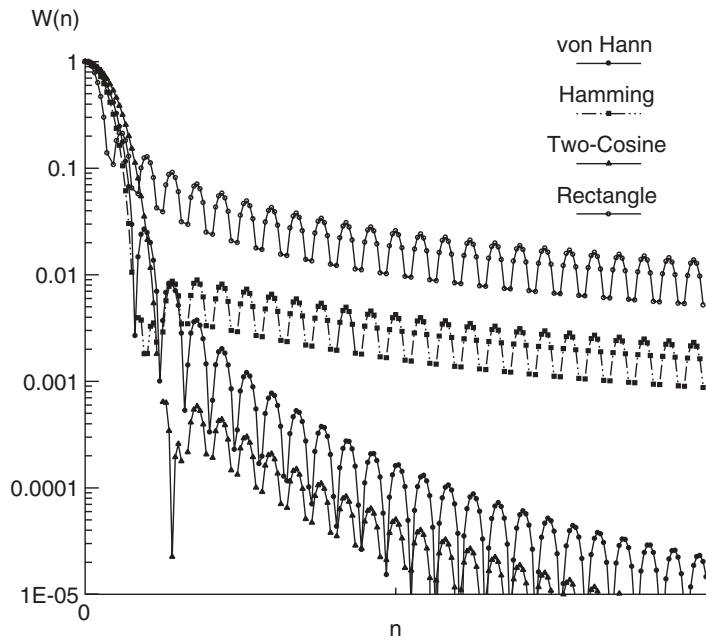


Figure 17.10. Transforms of the window functions of Figure 17.9 (log magnitude). The trade-off between main lobe width and sidelobe selectivity is apparent.

is the von Hann raised cosine

$$W_V(n) = 0.5 - 0.5 \cos\left(\frac{2\pi n}{N}\right), \quad (17.16)$$

which goes to 0 quadratically at the ends, $n = 0$ and $n = N$ (remember the periodicity). The main lobe of its transform is about twice as wide as the sinc function, but because of its smoothness, its transform falls off strongly at large offsets.[†]

The first sidelobe is about -30 dB, much better than the -18 dB of the second sidelobe of the rectangle (which is at the same frequency offset), but no great shakes if you need to separate two closely spaced sinusoidal components of very different amplitude. To get better close-in sidelobes at the expense of ultimate rejection, try the Hamming window,

$$W_H(n) = 0.54 - 0.46 \cos\left(\frac{2\pi n}{N}\right), \quad (17.17)$$

which is a linear combination of the rectangle and von Hann, chosen to cancel the first sidelobe of the von Hann. Its worst sidelobe is the third, which is more than 40 dB down. Elementary DSP books all have lots on windows, although usually with emphasis on their use in digital filter design. The two-cosine window is a linear combination the von Hann and its square, in the ratio of 26.5%:73.5%, chosen to minimize the largest sidelobe—it's an application of the Hamming technique to the second derivative at the patch point. This sort of window-cobbling procedure can help a lot in special situations.

The asymptotic falloff of the true continuous transform of the window function is governed by the window's order of continuity (including its ends); if the m th derivative is continuous, the transform falls off at least as $f^{-(m+1)}$. Removing a trend line from the data will get you an extra 1 in this exponent, at the expense of needing a separate scheme to keep track of what happened to that sawtooth wave in subsequent steps.

A properly windowed DFT, with enough samples that the transform has smoothly decreased by at least 60 dB by the time you reach $f = \pm\nu/2$, will give you good results most of the time; you can trust it about as well as a spectrum analyzer. Do pay attention to the normalization of your spectra, and remember that the spectral samples are no longer perfectly independent, as they were in the sinc function case; just because your one sinusoidal component is now spread over 3 bins doesn't necessarily mean that the total power is equal to the sum of the squares of the frequency samples, unless you've normalized your spectrum that way.

A final practical matter is that we don't really want to increase our data rate and computational burden N times by using a sliding DFT. Each frequency tap has its corresponding filter, and these filters are much narrower than $\nu/2$. Accordingly, they don't need to be sampled at the full rate ν . You can decide how often they need to be sampled by a straightforward aliasing argument, exactly as we did in the first place. A good window function will allow sampling at a rate of perhaps $4/P$ (i.e., $4/N$ times as fast as the incoming data), which corresponds to hopping the window over by $N/4$ samples at each DFT evaluation. This is somewhat surprising at first sight—we all know that you have to overlap lawn mower cuts by $1/4$ or so to do a good job, but that adds only 33% overhead; computing an FFT where $3/4$ of the points have been transformed before

[†]Here as elsewhere we assume that your data has not yet been rearranged for the DFT, if you're going to—always window first, or you'll get “an unexpected result” as they say in software manuals (i.e., garbage).

quadruples the computational burden, which seems very wasteful. Some windows and some problems allow overlaps of only $N/2$, but there's a significant accuracy trade-off, especially if you have a few strong components. If you aren't convinced, use your favorite math program to explore the consequences of reusing too few points, perhaps subtracting off a polynomial first to increase the order of continuity at the ends.

This leads to a piece of advice offered with great earnestness: whatever scheme you come up with, kick the tires hard, all four of them. Plot its expected performance out carefully in both time and frequency, SNR and signal amplitude, looking for aliasing due to the original sampling and to hopping the DFT, leakage due to inadequate window functions, scallop loss,[†] and (especially) blunders. Keep the results, and use them to check your real code for blunders and roundoff.

Aside: Negative Window Coefficients. The windows we have discussed are all positive. The frequency response of the equivalent filter can be sharpened up considerably by the use of negative coefficients as well (e.g., the sinc function, which has a brick wall characteristic in frequency). Should we use negative window coefficients? Well, yes, occasionally.

The noise gain of a digital filter (i.e., total power out over total power in, with a white noise input) is equal to the sum of the squares of its coefficients, while the signal gain goes as the square of the sum. Negative window coefficients thus tend to increase the relative error of the spectral estimates in the presence of noise, so you'll often hear it said that you should avoid them. There's less to that than meets the eye, though; Parseval's theorem says that the noise gain (see Section 13.1) in the time and frequency domains are equal, so for a good choice of window, this increased noise gain will mainly be due to squaring up the shoulder of the filter, which is what we want to happen (Figure 17.11).

In principle, a more general window can reduce the sampling rate required for the individual taps, and so reduce the total volume of data. You'll need to use longer transforms, though, because the main lobe of the window function will of course be much narrower to make room for the extra lobes. This increases the total computational burden, more or less eliminating the efficiency improvement we sought.

Sometimes, though, our data will characteristically exhibit a few very strong Fourier components, and our measurement requires measuring weak signals accurately, or we need to be able to measure the amplitudes of sine waves that aren't periodic in the sampling window with no scallop loss. In that case, the great improvement in the sidelobes that negative coefficients make possible will become very worthwhile, which is why they're used in FFT spectrum analyzers.

17.4.11 Interpolation of Spectra

One nice thing about the DFT is that you can get finer frequency sampling by zero-padding your array to the number of frequency samples you want, and then doing a bigger DFT (Figure 17.12). Since the frequency limits of the spectrum are always $\pm v/2$, the extra coefficients come from sampling the spectrum more finely. This introduces no

[†]Scallop loss is the error in the calculated channel power caused by the windowing, hopping, and the departure of the equivalent channel filters from being perfect rectangles. A plot of overlapped filter functions has a scalloped appearance about the peaks, as in Figure 17.8.

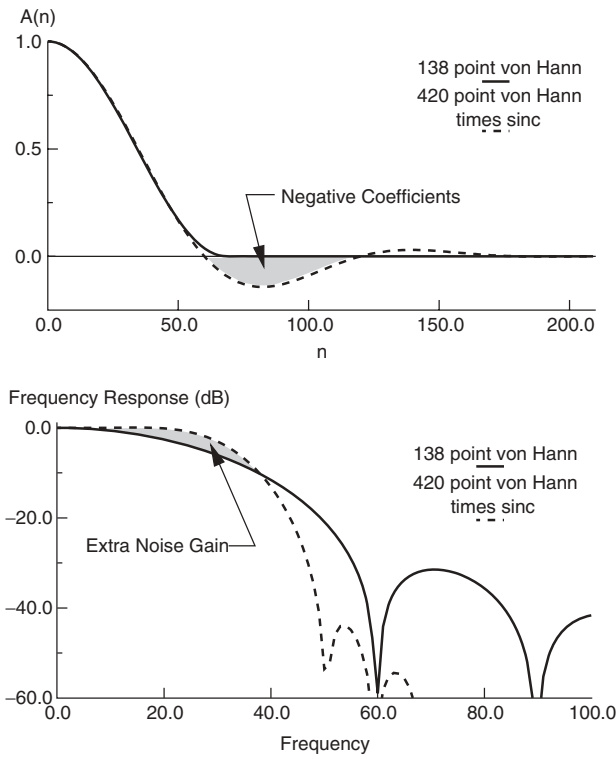


Figure 17.11. Negative coefficients increase the noise gain of the window, but can increase its selectivity. Here the extra noise gain comes from squaring up the shoulder of the frequency response and reducing its sidelobes.




Figure 17.12. Interpolating with a zero-padded DFT: (a) a 31-sample rectangle function and (b) its DFT of 128 points (boxes) and 512 points (solid line).

additional artifacts whatsoever, but it may reveal some that were hiding, especially sinc function leakage due to not having windowed the data.[†]

The innocuous nature of zero-padding makes the usual FFT restriction that N be a power of 2 no problem in practical cases. Zero-padding doesn't mean you can window your 128 data points with a 2048 point Hamming window (why?), nor does it improve the intrinsic resolution of the measurement, which is of course limited by the transform of the window function you use. Zero-padding does reduce scallop loss, since each frequency sample has a filter function centered right on it.

A related benefit is that if we have two data sets that are slightly out of registration, we can shift one with subpixel resolution by taking the FFT, putting the appropriate linear phase ramp on the frequency samples (from the shift theorem), and transforming back. This is especially useful in registering images, or in permitting signal averaging when the scan range is drifting around a bit, due e.g. to Doppler shift.

17.5 POWER SPECTRUM ESTIMATION

Estimating the power spectrum of some signal is a very common thing to want to do. Color temperature meters, spectrometers, RF spectrum analyzers, and frequency counters all perform power spectrum estimation in different ways, with different sets of assumptions about the input function. With proper attention to normalization, you can compute n samples of the power spectrum of $2n$ data points by taking the squared modulus of the DFT (negative and positive frequencies give the same modulus so you only need half of them). The normalization will often involve dividing by the number of data points, but that depends on just what you're trying to compute; if you're interested in the total energy of a certain run of data, you'll just take the sum of the squared moduli.

The key limitation here is noise. Taking the DFT of a short run of your signal will give you the correct Fourier series representation for that particular run, but unfortunately that is quite different from the power spectrum of the underlying random process (signal plus noise). If you take zero-mean white Gaussian noise and square it, the standard deviation in each sample is equal to the mean value. The same is true of its DFT power spectrum: the standard deviation in each frequency bin of the PSD is equal to its mean regardless of how many data points you take; all the extra information you supply goes into sampling the power spectrum more densely, not increasing its accuracy. This makes naive use of the DFT useless in estimating the power spectrum of a noisy signal. (See Figure 17.13.)

Aside: Is There a Right Answer?. Even without additive noise, power spectrum estimation is an ambiguous business. Say we have a continuous function of time that has no obvious beginning and ending, for example, Earth-surface solar irradiance on a 1 cm^2 patch at the North Pole. We can take only a finite number (though perhaps a large one) of samples of this function, with limited accuracy, and we want to know the power spectrum of the whole; we are willing to assume that our run of data is in some way typical. Typical in what way? A well-defined answer is not always easy to give. Our example has important time variations on all scales from the age of the Sun

[†]This may surprise some readers, who are used to hearing that the big jumps in the data from 0 to $g(0)$ and $g(N-1)$ cause artifacts—but consider taking the $2N$ point transform of N data points; (17.13) shows that sample $2n$ is identical with sample n of the N -point DFT.




Figure 17.13. DFT power spectrum estimation: (a) short run, (b) von Hann windowed, (c) rearranged for DFT, (d) short-time PSD, (e) PSD of $64\times$ longer run, and (f) summed $64\times$.

and continental drift, through cloud motions and the shifting shadows of snowflakes, a frequency range of perhaps $10^{19}:1$. We can perhaps sample 1 part in 10^9 of the whole time series. Looked at too closely, even that short run is somewhat ill defined, because the surface moves around, for example, with snowfall and ice pack motion. Our strategy will obviously have to depend on what we want to get out of the data that we have—what is noise to radio communications people may be signal to astronomers.

The mathematical power spectrum estimation problem assumes that we have a finite number of data points, quantized at some resolution; that we want to estimate the power spectrum of the true continuous random function from which they were drawn, without reference to times before and after; and that we know that the signal's time autocorrelation function is well represented by one member of some limited class of model functions (trigonometric polynomials in the case of the DFT). We usually assume that the function is ergodic. This is a well-defined problem, which can be solved efficiently by a number of methods. The key thing to remember is that this isn't the same problem as the one we really want to solve, and unless we choose the sampling strategy and the set of model functions with care, the answers it gives will be different as well. Watch out for reasonable-looking wrong answers, here as elsewhere.

17.5.1 DFT Power Spectrum Estimation: Periodograms

One way to combat the noisiness of the power spectrum estimate is by averaging. There are two main ways to do this, either summing the energy in n adjacent bins in a big long $2N$ point windowed transform, or by shift-and-add: split up your data points into n chunks of $2m$ points (remember to overlap the chunks by 75% or so), window, transform each one, compute the squared modulus to get n estimates of the m point power spectrum, and sum the results. Shift-and-add evaluation of power spectra is also called the *method of modified periodograms*.

Both procedures reduce the variance by a factor of the number of independent data points, more or less. Both are reminiscent of a signal averager; summing adjacent bins is something like signal averaging the autocorrelation of the function in the time domain, whereas shift-and-add is like signal averaging the power spectrum in the frequency domain. Of the two, shift-and-add is modestly cheaper in CPU cycles and less wasteful of data, since instead of a single big wide window, you're using many overlapping copies of a narrower one; thus almost all the samples have a chance to be in the middle of a window, and so contribute equally to the measurement. Binning the spectrum is a lot easier to code, though, and less blunder-prone as well. Use shift-and-add, but use binning as a sanity check during debug.

If your data are unequally spaced, or you have gaps in them, these simple algorithms are inadequate. For those cases, try Lomb's method (see, e.g., Rabiner and Gold).

17.5.2 Maximum Entropy (All Poles) Method

There's nothing intrinsically more accurate about using trig polynomials as fitting functions. In approximating functions we usually prefer rational functions, because they are a lot better at reproducing asymptotes (vertical or horizontal) or other non-polynomial-ish features. Power spectra often have similar behavior (e.g., isolated Lorentzian peaks), so rational functions of $\exp(j2\pi ft)$ seem like good candidates for model functions. If we make our rational function the reciprocal of a trig polynomial, this idea leads to the *maximum entropy method* (MEM) or *all-poles method* of Figure 17.14, where

$$f(t) \approx \frac{1}{\left(\sum_{k=-N}^N a_k e^{j2\pi kft}\right)}. \quad (17.18)$$

Note that this estimate is not band limited.

The maximum entropy method is a good way of pulling sharply peaked things out of noise; it amounts to putting in a priori knowledge that it's a few sharpish peaks in the frequency domain that you're looking for. If that's the case, you'll often get good estimates with relatively few terms, say, 5 or 10 times the number of peaks you expect. Don't go above the square root of the number of points, or you'll start getting a lot of spurious peaks—MEM is a bit like polynomial fitting that way. It is also very poor at fitting data sets containing pure tones—the condition number gets big. The *Numerical Recipes in C* MEM program doesn't tell you the condition number, so some numerical experimentation is required. Because the fit is nonlinear, the condition number depends on the data set. A reasonable check is to add some pseudorandom noise and recompute, looking for large deviations in the reconstructed spectrum.

If your data contains a strong pure tone which is cut off sharply at the ends of the interval, the true Fourier transform will have strong sidelobes, which will look like

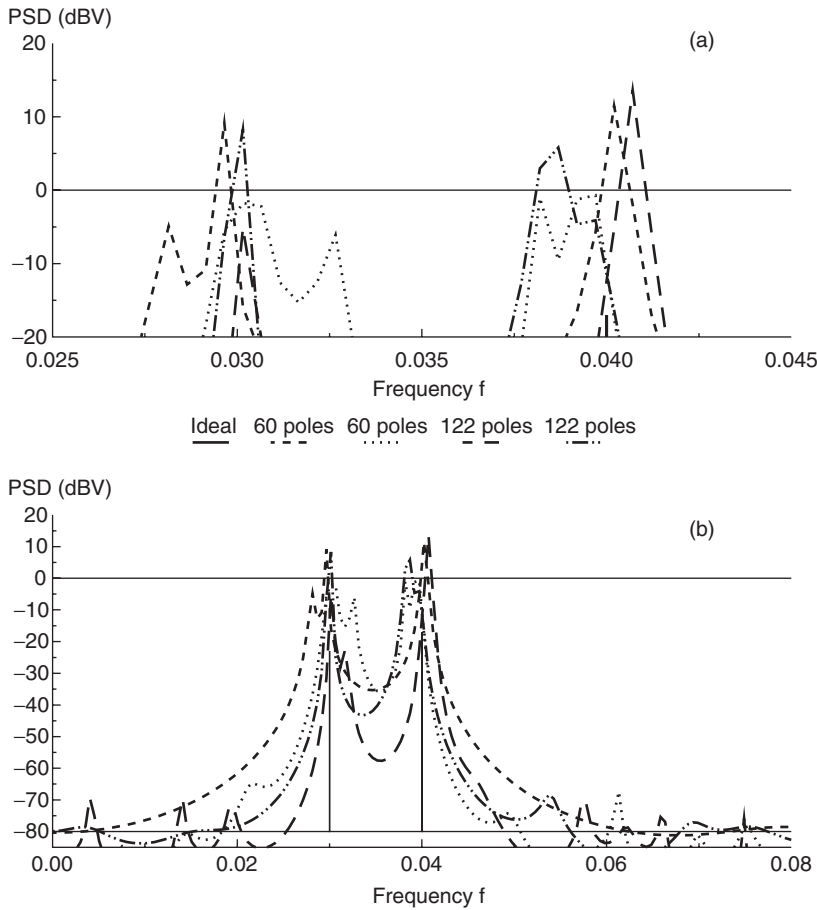


Figure 17.14. MEM power spectrum for 256 samples of a sum of sine waves, 0.1 V at $f = 0.03\nu$ and 0.2 V at $f = 0.04\nu$, plus 100 μV rms noise, with 60 and 122 poles and two starting phases: (a) peaks and (b) full range.

additional poles and which will fall off slowly with frequency. Accordingly, negative and positive frequency lobes will overlap, causing peak shifts that are sensitive to the phase of the tone in the data. A good test is to move the data around a bit; if you have 377 data points, try comparing the MEM results from points 19 to 376 with those from 0 to 357. The phase shifts that result will shift the sidelobe overlap in the frequency domain. If the peak position moves significantly, you have an overlap problem.

Besides being ill conditioned, MEM takes quite a lot of computation and is not easily extended by averaging, periodogram-style, because the average of two functions with M poles is in general a function with $2M$ poles. If you have lots of data and are CPU limited, it may pay you to use a simpler technique on a bigger data set.

Example 17.2: Film Thickness Measurement with a Fiber Spectrometer. Thin dielectric films tend to produce optical spectra with lots of peaks and nulls, where the film thickness is an even or odd multiple of $0.25\lambda/n$, respectively. One good way to measure

film thickness in situ is to use an optical multichannel analyzer (OMA) spectrometer, fed with a fiber bundle (see Section 10.5.4). From a signal processing point of view, the difficulty is that the detection bins in the OMA are not equally spaced in frequency; they tend to be at equal intervals in wavelength instead. Over the 2:1 range of a typical spectrometer (limited by grating orders overlapping), the reciprocal relationship stretches the low optical frequencies out over more bins than the high frequencies get. There's nothing intrinsically wrong with this.

It is inconvenient for film thickness measurements, though, because a film produces fringes that are equally spaced in optical frequency, and when sampled at equally spaced wavelengths, the frequency is not constant, but chirped.[†] There are several techniques possible for pulling such a chirp out of the noise. The most straightforward conceptually is to resample the data onto equally spaced frequencies and do a DFT. This is not as easy as it sounds, since an OMA spectrometer's output really doesn't represent band-limited data in all probability, and there aren't that many data points (often 512), and the spectrum is not always oversampled. Make sure you window the data correctly, or sidelobe leakage will be troublesome. Resampling is somewhat fraught with difficulties, as we saw earlier, and is prone to give badly wrong answers if you don't oversample enough.

Other techniques are possible. You can also use brute force correlation and fit a mathematical chirp to the data, using some reasonable criterion such as minimizing the absolute value of the residuals, or maximizing the (normalized) cross-correlation. This takes a lot of CPU cycles compared to the DFT, especially if you don't already know something about the thickness and refractive index you expect. It also suffers from the tendency of correlation methods to exhibit multiple peaks and to jump about as the height of one peak passes that of the previous champion, perhaps due to a phase shift in the chirp. It will work well if there are many cycles in the chirps (i.e., thick films), so that the subsidiary peaks are well suppressed.

Still another approach is to resample and use the maximum entropy method to pull out the peaks. This seems to be less sensitive to the artifacts from resampling than the FFT method, but is much more CPU intensive, which may matter if you're doing fast measurements with a slow computer (industrial process control applications often have to run on some overloaded old PC that makes an 8088 seem fast).

If there is a whole lot more data available than you can process in the time available, you can use a simpler approach: filter gently to get rid of the high frequency noise (use a symmetrical FIR filter to avoid phase shifting the chirp), and then just use local interpolation to find the peaks of the chirp. The peak positions provide independent estimates of the film thickness, so that a peak that has been estimated incorrectly (perhaps due to noise) will stick out. Averaging a whole lot of these values may get you better measurements than a fancier approach to a small data set, especially if there is intrinsic variability to the sampled data—for example, in polishing substrates, where the thickness is bound to be somewhat nonuniform.

The best advice in a case like this is to live with your data set for a while and try several things. Look for average-case accuracy, but use enough data that you know about how often rogues come along. Save all the rogues to disc for testing your refined algorithms. Above all, don't be a signal processing snob—use what works, know why it works, and be able to prove experimentally that it works.

[†]Sinusoids whose frequencies change with time are called chirps after bird songs, which have this characteristic.

17.6 DIGITAL FILTERING

Most modern optical systems produce digital data at some point, and this trend will only continue. Digital filtering is so cheap, flexible, and powerful that it is an indispensable part of the equipment of a modern electro-optical systems designer. Newcomers to this art should go look at Hamming's undergraduate book on digital filters. He goes through this stuff with lots of hand-holding and comforting detail, which is very important since there are lots of ways to get digital filters wrong.[†]

For designing digital filters, commercial software is available, but nevertheless it is highly worthwhile to learn about z -transforms. They are not difficult to learn or to use, and in providing a conceptual bridge between the continuous-time and sampled domains, they really make digital signal processing easier to understand.

In this book, we'll just dip our toe into the subject and talk about designing *finite impulse response* (FIR) filters[‡] using windows, which is a very broadly useful method. FIR filters whose responses are symmetrical in time (e.g., a symmetrically windowed portion of $\text{sinc } x$) have exactly constant delay with frequency, which is an enormous benefit; a symmetric N -point filter has a delay of exactly $(N - 1)\tau/2$ seconds. Unlike recursive filters,[§] FIR filters raise no stability issues and exhibit no *limit cycles*[¶] due to roundoff error. In situations where phase linearity is not a requirement, recursive filters are more efficient, and the efficiency gain increases with increasing filter order. On the other hand, optimal FIR filters are more efficient than delay-equalized recursive filters.

It is usually best to use filters of odd order, so that the delay can be an integral number of samples, but some types of filters (e.g., differentiators) work better in even order. Unless your system has an embedded DSP, it's usually best to do your filtering in software on a host PC—it's much easier and more flexible.

Aside: Zero-Order Hold. One sometimes overlooked effect of converting from analog to digital and back again is time delay. We usually don't ignore the one-clock delay between the CONVERT command to the ADC and latching the data into a register, but sometimes forgotten are the delays caused by finite track/hold bandwidth (so that our sampling time is slightly before the track-to-hold transition) and the zero-order hold on DACs. Remember that our whole treatment of sampled-data systems was based on a series of mathematical impulses, whereas real systems use rectangular pulses of width τ instead, the *zero-order hold* of Figure 17.15; in sampled-data terms, a DAC convolves the ideal output with $\text{rect}(t - \tau/2)$, which is a filtering operation. As we know, this produces a delay of $\tau/2$ and a sinc function rolloff with frequency; this is a large effect up near the Nyquist frequency, so don't neglect it in your designs (especially with actuators).

[†]For example, computing the FFT, chopping off the high frequencies, and transforming back, or automatically removing a trend line before filtering. If it isn't obvious why this is silly, go read Hamming.

[‡]Finite impulse response means "finite-length impulse response," that is, a filter that can be implemented by discrete convolution with a finite-length function h .

[§]Nonrecursive (FIR) filters produce a weighted sum of their inputs, whereas recursive (IIR) filters use the previous values of their outputs as well; this is the equivalent of a linear circuit with feedback, and leads to poles in the transfer function, whereas FIR filters have only zeros.

[¶]A limit cycle is a self-sustaining output sequence in the absence of an input signal.

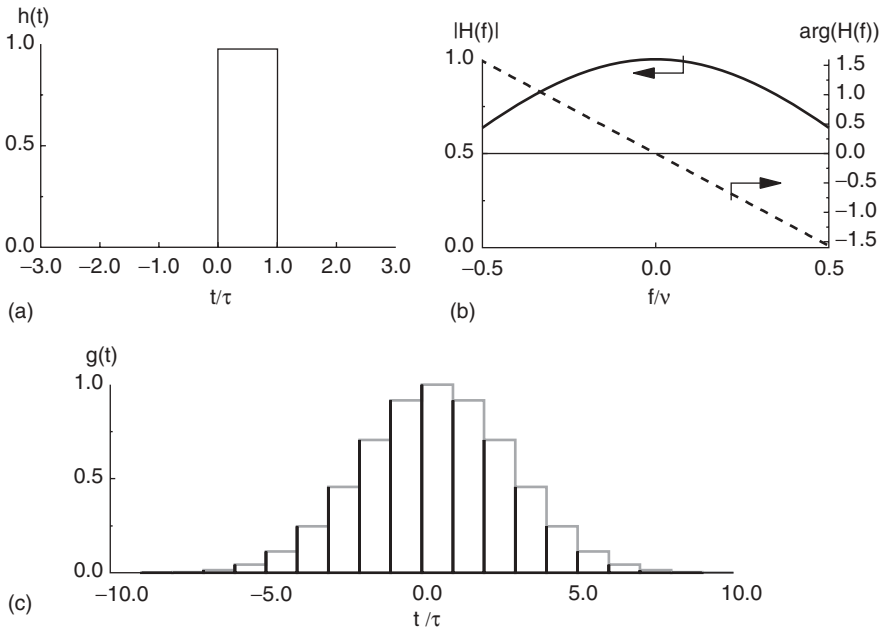


Figure 17.15. A zero-order hold produces a sinc rolloff and a half-sample delay: (a) the continuous-time smoothing function, (b) transfer function, and (c) result of the hold operation applied to delta-function samples.

17.6.1 Circular Convolution

We alluded to the use of the DFT plus multiplication to perform discrete convolutions, but we can't do an infinite sum with a DFT to evaluate (17.9), so how does that work?

Well, remember that the DFT gives correct samples of the aliased version of the Fourier transform of our continuous given function, assuming that the given function is periodic (the normalization is also changed to produce finite results). Thus we expect multiplying the DFTs of two functions together and taking the IDFT will produce correct samples of the convolution of the two periodic functions, which it does; the result is called the *circular convolution*. Since the width of the convolution is in general the sum of the widths of the functions, each period will smear out into two periods in general, producing serious wraparound errors everywhere.

In order for the circular convolution to produce the correct results for our discrete convolution, we'll need the functions to have compact support, and for them and their convolution to be able to fit in the N points available. Since the convolution of two sequences of lengths L and M is $L + M - 1$ samples long, the two must obey

$$L + M \leq N + 1, \tag{17.19}$$

or wraparound will occur (if it isn't obvious where the 1 comes from, consider convolving with a one-point function, which is the discrete version of a δ -function). Providing (17.19)

is obeyed, multiplying the N -point DFTs of two functions and taking the IDFT of the result yields their correct discrete convolution without wraparound.[†]

This discrete convolution algorithm is the basis of FIR digital filtering, which is convolution with the impulse response of the filter function, just as it was in analog. The way you do it in practice is to choose N large enough, zero-pad both functions up to length N , multiply the transforms, transform back, and unfold the resulting data. (Try it out with cases you know a priori before trusting that your code is correct—near enough is *not* good enough.)

Once again, apart from aliasing and the effects of windowing, the discrete convolution yields samples of the continuous convolution of the continuous functions g and h .

17.6.2 Windowed Filter Design

A brick-wall lowpass filter that cuts off at $\pm f_0$, $\text{BW}(f) = \text{rect}(f/(2f_0))$, has impulse response

$$g(t) = \frac{\sin(2\pi f_0 t)}{\pi t}. \quad (17.20)$$

Since this is a band-limited function, we can sample it without aliasing at any rate $f_s > 2f_0$, so obtaining an infinitely long impulse response. As a practical matter, we need to cut it off at a finite number N of coefficients. Chopping off the impulse response sharply leads to artifacts and ringing in the frequency domain as we saw when we discussed data windowing in Section 17.4.9, so as we did there, we make $g(t)$ go smoothly to 0 at $t = \pm T$, by multiplying by a window function (Figure 17.16). Providing T is at least a few times $1/f_0$, the approximation is quite reasonable, and the effects of windowing on the transfer function are easily checked via DFT. This is a good technique for generating FIR filters; the filters so obtained are not optimal, but they are great for getting something to work quickly and for making trade-offs of bandwidth, measurement speed, and noise. Once you get close, you can use some optimizing program to make the perfect filter.

17.6.3 Filtering in the Frequency Domain

We saw how to specify a filter in the frequency domain in Sections 1.8.1 and 13.8. Digital filter specs are similar, except that more attention is often paid to phase linearity, because we don't have the tuning problem with digital filters—what comes out of the design program is the filter itself, not just the L and C values. Also, at least with FIR filters, it's easy to get perfect phase linearity: just make the filter coefficients symmetrical, so that in an $(N + 1)$ -point filter, $C_n = C_{N-n}$.

Because the FFT convolution algorithm is fast, frequency-domain filtering is generally preferable. There are important special cases, though, where you should do it in the time domain: real-time systems, where it is unacceptable to wait N sample periods before producing output; filters with only a few coefficients, which take less than the several times $\log_2 N$ floating-point operations per point of FFT convolution; or where you need very sharp skirts, and don't care about the group delay, which strongly favors elliptic IIR filters with very long tails in time.

[†]You can also use a shift-and-add technique, which fixes up the wrapped parts afterwards, but there's no significant efficiency gain, and it's more blunder-prone.

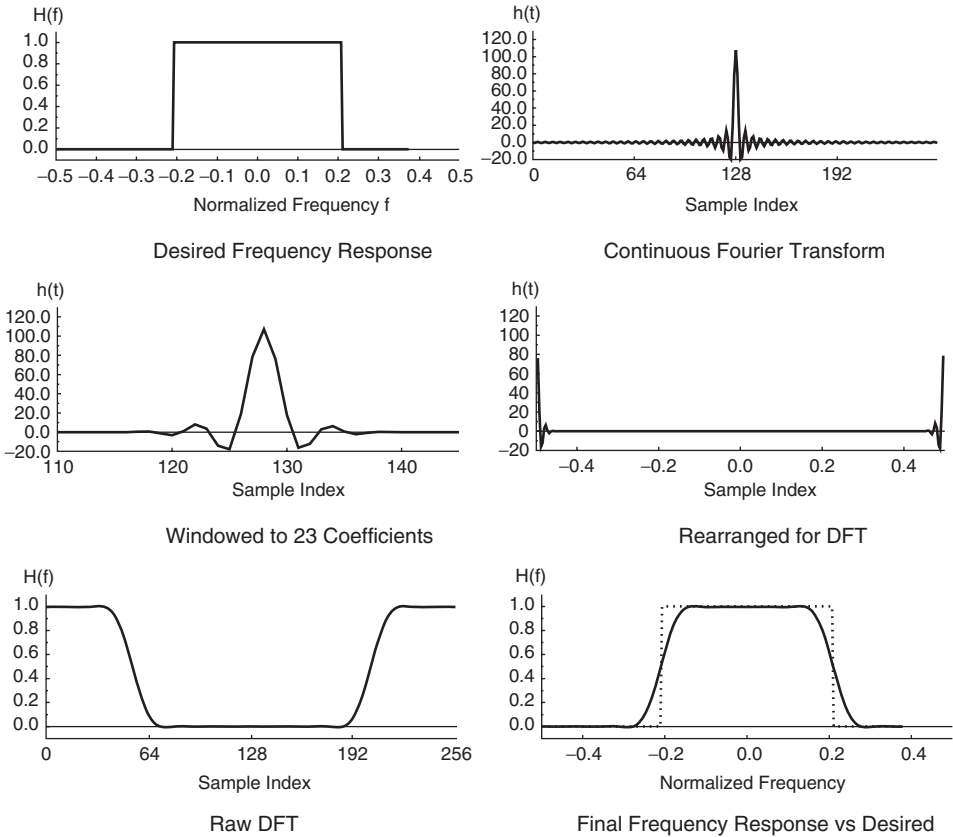


Figure 17.16. Windowed FIR filter design.

17.6.4 Optimal Filter Design

There are packages for digital filter design, which produce results that are optimal in some precise sense. Simple ones use a sampled-frequency representation, with the coefficients in the transition band left indeterminate; a numerical optimizer then optimizes the impulse response (i.e., the filter function in the time domain) by minimizing the number of coefficients required, handling ringing and overshoot, minimizing time delays, or whatever combination you specify. The most widely used FIR filter design program is the classical one by Parks and McClellan, originally published in Fortran but available as public-domain C code in various places on the Internet; search on “Parks–McClellan” and you’ll find it. It’s quite a good program, supporting multiple passbands and stopbands with different (specified) gains, and different weights for the rms error in each band.

More advanced packages can find the true optimum filter, including minimizing sensitivity to roundoff error. You just have to be careful what you wish for, because some of these theoretically optimal filters may have properties you don’t like, for example, deconvolution filters with sharp frequency peaks that make noise look like ringing.

17.7 DECONVOLUTION

17.7.1 Inverse Filters

Ordinary lowpass and bandpass digital filtering are the workhorses of DSP, but especially in electro-optical instruments, there is a uniquely valuable function digital filters can perform: eliminating obnoxious instrument functions. Calibration can be looked at as doing this in the real-space domain, for example, by removing the effects of etalon fringes or the 2.2 μm absorption band in the quartz bulb of a tungsten-halogen lamp. There are other systems such as phase-sensitive confocal microscopes (whose coherent transfer functions are nearly triangular instead of flat), or phase-sensitive systems suffering from defocus, where the instrument function is well known a priori, and the full complex data are available.

An inverse filter can produce results close to magic in these sorts of cases (Figure 17.17): just take the transfer function you want, divide it by the one you've got (window the result appropriately to avoid dividing by very small numbers), compute the Fourier transform of the result, take a whole bunch of samples, window appropriately, and you are in possession of a digital filter that can turn some ugly transfer function

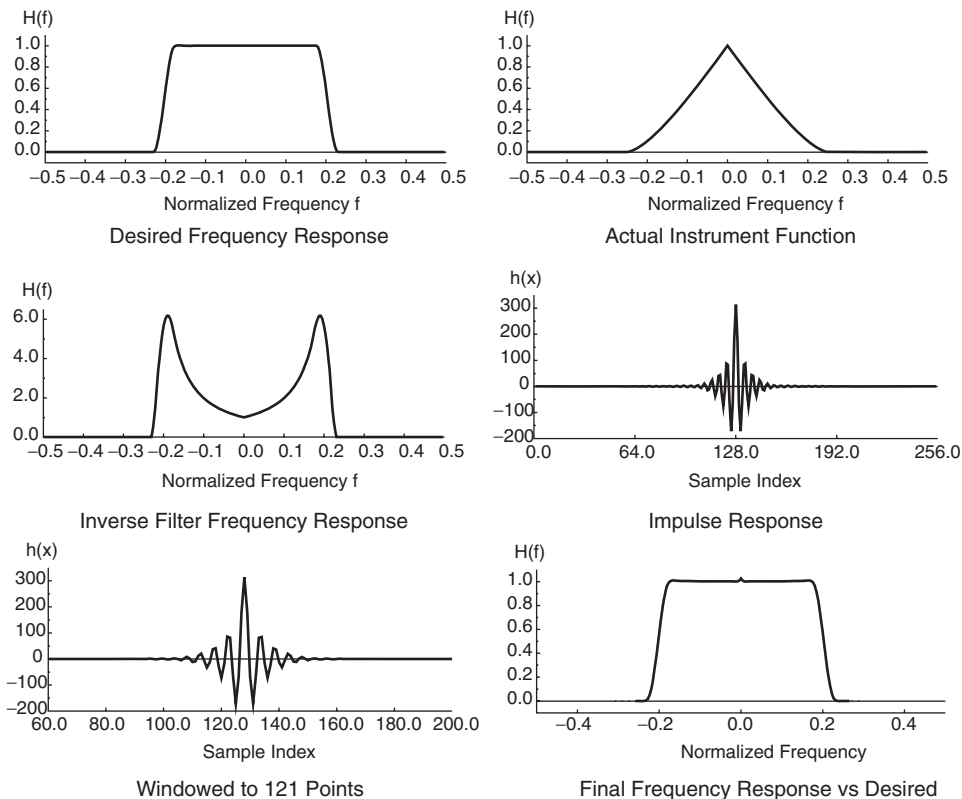


Figure 17.17. Inverse filter for a phase-sensitive confocal microscope.

to sheer beauty, at least part of the time. Used cautiously, this can make a factor of 2 difference to already-good microscope data,[†] or extend your depth of focus by 10×.

Be really careful about what you do with the edges of your instrument function if it's not fairly uniform; multiplying by the inverse of a Gaussian smoothing, for example, will lead to tears. Keep an eye on the noise gain, and watch especially for tall, narrow gain peaks—they make noise look like signal, which is a particularly obnoxious trait. Use numerical experiments with canned data and noisy fake data to make sure you've got this right.

17.7.2 Wiener Filters

Wiener filtering is a generalization of the inverse filter idea that explicitly takes account of finite SNR. It concerns how to get the best estimate of a function g that has been smeared by some (undesired) filter h and then afterwards corrupted by additive random noise r . The Wiener filter h_W is chosen to minimize the mean square discrepancy between g and $h_W \star (r + h \star g)$. Sometimes h is a δ -function, but often it isn't.

Construct a filter function

$$H_W = \frac{|G|^2}{H(|G|^2 + |R|^2)}, \quad (17.21)$$

and use that for your deconvolution. The Wiener filter provides most of the advantages of the inverse filter, with an automatic way of controlling the noise gain. (It is interesting to note that h_W , which acts on the signal amplitude, is obtained from the ratio of power spectra.) If your signal has a sharp falloff with frequency, this technique can help a fair amount and is well worth trying.

The Wiener filter is a least-squares optimum, and like most least-squares things it is nice for proving theorems but less useful in real life except to generate a reasonable first guess; controlling the deconvolution filter's shape by varying the length of the window used on the inverse filter is often a better bet unless your SNR changes a lot—and if it does, you'll need to think and test very carefully to make sure your subsequent processing doesn't go wrong when h_W changes.

17.8 RESAMPLING

Much of the time, the natural time synchronization for sampling the data is not the natural one for processing it. Correctly interpolating the data from one equally spaced mesh to another is easy—take the DFT, zero-pad it to the right length, and transform back; if you use the chirp- z transform algorithm (see e.g., Oppenheim & Shafer), N can even be a prime number (remember how the sinc functions overlap away from the original frequency samples, though). The chirp- z transform allows you to choose your frequency samples on any arc of a circular spiral in the complex f plane, so you can have any sampling resolution you like, and fine frequency samples can be had without computing the DFT over the entire unit circle at the same resolution, as you must with zero-padding.

[†]Philip C. D. Hobbs and Gordon S. Kino, Generalizing the confocal microscope with heterodyne interferometry and digital filtering. *J. Microsc.* **160** (3), 245–264 (December 1990), <http://electrooptical.net/www/confocal/confocal.pdf>

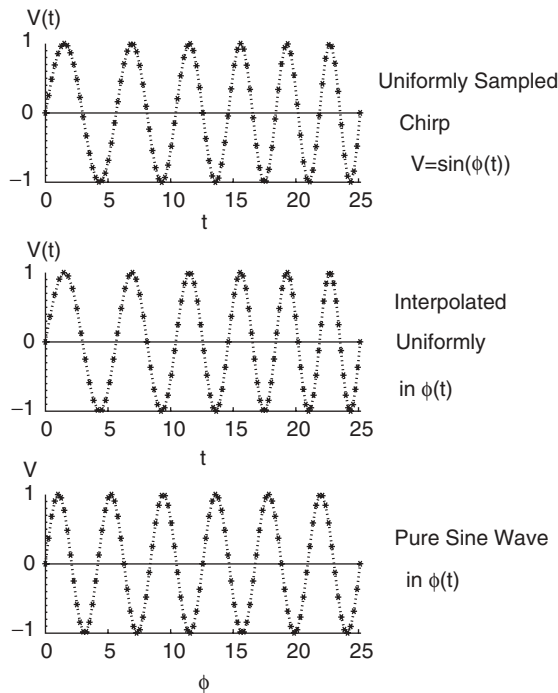


Figure 17.18. Resampling.

A seemingly similar but actually much harder problem arises when the data are not equally spaced in the right independent variable; an example is an OMA spectrometer of Example 17.2, where the output bins are equally spaced in λ and not f . This leads to the idea of resampling: using interpolation to move the sampling points from where they are to where we would like them. It is often possible to do this well, but great care must be used—resampling is an even more effective generator of reasonable-looking wrong answers than DFT power spectrum estimation. You have to multiply by the Jacobian, that is, the ratio of the bin widths before and after, and even that only levels the mean value; once you've stretched the bin widths, the noise floor will no longer be flat.

There are good and bad reasons for resampling. The three most common are: (1) it saves labor and expense in making the instrument linear in the first place (e.g., by using a rather costly f - θ lens on a scanner); (2) the signal has nice properties when resampled (e.g., a chirp could turn into a CW sinusoid as in Figure 17.18); or (3) the independent variable is inherently nonlinear in a way that must be compensated for in order to perform further processing.

Before using resampling, look hard at whether you can change the apparatus slightly to eliminate the need for it. You can vary the sampling clock rate to force scan data to be equally spaced, for example, or use a Fourier algorithm for unequally spaced data.[†]

The main danger in resampling is *truncation error* (see Section 17.4.2) due to the interpolation inherent in the procedure. Equally spaced data have a built-in translationally invariant interpolation method, that is, summing the sinc functions for each data point

[†]*Numerical Recipes in C* has a section on this.

in the set; as soon as you change the mesh, that stops being true, and your subsequent Fourier processing steps lose their mathematical foundation. Unequally spaced data have very different (and rather unintuitive) properties from equally spaced ones.

One way of seeing this is in the OMA example. The widths of the bins change under resampling, so each bin must be multiplied by its width before resampling occurs so that energy is conserved (this is the same as multiplying by the Jacobian). But this amplitude adjustment depends on where in the bin the real frequency component actually was; and that information is not present in the OMA output. This is obviously more serious when the original binning was coarse. You should thus ensure that your original data are well oversampled at the shortest frequency present in the signal, and filter really well in analog ahead of time to make doubly sure. Make sure that the same is true in the resampled data. Sometimes you can't do that, but do calculate the theoretical error bounds before proceeding.

Remember, you're making assumptions about how the data are to be reconstructed; the unequally spaced mesh may not have nice properties for reconstruction. Beware, test using many data sets, and once again, use *lots* of sample points.

17.8.1 Decimation

The previous section cautioned against ill-considered use of resampling. One safer special case is decimation, where the sample rate is divided by an integer N .[†] This is safer, since no interpolation is required. The samples of the fast data stream are first filtered digitally, so that the Nyquist criterion is satisfied at the new, slower, sampling rate, and then in every run of N samples, $N - 1$ are simply discarded.

17.9 FIXING SPACE-VARIANT INSTRUMENT FUNCTIONS

Sometimes our instrument function isn't even close to translationally invariant. Consider, for instance, a crude spectrometer using a tungsten-halogen bulb operated at different temperatures plus an unselective detector. The idea is that different bulb temperatures give different-colored light, so that, for instance, we might be able to subtract a measurement at 1500 K from one at 3400 K and get a measurement of sample absorption in the blue. However, not only is the Planck curve wide, but its width is proportional to its peak wavelength, so we can't just do a deconvolution.

In this case, we're back to needing a matrix operator; the smearing is some known matrix with no nice special properties. Okay, it will need some more CPU time, but we just multiply by the inverse, right? Well, sure, if the smearing is very gentle. Unfortunately, in the light bulb spectrometer example, the smearing function is $1\frac{1}{2}$ octaves wide[‡] and very smooth, so trying to get 512 samples over the visible spectrum from measuring at 512 different bulb temperatures is a tiny bit ill-conditioned: even for 10 bands over the visible, the condition number is on the order of 10^{15} . A 300 dB noise gain is not quite what we were looking for.

[†]The original Roman decimation was punishing a military unit for cowardice or dereliction by executing every tenth man. In digital signal processing, division by an integer works much better.

[‡]Why isn't it 2 octaves wide as in Section 2.4.1?

There is a method called *singular value decomposition* (SVD), which is a very general and well-conditioned way to solve linear systems (not necessarily square) by factoring the matrix into a row-orthogonal matrix, a diagonal matrix of singular values, and a column-orthogonal matrix. It can solve nonsingular square systems exactly, though somewhat less efficiently than *LU* decomposition. More importantly, it can also solve overdetermined, singular, or ill-conditioned systems approximately, yielding the minimum squared error solution to any rank we choose; by appropriate choice of the rank, the condition number can be controlled well.

If we are prepared to be less ambitious and abandon our idea of 512 spectral channels in favor of the three human cone pigment curves instead of 512 (see Figure 17.19), SVD will give us a pseudoinverse operator (the Moore–Penrose pseudoinverse) to solve the equation. The condition number is now more like 10^4 to 10^5 , depending on how close a fit we demand, and on how wide a range of black body temperatures we can use. These operators were calculated assuming that our detector cuts off exponentially near 800 nm.

The condition number is then equal to the ratio of the largest singular value to the smallest one; we select the rank of the pseudoinverse by discarding too-small singular values. The resulting pseudoinverse operator is the least-squares optimum for that rank. *Numerical Recipes in C* has some working code and a very good discussion of this technique, which is well worth your time if you have a variable instrument function like this.

Applying the technique is simple, though since it's an N^3 algorithm, it's expensive in CPU cycles: take your instrument function matrix, apply SVD, edit the singular values, calculate the Moore–Penrose operator, and left-multiply it by your desired bandshape. The result will be a matrix that takes your nonuniformly smeared data and produces one filter band output (for more bands, you have to left-multiply by each bandshape in succession).




Figure 17.19. Least-squares operator to transform measurements taken with black body source into human cone responses. As you can see from the size of the vertical scale, this isn't too practical.

17.10 FINITE PRECISION EFFECTS

17.10.1 Quantization

The study of the effects of quantization is somewhat subtle, but most of the subtlety concerns doing arithmetic with short word lengths rather than signal quantization per se, so we can usually ignore it if we're using lots of bits, for example, 32 bit integers, or at least single precision floating point to do our arithmetic.

Provided the signal swing is at least several ADU p-p, the quantization noise is well represented as additive Gaussian white noise of rms value $1/\sqrt{12}$ ADU. This is intuitively reasonable, since we expect the LSBs to be less and less correlated between samples as the number of bits goes up. In fact, if our digitizer is good enough, we can get the effect of extra bits by adding analog white noise to smear out the steps and then averaging for a long time to get rid of the noise again. This puts fairly stringent limits on jitter and the DNL of our digitizer. Beware!

If you are using integers, then your adds and subtracts are exact (and hence noiseless), but your multiplies involve rounding the results to the nearest multiple of 2^{-b} , which (assuming the errors are uncorrelated) adds another $2^{-b}/\sqrt{12}$ noise signal after the multiplication. All these noise sources can be carried along to the filter output, and the rms sum of their values taken; this will yield a reasonable estimate of the output noise, unless you have limit cycles and so forth to contend with (see Section 17.6).

17.10.2 Roundoff

Some signal processing algorithms are very vulnerable to roundoff, because roundoff errors during the computation can be magnified greatly by subsequent steps. Ill-conditioned matrix operations such as least-squares fitting of polynomials in x^n are one example, and maximum entropy method power spectrum estimation is another. For a given algorithm, the maximum possible amplification of roundoff error relative to the signal is called its condition number, and is a good number to know. It is equal to the ratio of the magnitudes of the largest and smallest magnitude eigenvalues of the operator.

17.10.3 Overflow

The overflow behavior of algorithms depends on whether they are implemented in integer or floating-point arithmetic. Integer overflow is not an error that will get caught; floating-point overflow will cause an exception unless you mask it off in your software. Fixed-point algorithms are more efficient and can use cheaper CPUs, but need close attention.

Integer overflow can lead to large, self-sustaining oscillations in the signal processor—that is, to disaster. Special-purpose processors sometimes use saturating adders for this reason (a saturating adder returns its maximum positive value for any positive-going overflow, and its maximum negative value for any negative-going one). This is a sensible approach that mimics the behavior of good analog networks, but adds a lot of software overhead unless your hardware supports it. (Customized arithmetic is a good application for an FPGA.)

17.11 PULLING DATA OUT OF NOISE

We're always trying to squeeze the most out of our measurements with digital post-processing, which is a great idea. Unfortunately, lots of people neglect the essential precondition, dynamic range reduction in analog (see Chapters 10 and 15). Digitizing prematurely does serious violence to your dynamic range, especially in slow measurements. Most of us have had signals that were perfectly adequate for a measurement but were invisible on a scope—that's your measurement if you digitize too soon.

17.11.1 Shannon's Theorem

It is intuitively reasonable that a given, fixed-bandwidth channel can carry more information if the SNR is high. Even a 1 Hz wide channel can be used to transmit a lot of data if the data are encoded as very small frequency shifts; for example, a frequency counter can measure the frequency of a 1 Hz signal to 8 decimal digits in 1 s, by relying on the stability of the zero crossings in high SNR situations, an equivalent data rate of $8 \log_2(10) \approx 26.6$ bits/s. Remember from Section 13.6.9 that the phase error in radians is

$$\langle \Delta\phi \rangle = \frac{1}{\sqrt{2}} \sqrt{\frac{P_N}{P_C}}, \quad (17.22)$$

so a frequency measurement of 1 cycle at a given SNR (equivalent to CNR in this case) will yield b bits' accuracy, where b is

$$b = \frac{1}{2} \log_2(\text{SNR}) + \log_2(2\pi). \quad (17.23)$$

The time taken for the measurement will depend on how a particular filter settles to b bits' accuracy, but for a given filter shape, it will scale linearly with bandwidth. This is not the optimal measurement situation, of course, since it ignores what happens in between zero crossings, but it shows that at high SNR the channel capacity C (the maximum data rate in bits/s) is proportional to $\text{BW} \cdot \log_2(\text{SNR})$. On the other hand, a signal averager needs N sweeps to improve the SNR by a factor of N , so at very low SNR we expect C to go as the SNR. Shannon's theorem is the mathematical statement that, for any network, the maximum rate of information transfer cannot exceed

$$C = \text{BW} \cdot \log_2(1 + \text{SNR}), \quad (17.24)$$

which has the right asymptotic behavior (Figure 17.20) and also takes into account the time–bandwidth product issues that govern how often we can make independent measurements of a given accuracy. (Note that this is the electrical and not the optical SNR—see Section 3.3). Shannon's theorem sets the upper limit for the amount of information our measurement contains, and so tells us how close we are to the theoretical optimum. This is very valuable, since in the design of a complicated instrument, we may make a blunder somewhere and wind up far from the optimum; Shannon's theorem helps us catch it, just as keeping an eye on the shot noise limit helps us in Chapter 18. If you don't watch the Shannon limit closely, you'll wind up attempting recreational impossibilities on one hand, or turning a silk purse back into a sow's ear on the other.




Figure 17.20. Shannon limit.

Aside: Shannon's Theorem and the Bode Limit. In Section 18.5.3, we saw that for a given RC product, there was a trade-off between bandwidth and the logarithm of the matching efficiency, which is analogous to Shannon's theorem for signaling.

17.11.2 Model Dependence

The optimal data recovery technique is one in which the signal energy is all concentrated into one sample when you're done, and the noise is spread out. An example would be a chirped signal in additive noise; resampling it so that the chirp became a pure sinusoid right at one of the sample frequencies of a DFT will do this. Fitting some sine wave to the data and extracting its amplitude, frequency, and phase is another example. Both of these are, of course, nonsense unless the model represents the signal well.

17.11.3 Correlation Techniques

In Section 13.8.10 we encountered the idea of a matched filter, whose transfer function was the complex conjugate of our signal; this undid any phase funnies and resulted in the highest received SNR. In the digital domain, there's a cheat we can use: flip the actual data around in time, and use that as the filter—that is, compute its time autocorrelation. Noise has good autocorrelation properties, but we're not exactly computing the true statistical autocorrelation, since there's no ensemble to average over, so we have to be a bit careful here. *Uncorrelated* noise sources have zero ensemble-averaged cross-correlations, but because of finite bandwidth, that is not true of the time autocorrelations of samples of noise.

With sufficiently small noise, this would be close to the matched filter and would therefore provide near-optimal response. It gets worse fairly fast at low SNR, of course, since the cross-correlation of the noise and the data is unlikely to be 0, and the cross-correlation will shrink more slowly with signal amplitude than the desired signal's autocorrelation.

On the other hand, it is a very good way of pulling out a narrowband signal in wideband noise; for example, in a backscatter interferometer where the Doppler frequency is unknown in advance, you can find tone bursts in white Gaussian noise. On the other hand, if your data set contains lots of nonrandom background signals, or your noise has structure of its own, autocorrelation methods will generate artifacts all over the place.

Cross-correlations between what you know the signal ought to look like and the actual data can be useful in estimating the delay of your signal; this is often used for range measurements in radar. If you don't have the luxury of designing your own signal shape (as radar people generally do), then when you compute the cross-correlation, there will usually be more than one peak—especially with short data runs. Reliably finding the right peak is nontrivial in general, because of the large discontinuities that occur when one peak grows up to pass its neighbor. Usually you're really interested in the envelope of the cross-correlation, which means you have to fit a function to the peaks, or else try Hilbert transforming, which is hard with short data runs. (You can design a Hilbert transformer using windows, as we did with lowpass filters.)

17.11.4 Numerical Experiments

It is not always trivial to know what the condition number of your algorithm is. We've managed to preserve our continuous-time ideas of bandwidth and SNR through all our Fourier-type processes; DFTs and digital filters have well-defined noise gains (Section 13.1), so that's simple enough. However, once you start cascading signal processing components and fancy digital algorithms (e.g., SVD), things become less clear very rapidly. For example, an analog filter with a big peak in it at a frequency that doesn't occur in the data will in practice boost the total noise much less than you might expect from its noise gain. This can be estimated analytically, and simulated afterwards, but both of these approaches can hide subtle blunders. A combination of physical and numerical experiments is especially necessary in these situations.

17.11.5 Signal Averaging

Signal averaging was encountered in Chapter 13 as a good way to narrow the measurement bandwidth without concentrating all the signal power down in the $1/f$ noise region. Now that we're a bit more sophisticated about digital processing, we can investigate what the actual passband of a signal averager looks like.

Suppose we want to average K scans of N points each, and imagine that we have all NK data points in a big array like a TV raster, as shown in Figure 17.21. At each scan point x_n , the raw data are a time series, $y_n^k(t_k)$, where $t_n^k = t_n^0 + kP$ and P is the scan period; in other words, y_n^k is the n th data value from the k th scan.

Now let's look at the n th bin, as a function of k , from a Fourier analysis perspective, with an eye to keeping it uncorrupted by strong spurious signals (e.g., harmonics of 60 Hz). The n th bin is sampled with period P , which of course is horribly undersampled; however, because the true measurement data are periodic with period P , all the desired components of bin n are aliased right down to DC (which is just what we want, since we're about to average them). The filtering operation is a rectangular-windowed average, whose frequency response is a sinc function,

$$|H(f)| = \text{sinc}(fT), \quad (17.25)$$




Figure 17.21. Partial tableau of a signal averaging scheme.

which is replicated at every harmonic of $f_{\text{scan}} = 1/P$. This is a lousy window, as we saw, so for better spur rejection we can use data windowing.

This is a different application of windowing than in spectrum analysis or FIR filter design. We want to multiply the points y_n^k by a window function W —but now $W(k)$ the scan line index, *not* $W(n)$, the index of the point within the scan line. We’re weighting the average of the scan lines. Figure 17.22 shows one of the passbands near 60 Hz of 10 averages, 100 points/scan, $\nu = 1061$ Hz, and $1/P = 10.61$ Hz (i.e., no retrace delay);

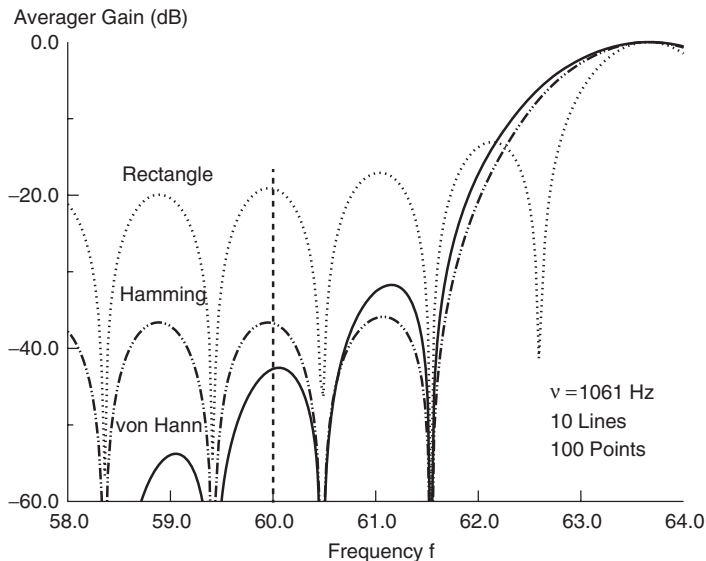


Figure 17.22. Detail of the frequency response of a signal averager, with spur rejection improved by *transverse* data windowing.

note the improvement in the sidelobes due to windowing in k , at the price of a wider main lobe, that is, poorer wideband noise rejection.[†] Another way of looking at this is that windowing in n will force the averaged data to go to 0 at the ends of the scan, and so improve the resolution of spectral estimates of the signal and residual noise and aliased interference, which isn't what we want. On the other hand, windowing in k will affect each averaged point the same way, and will improve the *rejection* of isolated spurious signals such as hum.

17.11.6 Two-Point Correlation

If you're trying to measure an aperiodic signal (e.g., noise), so that signal averaging isn't applicable, it isn't easy to separate out the instrument contribution. If the statistics of the signal are stationary, the signal magnitude can be measured using *two-point correlation* (see Section 15.3). In this trick, we measure the signal using two or more voltmeters simultaneously. The instrument noise is assumed to be uncorrelated. With M voltmeters in parallel measuring a true signal $v(t)$, the measurements that result are

$$v_j(t) = v(t) + v_{Nj}(t) + Z \cdot \sum_{j=1}^M i_{Nj}(t), \quad (17.26)$$

where $i_{Nj}(t)$ and $v_{Nj}(t)$ are the instantaneous noise current and voltage of the j th voltmeter. The current noises add because they're wired in parallel. The signal power can be obtained as the time average of $v_j(t)^2$, but that's polluted by voltage and current noise. However, if we compute the time average of $v_j \cdot v_k$, the cross terms in v_N drop out, and only the current noise is left. The current noise can be reduced to very low levels with MOSFET buffers, and so signal statistics can be measured accurately with the correlation technique even if they're below the level of the technical noise.[‡] This is a good way to measure the absolute noise of a low value resistor, for example. We can do better still by using more than two meters; the number of independent measurements is $M(M-1)/2$, so for large M our rms noise voltage goes down like $M/\sqrt{2}$, which is faster than the usual square-root dependence of averaging. (In Section 10.10 we discussed a generalization of this idea, called *closure*.)

17.12 PHASE RECOVERY TECHNIQUES

17.12.1 Unwrapping

Phase data are nearly always obtained *modulo* 2π , unless some tracking device was employed, or other information is available, for example, two phase detectors, one with frequency dividers ahead of it so that its range is not $\pm\pi/2$ but $\pm N\pi/2$. Such phase data are said to have been wrapped into the fundamental range, usually $[-\pi, \pi)$, and must be unwrapped to obtain the true phase data.[§]

[†]The Hamming window is unimpressive here because there aren't enough data samples. It's much better with 100 lines than with 10.

[‡]This obviously fails to get rid of the DC contributions, which are precisely those that survive time averaging.

[§]Some instruments use $[0, 2\pi)$, which makes life interesting when (as often) we want to servo around 0—it wraps continually, producing an average value of π , which is 180° from the right answer.

Unwrapping 1D phase data is usually easy; a phase wrap has occurred whenever two adjacent points differ by more than $3\pi/2$ radians (or π radians if you can be sure there's no zero in between, which will produce a sudden π phase shift but is not a wrap). The only times you get into big trouble are when the data are undersampled, when a tracking phase detector has inadequate bandwidth and so smears the wrap out, or when the data contains big jumps that are real, for example, a phase shifting interferometer looking at a sample with steep slopes (where the data drop out) that may be several wavelengths tall.

The 2D case is considerably harder and is usually handled by minimizing the L2 norm of the slope of the phase, via an iterative algorithm, or by a Green's function approach.[†]

Some kinds of phase data, such as synthetic-aperture radar and full-field phase shifting interferometry, often lead to phase maps with inconsistencies; these can be handled with the method of residues.[‡] This task will burn lots of cycles but will give you the best results for noisy or inconsistent data, as you might get from an undersampled data set, or one that was changing with time as the data were taken (as it often will be).

If the data are of good quality, you can do it by numerical crystallization, which is like tiling a floor: you start from somewhere in the interior, and add points to the periphery of your unwrapped data set by forcing the phase jumps at the edges to be less than $3\pi/2$ radians. You can often find inconsistencies this way, especially at boundaries where no data are available.

If you have an undersampled measurement, for example, a phase shifting interferometer using a CCD with a very small fill factor and a pixel pitch $\Delta > 0.5\lambda/NA$, you can still unwrap the phase successfully in the presence of multiples of 2π shift per pixel, if you know a priori that the wavefront phase is smooth. In that case, instead of keeping $|\partial\phi/\partial x| < 2\pi/\Delta$, as we did before, we require that

$$\left| \frac{\partial^2\phi}{\partial x^2} \right| < \frac{2\pi}{\Delta^2}. \quad (17.27)$$

17.12.2 Phase Shifting Measurements

If you have measurements of the quadrature amplitudes I and Q , you can turn them into amplitude and phase easily; most compilers provide canned functions for this. When I and Q have been corrupted by noise, spurs, and imperfect phase shifting between them, life gets more difficult; the arctangent is a pretty well-behaved function, but its error sensitivity is not 0. Many algorithms have been proposed for getting A and ϕ from N measurements spaced around the unit circle, each with different sensitivity to noise, harmonics, and errors in the phase step size, of which the best known are the Carré and Hariharan algorithms. If you're doing phase shifting interferometry, your algorithm needs some attention. Which is best depends very much on your situation, and on what the dominant source of noise and errors.

If your system has significant distortion of the sinusoidal phase dependence of the signal, the situation becomes much more difficult, as this distortion generally becomes worse as the signal level increases. This is the case, for example, when using RF mixers as phase detectors over a wide range of phases; all the harmonics and spurious mixing

[†]Gianfranco Fornaro, Giorgio Franceschetti, Riccardo Lanari, and Eugenio Sansosti, Robust phase unwrapping techniques: a comparison. *J. Opt. Soc. Am. A* **13**(12), 2355–2366 (1996)

[‡]Dennis C. Ghiglia and Mark D. Pritt, *Two Dimensional Phase Unwrapping*. Wiley, Hoboken, NJ, 1998.

products that would normally be filtered out are now right on top of your signal, since any harmonic of DC is still DC.

Be suspicious of optimal algorithms, because one of two things will be true: either the optimum is not sharp, in which case optimality doesn't buy you that much; or it is sharp, in which case minor failures to observe the assumptions on which it was derived will have large consequences. A big calibration table is probably better than any of these algorithms, if you have the time and space.[†]

17.12.3 Fienup's Algorithm

We're right at the edge of our scope, but there is another class of phase recovery algorithms that's worth mentioning. They use a priori information about an (intensity) image to recover its corresponding phase image. There are several algorithms for this, of which the best known is the Fienup phase retrieval algorithm.[‡] They all work by iteratively applying constraints in real space and \mathbf{k} -space; they're all a bit flaky, but when they work, they're amazing.

[†]There are good discussions of these algorithms in Daniel Malacara-Doblado et al., *Optical Eng.* **36**(7), 2086–2091 (1997); and Yves Surrel, *Appl. Optics* **36**(1), 271–276 (January 1, 1997).

[‡]J. R. Fienup, *Appl. Optics* **21**(15), 2758–2769 (1982).

Front Ends

Life is like a sewer. What you get out of it depends on what you put into it.

—Tom Lehrer[†]

18.1 INTRODUCTION

In Chapter 3, we dealt with optical detectors and their uses, from the output of the optical system to the detector leads. Now it's time to discuss the electronic front end—a ticklish place, between the detector leads and the signal processing system. The front end's job is faithfully turning the detector output into a buffered, filtered electronic replica. Like maintaining sewers, this is not glamorous work, but failure is very noticeable.

Bad front ends are too noisy, too slow, or both. The two are not unrelated; it's easy to make the front end fast if you are prepared to sacrifice signal-to-noise ratio, or if you have lots of light. People tend to give up much too soon—it really is possible to do fast measurements, *at* the shot noise limit, *at* low light intensities, *with* ordinary components. This is most of what this chapter is about. It has some pretty heavily technical stuff in it, so don't worry too much if it doesn't stick when you read through it; if you make the same struggle for SNR yourself, it will become clearer very quickly.

A basic front end is just a transimpedance amplifier (current-to-voltage converter). More advanced ones perform linear and nonlinear combinations of the signals from more than one detector, as in differential measurements, and these operations must be very accurate. Such advanced front ends allow you to make high stability measurements in bright field, make simultaneous amplitude and phase measurements, or reject laser noise with extremely high efficiency. A good one can spectacularly improve your measurement. Throughout this chapter, we'll be tossing around ultraquiet voltage and current sources, and most of the time we'll be behaving as though light source noise doesn't exist. Don't be concerned about where we're getting these magic parts—quiet voltages and currents are constructed in Section 14.6.5 and Example 14.1, and source intensity noise is largely vanquished in Section 10.8.6.

[†]Preamble to “We Will All Go Together When We Go,” in *An Evening Wasted with Tom Lehrer*, private label recording 1959, reissued by Reprise Records, 1966.

TABLE 18.1. Major Sources of Noise in Front Ends

Source	Type	Formula	Dominates When
Photocurrent	Shot noise	$i_{N\text{shot}} = (2eI_d)^{1/2}$	Bright light, large R_L
Load resistor	Johnson noise	$i_{N\text{th}} = (4kT/R)^{1/2}$	Dim light, small R
Amplifier	Input current noise	i_N as specified	Ideally never
Amplifier	Input voltage noise	v_N as specified	Dim light, large RC , or a fast, noisy amp
Power supply	Ripple, regulator noise		Only by blunders

18.1.1 Noise Sources

Since good front end design is largely a matter of balancing noise sources, it would be worth beginning by reviewing the discussion of noise sources and calculations in Section 13.6.2. Table 18.1 summarizes the major sources of electronic noise encountered in front end design.

18.1.2 Sanity Checking

Since the first edition of this book was published, the author has been receiving a certain number of e-mails from people with detection problems, which are welcome. One common feature that has emerged is that specifications for optical instruments are often set by people who are, *ahem*, not expert in optical measurements.[†] One of the most common is to insist on wide bandwidth with high SNR at low light levels, for example, 50 MHz at 20 pA of photocurrent, which cannot be done for reasons having nothing to do with circuit design. Accordingly, here are a few representative rules of thumb for frequently asked questions:

1. If you have N photons/s, your SNR will drop to 0 dB at a bandwidth $B = N/2$ Hz. This is an inescapable limit based on counting statistics. Your maximum achievable SNR is $N/(2B)$, so since 20 pA is 1.24×10^8 electrons/s, counting statistics limit the SNR to 1.24 (0.9 dB) in 50 MHz.[‡]
2. Using a few high precision parts doesn't get you a precise measurement. You can measure a photocurrent very accurately, but accurate measurements of light intensity are very hard, and 24 bit A/D converters don't help. It isn't that photodiodes aren't good transducers—there are none better—but that the problem isn't well posed. The mapping of what you actually care about onto the photocurrent is just about always imprecise at the level of a percent or two, due, for example, to etalon fringes, calibration drift, background light, glints, and so on. There are honorable exceptions, but not that many, and no measurement whatsoever can give a clear answer to a fuzzy question.
3. Physicists (such as the author) are often prone to oversimplifying circuit requirements. It is not enough to aim at being “shot noise limited” or “ R_f limited” and

[†]I'd have been less complimentary, but then you couldn't show them this section when the problem comes up. Miller and Friedman's book, *Photonics Rules of Thumb*, was written partly to help cure this problem.

[‡]Time–bandwidth product issues like this show up in digital signal processing too—see Section 17.5.

stop there. Filters do not cut off everything outside their passbands. Not everything has a one-pole rolloff. Shot noise limited SNR can be improved by getting more photocurrent. There's no substitute for calculating the SNR and frequency response.

18.2 PHOTODIODE FRONT ENDS

This section is really an extended example of how to design a front end amplifier for a visible or near IR photodiode, and how to get a factor of 4000 improvement in bandwidth over the naive approach without sacrificing SNR. All the signal-to-noise comparisons we'll be making will be the DC signal power to the rms noise, which is really a carrier-to-noise ratio (CNR) since what we think of as the signal will usually be much smaller than the DC. The SNR is what we care about, so we'll use that for rhetorical purposes.

18.2.1 The Simplest Front End: A Resistor

Say we need with a detector subsystem whose 3 dB bandwidth is 1 MHz, a photocurrent of $2 \mu\text{A}$ from a silicon photodiode whose capacitance is 600 pF at zero bias. Given a detector whose output is a current, the easiest way to form a voltage from it is to shove it into a resistor, say, $1 \text{ M}\Omega$, as shown in Figure 18.1. While this circuit generates an output voltage $V_o = I_d R$ with admirable linearity (at least until we forward-bias the PD too far), problems arise as soon as we ask about AC performance. Since the full signal swing appears across the detector capacitance C_d , the output rolls off starting at

$$f_{RC} = \frac{1}{2\pi R_L C_d}, \quad (18.1)$$

which is 265 Hz at zero bias. This is a mere factor of 3800 slower than our 1 MHz design point. As we saw in Section 3.4.5, most visible and NIR detectors can be operated at reverse bias, which will reduce C_d (by as much as 7–10 times) while increasing the leakage current slightly. This is nearly always an excellent trade, contrary to what you'll often read elsewhere. This diode's data sheet says that its leakage current is about 0.5 nA at room temperature, for a 12 V reverse bias, and that this bias will reduce C_d by a factor of 6, to 100 pF. That gets us to 1600 Hz, still not blazing fast. (We get to keep




Figure 18.1. The world's simplest front end: a load resistor.

that factor of 6 all through the design, however, so eventually it'll take us from 170 kHz to 1 MHz, which is a bit more impressive sounding.) Since the noise of the bias current will be more than 30 dB below the photocurrent shot noise, this seems like a good thing to do: we get a factor of 6 in bandwidth for a shot noise increase of 0.004 dB—which is too small even to measure. (See Section 3.5.2.)

The signal voltage V_o goes as

$$V_o(f) = \frac{i_d(f)R_L}{1 + j2\pi R_L C_d f}, \quad (18.2)$$

as shown in Figure 18.2.

Somewhat surprisingly, though, the signal-to-noise ratio does not deteriorate at all. The resistor's i_N and the shot noise current are both treated exactly as the signal is. The reason for this is apparent from Figure 18.1: the signal and noise sources are all connected in parallel.[†] Thus they all roll off together with increasing frequency, which makes their ratios frequency independent, as Figure 18.2 shows. Any deterioration of the signal-to-noise ratio of the measurement is due to the subsequent amplifier stages. It's not the poor amplifier's fault, though—a source whose impedance changes by a factor of 600 over the band of interest is not the best of neighbors.

As is usual when it's circuit constants and not laws of nature which are in the way, with a bit of ingenuity we can find circuit hacks to get around the rolloff without messing up the SNR too badly.




Figure 18.2. Frequency response and narrowband CNR of the photodiode/load resistor combination of Figure 18.1, with $R_L = 1 \text{ M}\Omega$ and $C_d = 100 \text{ pF}$.

[†]Why is it OK to move the bottom of C_d and i_d to ground?

TABLE 18.2. Noise Degradation Due to the Johnson Noise of a 300 K Resistor

$I_d R(\text{V})$	$i_{N\text{th}}/i_{N\text{shot}}$	ΔSNR	$I_d R(\text{V})$	$i_{N\text{th}}/i_{N\text{shot}}$	$\Delta\text{SNR} (\text{dB})$
5.1	0.1	-0.04	0.14	0.6	-1.3
1.3	0.2	-0.17	0.10	0.7	-1.7
0.57	0.3	-0.4	0.080	0.8	-2.1
0.32	0.4	-0.6	0.063	0.9	-2.6
0.20	0.5	-1.0	0.051	1.0	-3.0

18.2.2 Reducing the Load Resistance

After reverse biasing, the first thing everyone thinks of is reducing the load resistance, because that reduces the RC product and speeds things up. This *does* reduce the signal-to-noise ratio, because unlike the previous case, the resistor's noise current goes up as its value is reduced. There is nothing much lost by reducing the resistance while shot noise still dominates—when the shot noise current is larger than the Johnson noise current. The shot noise ceases to dominate when the two become equal, that is, when the DC voltage drop across R_L is $2kT/e$ (51 mV at room temperature), and we enter the *Johnson noise limit*. Good instrument designers grind their teeth if they're stuck in the Johnson noise regime, since the data from an expensive optical system are being seriously damaged by circuit limitations. In particular, running a photodiode into a room temperature $50\ \Omega$ load is always a mistake unless the light level is very high (milliwatts in the visible). There are lots of things you can do to get decent bandwidth, so resist the $50\ \Omega$ temptation. (As we'll see later, it's also possible to achieve an effective temperature of the load resistance as low as 35 K at room temperature, so all is not lost.)

Remember too that the SNR versus thermal noise curve doesn't have a sharp corner. Table 18.2 shows the SNR degradation due to load resistor Johnson noise as a function of the DC voltage across R_L , which is a convenient way to remember it. If $I_d R = 0.2$ V, then $i_{N\text{th}} = 0.5i_{N\text{shot}}$ (6 dB down), and we've already lost 1 dB in SNR.

Making R_L too big wastes both bandwidth and dynamic range, so it is usually best to choose a value that drops 100 mV to 1 V. (Later we'll do this with transimpedance amps too.) For the present circuit, we will assume that a 1 dB loss is acceptable, so we'll shoot for a voltage drop of 0.2 V. With our $2\ \mu\text{A}$ photocurrent, we'll need a $100\ \text{k}\Omega$ resistor, which will improve the RC bandwidth to 16 kHz, a mere factor of 60 away from our goal.

18.3 KEY IDEA: REDUCE THE SWING ACROSS C_d

Once we have carefully chosen R_L and reverse-biased the photodiode, the circuit will probably still be too slow, as we've seen. It's time to change the circuit topology and see if that works well enough. We may observe that the source of the poor bandwidth of the load resistor approach is that the full signal swing appears across C_d . If we make both ends of the photodiode work at constant voltage, then there will be no swing across C_d , and hence no capacitive current (see Section 15.3). Making the swing small requires making the load impedance small. How can we do that without degrading the noise?

18.4 TRANSIMPEDANCE AMPLIFIERS

One way to do it is to make the detector work into a virtual ground, as shown in Figure 18.3. Although the inverting input of A_1 draws no current, feedback forces the voltage there to be close to zero at all times. The way this works is that A_1 senses the voltage across C_d and wiggles the other end of R_f to zero it out. Provided A_1 has high open-loop gain A_{VOL} , the swing across C_d is greatly reduced, and the bandwidth greatly improved. The amplifier input adds a significant amount (2–20 pF) of its own capacitance C_{in} , which must be added to C_d . Because this circuit is so important in applications, it's worth spending a little time analyzing its bandwidth and noise.

The voltage gain of A_1 is not infinite, so that the swing is not exactly zero; to produce an output voltage V_o , A_1 requires an input voltage $V_i = V_o/A_{VOL}$. A_{VOL} rolls off at high frequency, which limits the bandwidth improvement. Prepackaged op amps have their open-loop frequency responses carefully tailored to make them easy to use, which in practice means that they roll off like $1/f$ (6 dB per octave), with a nearly constant 90° phase shift from a low frequency all the way to their unity gain crossover at f_T . The uppermost curve of Figure 18.4 shows the response of an LF356 (105 dB DC gain, 4 MHz f_T), which is of this character. The advantage of this is that any closed-loop gain will result in a stable and well-behaved circuit that settles quickly. This approach is called *dominant pole* compensation; its drawback is wasted bandwidth at high closed-loop gain, which does not greatly concern us here. Mathematically, A_{VOL} is approximately

$$A_{VOL}(f) = \frac{A_{DC}}{(1 + jf/f_{dom})(1 + jf/f_2)} \quad (18.3)$$

The exact values of the DC gain A_{DC} and the dominant pole frequency f_{dom} are not well controlled from unit to unit. Their product, known as the gain–bandwidth product (GBW), is approximately equal to the unity gain crossover frequency f_T and is a well-controlled parameter. The other term in the denominator, which is a pole at frequency f_2 , represents the effects of limited bandwidth in other stages of the amplifier. In amplifiers intended for use at unity gain, f_2 is always higher than f_T , but not by much—a factor of 1.2 to 4, thus contributing an additional phase shift at f_T from 40° down to 15° .




Figure 18.3. Op amp transimpedance amplifier.




Figure 18.4. Frequency responses of parts of the transimpedance amplifier loop. A_1 is an LF356 op amp, $R_f = 100\text{ k}\Omega$, $C_f = 6.3\text{ pF}$, $C_d = 100\text{ pF}$. See Section 18.4.1 for the choice of C_f .

When we close the feedback loop by connecting some network between the output and input of the amplifier, we can predict the frequency response of the resulting circuit from the open-loop responses of the amplifier and the feedback network.

The feedback network here is the series combination of R_f and C_d , whose voltage gain is

$$H_{fb}(f) = \frac{1}{1 + j2\pi f R_f C_d}. \tag{18.4}$$

Roughly speaking, the closed-loop gain of an amplifier starts to roll off at about the point where the product of the open-loop gains of the amplifier and feedback network falls to 0 dB. Extra phase shifts due to the other poles in the circuit can modify this somewhat, as we'll see below, but it's within a factor of 2.

Far down on their slopes, the responses of the feedback network and the amplifier are approximately $-jfRC/f$ and $-jfT/f$, respectively. Their product is approximately $-fRCfT/f^2$, and the loop bandwidth of the resulting transimpedance amplifier is therefore approximately

$$f_{CL} \approx \sqrt{fRCfT}, \tag{18.5}$$

which for the LF356/100 kΩ/100 pF combination is $(16\text{ kHz} \cdot 4\text{ MHz})^{1/2}$, or about 250 kHz. The transimpedance rolls off somewhat earlier than this, since it depends on the magnitude of the impedances of the feedback elements, and not merely on their ratio. Calculating the transimpedance bandwidth is a straightforward exercise—you put a current into the summing junction and calculate how much goes through R_f and how much through C_d . Without going hip deep into algebra, you lose a factor of between $\sqrt{2}$ and 2 in bandwidth, depending on the details of the frequency compensation scheme, so for a rule of thumb we'll say that

$$f_{-3\text{ dB}} \approx \frac{\sqrt{fRCfT}}{2}. \tag{18.6}$$

We'll actually get around 130 kHz transimpedance bandwidth from the LF356 circuit, a factor of more than 8 improvement.

This is still fairly far from 1 MHz, but getting a lot closer. We need about 8 times more bandwidth, so if we choose an amp with a bandwidth 60 or so times higher (i.e., 250 MHz), then we ought to get there. Right? Well, sort of. There are two things we've left out of this picture. One is noise, and the other is frequency compensation. Frequency compensation is easier, so let's knock that off first.

18.4.1 Frequency Compensation

The equation for the closed-loop noninverting gain of a feedback amplifier is

$$A_{VCL}(f) = \frac{A_{VOL}(f)}{1 + A_{VOL}(f)H_{fb}(f)}, \quad (18.7)$$

where H_{fb} is the gain of the feedback network (usually a voltage divider). For frequencies where the *loop gain* $A_{VL} = H_{fb}A_{VOL} \gg 1$, this simplifies to $1/H_{fb}$. Looking at the denominator, clearly what happens when $A_{VL} \approx 1$ will have a great effect on the closed-loop behavior. If the loop gain has a phase of -180° when it crosses unity magnitude, the denominator will go to zero and the closed-loop gain will be infinite, which means that the circuit will oscillate fiercely near there (this is a sufficient but not necessary condition for oscillation). On the other hand, if the phase is -90° , then there will be a well-behaved *RC*-type corner there. The difference between the actual open-loop phase and -180° is called the *phase margin*. In practice, as long as the worst-case phase margin is greater than 45° or so, the closed-loop response will not exhibit undesirable peaking, and the time-domain step response will not overshoot too much.

From (18.7), we can show that an amplifier whose phase margin is 90° , that is, a single *RC* rolloff, has a closed-loop 3 dB corner frequency f_c at exactly the open-loop unity gain crossover, whereas one with 45° margin has its corner at an open-loop gain of only 0.52.

In order to achieve a 45° phase margin, we need to stop the rolloff of the feedback network at a frequency about equal to the closed-loop corner. We can do this by putting a capacitor C_f across R_f , where

$$C_f = \frac{1}{2\pi R_f \sqrt{f_T f_{RC}}}, \quad (18.8)$$

which gives a phase margin of between 45° and 60° , depending on how fast the amplifier is. An alternative is to put a small resistor R_s in series with the photodiode, where R_s is

$$R_s = \frac{1}{2\pi C_d \sqrt{f_T f_{RC}}}. \quad (18.9)$$

In complex-variables language, these additions put a zero into the transfer function (see Section 15.4.3). The exact value of R_s or C_f that gives the optimal trade-off of peaking versus bandwidth for your application depends on what you are most interested in, so take these values as starting points. Beware of device-to-device variations in C_d and GBW if you're making more than one or two copies. If you crank C_f down to the

absolute lowest tolerable value in your lab prototype, Murphy’s law dictates that the next 100 photodiodes will be at the upper spec limit for capacitance, and all 100 circuits will oscillate merrily. It is axiomatic that all prototypes contain at least one perfect component, which works beautifully but is totally unrepresentative of normal production units.

Although the approximation (18.6) ($f_{-3\text{ dB}} \approx f_{\text{CL}}/2$) is good enough for early design purposes, it is worthwhile to carefully plot the frequency response of the transimpedance, because it hasn’t the same shape as the closed-loop gain. With the approximate expression (18.3) for the op amp’s gain, the transimpedance is given by

$$Z_m = \frac{A_{\text{VOL}}Z_f}{1 + A_{\text{VOL}} + j2\pi fC_dZ_f}, \tag{18.10}$$

where Z_f is the complex impedance of the parallel combination of R_f and C_f .

Figure 18.5 shows the performance of the transimpedance amplifier, with frequency compensation by $C_f = 6.3\text{ pF}$, as calculated from (18.8). The transimpedance bandwidth is only about half f_0 , and it rolls off very steeply (approximately 18 dB/octave, equivalent to 3 poles). Also shown are the open-loop gain and the closed-loop noninverting gain, which we will encounter in the next section.

18.4.2 Noise in the Transimpedance Amp

Confusion reigns supreme in discussions of noise in transimpedance amps. Let’s try to boil it down to something reasonably memorable. Figure 18.6 shows a simple but adequate noise model of a transimpedance amp plus a photodiode. It is visually obvious that all the current sources are treated identically: I_d , i_{Nshot} , i_{Nth} , and I_{Namp} appear in parallel. The Johnson noise i_{Nth} of R_f really appears across R_f , of course, but because the impedance of the op amp output is very low, the other end of the noise current source is at ground for noise purposes. The signal current thus appears in parallel with the current noise sources, just as in the simple load resistor case, so the rolloff in the frequency response will once again not degrade the signal to current noise ratio.




Figure 18.5. Performance of the transimpedance amplifier.




Figure 18.6. Simplified noise model of the transimpedance amplifier. All noise sources except e_{Namp} are treated exactly as the photocurrent and shot noise, so that only e_{Namp} changes the SNR.

The only noise source that is treated differently is the amplifier’s voltage noise, e_{Namp} . Because the amplifier amplifies only differential signals (i.e., those in which its inputs move in opposite directions), the model noise source can be put in either input lead. Here we put it in the noninverting lead, which simplifies the analysis: clearly, e_{Namp} will be multiplied by the noninverting gain of the amplifier, A_{VOL} —which is therefore the noise gain of the stage. (See Section 13.1.) Taking Z_f as the complex impedance of the feedback element (R_f in parallel with C_f),

$$A_{VOL} = \frac{A_{VOL}}{1 + \frac{A_{VOL}}{1 + j\omega C_d Z_f}}. \tag{18.11}$$

For frequencies well within the loop bandwidth, the resulting equivalent noise current is approximately

$$i_N \approx (2\pi f C_D) e_{Namp}. \tag{18.12}$$

This gain begins to rise at the RC corner frequency of C_d and R_f , just where the signal rolloff would have begun if we were using a simple load resistor approach; in fact, the SNR as a function of frequency is identical to that of the same amplifier connected as a buffer following a photodiode and load resistor, which is reassuringly reasonable. All we’ve done is to tailor the frequency response by using feedback to jiggle the far end of R_f ; this shouldn’t get us something for nothing. The addition of C_f or R_s doesn’t fundamentally change this, but it causes the input referred noise to level off at the frequency of the feedback zero.

If the op amp’s voltage noise is very low, or if we are not trying to get a huge bandwidth improvement through the $(f_T \cdot f_{RC})^{1/2}$ mechanism, this rising noise contribution will not limit us. If we are relying heavily on this mechanism, though, the noise may increase catastrophically: it will begin to dominate all other noise sources at approximately f_3 , where

$$f_3 = \frac{1}{2\pi e_{Namp} C_d} \sqrt{2eI_d + i_{Namp}^2 + \frac{4kT}{R_f}}. \tag{18.13}$$

We can see this nefarious gotcha in action in Figure 18.7, which is a plot of the noise power spectral density of our LF356 circuit. The voltage noise is unimportant at low frequency but rises to dominate the entire noise budget. The log-log plot is a bit deceiving; plotting the noise power versus frequency on linear scales, as in Figure 18.7, gives a more visceral feel for the problem. It only gets worse when we try to go faster. One reason for the confusion is that most of us have a persistent idea that noise spectra are flat or nearly so—which is often true, but not here.




Figure 18.7. Noise performance of the transimpedance amplifier of Figure 18.3, showing the dominance of dark noise (i.e., additive circuit noise) at high frequency. At right, the same data plotted on linear scales. This shows the true character of the $e_{N_{amp}}$ problem.

18.4.3 Choosing the Right Op Amp

In order that the op amp not dominate the noise, choose it by the following rules (worst-case specifications apply):

1. $i_{N\text{amp}} < 0.5i_{N\text{th}}$. We've chosen R_f so as to lose no more than 1 dB to Johnson noise, so don't mess it up by choosing an amp whose current noise is as big or bigger than R_f 's.
2. $e_{N\text{amp}} < 0.5e_{N\text{th}}$. Similarly, we don't want the amplifier's voltage noise to dominate under any circumstances.
3. $e_{N\text{amp}} < 0.5i_{N\text{th}}/(2\pi f_{-3\text{dB}}(C_d + C_{\text{in}}))$. This ensures that the rising noise current due to $e_{N\text{amp}}$ doesn't begin to dominate anywhere in the band we care about. (C_{in} is included explicitly here as a memory jogger—it always has to be added in.)
4. $f_T > 2f_{-3\text{dB}}^2/f_{\text{RC}}$. The amplifier has to be fast enough to raise the bandwidth sufficiently.
5. $f_T < 10f_{-3\text{dB}}^2/f_{\text{RC}}$. Going too fast is asking for trouble. The size of the noise peak will be so large that extensive filtering will be needed to get rid of it, and the circuit may even oscillate.
6. If finding an amp that satisfies Rules 1–4 runs into money, either spend it, or use a circuit hack to get round it. Don't economize here.
7. It is not always necessary to use a unity-gain stable amplifier because of the $C_d R_f$ gain peak, but watch the frequency compensation extra carefully if you don't.

Using a decent part, which has guaranteed specifications for noise and gain bandwidth, usually pays for itself many times over by the relaxation this permits in the specs of the optical system. Use worst-case design here, and leave a safety margin, especially on $e_{N\text{amp}}$. The log–log plots are deceiving; if $e_{N\text{amp}}$ dominates only near the high end, that is as bad as dominating over the whole band, because there's a lot more high end than there is low end.

Table 18.3 lists some good op amps for use in photodiode front ends. The exact circumstances in which each is superior are somewhat complicated: try the low voltage noise ones with bright light, and the low current noise ones at low light. Remember that op amp input capacitance has the same effect as photodiode capacitance, so that the musclebound FET units with the high C_{in} are not as universally useful as their low noise specs would suggest. This remains true even if we put them on steroids, as the following example shows.

Example 18.1: External JFET Differential Pair. It's often possible to improve the noise performance of op amp TIAs by adding a discrete JFET front end. By adding a pair of 2SK369s with a voltage gain of 20 or so, running at $I_D = 10$ mA ($V_{GS} = -0.1$ V), we get 1 Hz $v_N = 0.7$ nV, which is 1.0 nV for the pair combined. Keeping V_{DS} down to 6 V or so keeps the gate current below 5 pA, so i_N should be in the femtoamps. Most of the input current noise comes in via the drain-gate capacitance, so we'll assume a cascode stage. Getting a stage gain of 20 requires $R_L = 20/g_m = 400$ Ω .

The resulting amplifier has about 90 pF of inverting input capacitance, unfortunately, due to the very large die size of the JFETs. With the BJT cascode, the 50 pF feedback

TABLE 18.3. Suggested Op Amps for Front Ends^a

Device	Manu- facturer	f_T (MHz)	$v_N@10\text{ kHz}$ (nV/Hz ^{1/2})	$i_N@10\text{ kHz}$ (pA/Hz ^{1/2})	C_{in} (pF)	$C_{in} \cdot v_n$ (typ)	Remarks
FET							
LF356	NS	4	12t	0.01t	3t	36	Cheap, good
OPA627	TI	20	6	0.0025	7t	32	OPA637 decomp version
OPA129	TI	1t	15t	0.0001	1.5t	22	Good in low light, expensive
AD795	AD	1.6t	11	0.0007t	2t	18	Good at low I_d
OPA657	TI	1600	4.8t	0.0013t	2.6t	12	$\pm 5\text{ V}$, $A_V > 10$ but fast and quiet
OPA656	TI	230	7t	0.0013t	1.7t	12	Unity-gain compensated 657
Bipolar							
OP-27	AD	8	3.8	0.6	2t	7	Good for $I_d > 5\ \mu\text{A}$
OP-37	AD	45	3.8	0.6	2t	7	$A_V > 5$
OPA687	TI	3800	1.1	3.3	1.8t	2	$A_V > 40$, $35\ \mu\text{A}$ I_b , 5 V
LT1028	LT	50	1.1	1.6	5t	4.8t	Good for low R_p diodes (e.g., InAs)
AD829	AD	600t	2	1.5	3.2t	5.5	
AD8397	AD	69t	4.5t	1.5t	1.4t	6.3	3600 V/ μs slewing
LM6311	NS	80t	2.3t	3.5t	2.5t	8.8t	Poor data sheet
LM7332	NS	21t	15.5t	1t			Drives unlimited C_L
LMH6624	NS	1500t	0.9t	2.3t	2.2t	2t	$\pm 6\text{V}$, poor datasheet
LME49710	NS	55t	4.7	1.6t	3t	7.5	LM4562 dual, LME49740 quad
OP-470	AD	6t	5	0.4t	2t	12	Quad; OP270 dual
AD8397	AD	33t	4.5t	1.5t	1.4t	6.3	$\pm 12\text{V}$
ADA4898-1	AD	33t	0.9t	2.4t	1.5t	1.4	$\pm 12\text{V}$ —amazing part
MC33078	ST	10	4.5t	0.5	12t	54	Dual, largish C_{in}

^aMostly $\pm 15\text{ V}$ devices due to their much greater dynamic range. Unless noted, devices can use $\pm 15\text{ V}$ supplies and are unity gain stable.

capacitance times the noise at its emitter contributes a 1 Hz current noise

$$i_{NC}(f) = 2\pi f C_{dg} kT \sqrt{\frac{2}{eI_D}}, \quad (18.14)$$

which is 0.05 pA at 1 MHz, rising 20 dB/decade. With a total summing-junction capacitance of 190 pF including C_d , Rule 3 shows that even a nice FET like this one is 16 dB too noisy for our application. (A pair of BF862s would be a little closer, but still at least 12 dB too noisy.) The biggest benefit of this approach is the very low C_{in} you can get by cascoding, which allows the use of much bigger feedback resistors, as we'll see below.

18.4.4 No Such Amp Exists: Cascode Transimpedance Amplifiers

In our design, if we are aiming at getting to 1 MHz transimpedance bandwidth, Rules 1–5 lead to an amplifier with the following characteristics: $I_{Namp} < 0.20 \text{ pA/Hz}^{1/2}$; $e_{Namp} < 0.32 \text{ nV/Hz}^{1/2}$; $250\text{MHz} < f_T < 1250 \text{ MHz}$. No such amplifier exists, not even with an external input stage. Now what? Another circuit hack, of course.

Recall that our reason for using the transimpedance amplifier was to get rid of the voltage swing across C_d . We can do this another way, by using a common-base transistor amplifier, as shown in Figure 18.8 (Resistor R_E will come in later—ignore it for now.) The transistor faithfully transmits its emitter current to its collector, while keeping its emitter at a roughly constant voltage. This idea is used in common-emitter transistor amplifier design to eliminate severe bandwidth limitations due to collector–base feedback capacitance (the Miller effect). The resulting amplifier configuration resembles a two-layer cake and is called a *cascode*. The cascode idea works here as well. In the Ebers–Moll transistor model, the small signal resistance r_E of the transistor’s emitter circuit is

$$r_E = \frac{kT}{eI_C}, \tag{18.15}$$

where kT/e is 25 mV at room temperature (r_E is intrinsic to the transistor and should not be confused with the real metal film resistor R_E). Thus our $2 \mu\text{A}$ photocurrent sees a resistance of 12.5 k Ω , so that the RC bandwidth increases by a factor of 8 immediately, to about 130 kHz. What is more, the collector circuit has a shunt capacitance set only by the output capacitance C_{ob} of the transistor and C_{in} of the op amp, which can be chosen to be much less than C_d , so that we can raise R_f if we choose without losing bandwidth or suffering from serious e_{Namp} multiplication.

Aside: r_E Auto-scaling. An interesting feature of the nonlinearity of r_E is that it automatically adjusts to an increase in photocurrent by reducing the RC product. In Section 18.2.2, we chose a value of R_f proportional to $1/I_d$; that’s just what Q_1 does, while maintaining the $8\times$ bandwidth improvement.

On the other hand, we can no longer improve the bandwidth by simply using a faster amplifier, and besides, the bandwidth of the circuit depends on I_d . Sure, the limitations of the transimpedance stage are less of a worry, but if we can’t get the $(f_{RC}f_T)^{1/2}$ bandwidth improvement, have we really gained anything? The answer is yes. First, remember that the RC rolloff moves 8 times higher in frequency, which by itself often makes the bandwidth adequate. Second, we are not powerless to improve the bandwidth further: in

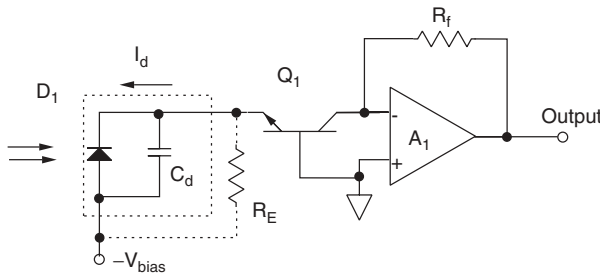


Figure 18.8. Using a common-base amplifier greatly reduces the effects of C_d and significantly improves the SNR as well. External biasing via R_E provides even more improvement.

fact, there are two ways to fix these minor warts while gaining even more bandwidth. First of all, though, let's look at the SNR of the cascode to see what this bandwidth improvement costs us.

18.4.5 Noise in the Cascode

In the simple load resistor case, the signal, shot noise, and Johnson noise contributions rolled off together, resulting in a constant SNR. Here we're not quite that lucky, because there is an additional noise contribution from Q_1 , which rises with frequency; it is much more benign than the e_{Namp} problem with transimpedance amplifiers, however.

Any transistor has some noise of its own. A simple noise model of a BJT is shown in Figure 18.9, which neglects only the Johnson noise of the base resistance $r_{B'}$ (normally only a problem when $I_C \gtrsim 1$ mA). The active device in the model has infinite transconductance (i.e., emitter impedance of 0Ω) and no noise.

Noise current i_{nB} is the shot noise of the base current $I_B = I_C/\beta$, which is inescapable, while i_{nC} is the shot noise of the collector current, which shows up in parallel with the small signal emitter resistance r_E (we ignore the difference between I_C and I_E for now, and just talk about the collector current I_C). If the emitter is grounded, all of i_{Nbias} goes from ground into the emitter, and so contributes full shot noise to the collector current. On the other hand, if the emitter is biased by a current source (e.g., a resistor many times bigger than r_E), all the noise current has to go through r_E , and none at all winds up in the collector current (the real emitter lead does jiggle up and down slightly, though). It may be more comforting to talk about the Thévenin model, where the shot noise is converted to an emitter-base voltage by dividing by the transconductance, so that the voltage noise is

$$e_{Nshot} = \frac{\sqrt{2eI_C}}{g_m} = kT \sqrt{\frac{2}{eI_C}}, \quad (18.16)$$

to which must be added the Johnson noise of the extrinsic base resistance $R_{B'}$, usually 40–100 Ω . (They are added in RMS, of course.†)




Figure 18.9. Simplified noise model of a bipolar junction transistor (BJT).

†Don't try to calculate this noise contribution by applying the Johnson noise formula to r_E —it's 3 dB lower than that, and the physics is completely different.

Whichever way you prefer, if the photodiode impedance is infinite, the transistor does not contribute noise to the collector current. Referring back to Figure 18.8, we see that our diode really has capacitance, so the finite impedance of C_d makes i_{NC} split between C_d and r_E by the ratio of their admittances (remember we're dealing with RMS averages of noise here, so only the magnitude matters).

A couple of lines of algebra then give the Q_1 contribution to the noise:

$$i_{NQ_1} = V_{NQ_1} \frac{\omega C_d}{\sqrt{1 + (\omega C_d r_E)^2}} \approx \sqrt{2eI_C r_E^2} \frac{\omega C_d}{\sqrt{1 + (\omega C_d r_E)^2}} \quad (18.17)$$

(you can instantly see this if you use the Thévenin equivalent circuit for i_{NC} and r_E). In an unbiased cascode, where I_C is all from photocurrent, this contribution exactly cancels the rolloff of the photocurrent shot noise, so that the collector current of Q_1 has full shot noise at all frequencies. Thus the 1 Hz SNR rolls off exactly as the signal does and is 3 dB down at the signal corner frequency f_c . This is at least easy to remember.

On the other hand, if the applied emitter current I_{Eq} has only δ times full shot noise power, as it will in a minute, the i_{NC} contribution will start to dominate the bias current noise at a somewhat lower frequency

$$f_{SNR} = f_c \sqrt{\delta}, \quad (18.18)$$

which turns out to be a serious limitation.

Aside: Q_1 Voltage Noise. If the $R_{B'}$ noise of the transistor is important, its voltage noise has to be added in to the numerator surd of (18.17), which becomes

$$i_{NQ_1} = \sqrt{2eI_C r_E^2 + 4kT R_{B'}} \frac{\omega C_d}{\sqrt{1 + (\omega C_d r_E)^2}}. \quad (18.19)$$

(BJT noise models have a good handle on the fundamental physics, so BJT circuits actually follow the model.)

18.4.6 Externally Biased Cascode

Of the two promised ways we can improve the bandwidth, the simpler one is to apply a very quiet DC bias current I_{Eq} to the emitter of Q_1 , in addition to I_d . The value of r_E can be reduced considerably this way, further improving f_{RC} . For example, if we use $I_{Eq} = 20 \mu\text{A}$, r_E drops to 1.25 k Ω and f_{RC} is 1.27 MHz—quite a bit better than our original 1600 Hz, and enough for the circuit requirement. We start running into the input capacitance C_{in} of the op amp, which limits how fast we can go just the way C_d did before. Switching to a slightly faster op amp such as an LF357 and using $C_f = 0.5 \text{ pF}$ overcomes C_{in} and gets us to a 1.1 MHz 3 dB bandwidth for the whole circuit. The net bias current now has 10 times less than full shot noise, so (18.18) predicts that the SNR will be down 3 dB at only 330 kHz, which is not good enough. We could just as easily use $I_{Eq} = 200 \mu\text{A}$, so that the shot noise corner will be at 1.3 MHz, but this starts to get us into trouble. Let's look at why.

18.4.7 Noise Considerations

Generating that very quiet (much less than full shot noise) current is easy: a metal film resistor R_E to a well filtered supply is all that is required, provided that the resistor drops a large enough voltage.

If we make R_E drop N times kT/e , the noise power due to $i_{N\text{bias}}$ will be reduced by a factor of N^2 . In our example, with $I_{\text{Eq}} = 200 \mu\text{A}$ and $I_d = 2 \mu\text{A}$, if $R_E I_{\text{Eq}} = 2.5 \text{ V}$, the shot noise power from the bias current will be reduced to 10^{-2} times the photocurrent shot noise power, a negligible addition.

The real limitations come from base current shot noise and Johnson noise in R_E . The base current I_B has full shot noise, which limits the bias current noise to at least $1/\sqrt{\beta_0}$ times full shot noise current. (That β_0 is the DC current gain, which is what's relevant even at high frequency, since it's the DC value of I_B that sets the shot noise level.) Start out with a BFG25A/X or BFT25A, but consider using a superbeta transistor ($\beta \approx 1000$) like an MPSA18. Darlingtons can have huge betas, but they're fairly noisy, and so they often make disappointing bootstraps. From (18.16), the voltage noise of the driver stage is high, and the transconductance of the output stage is also high, leading to lots of noise current unless the driver stage is run at a pretty high current itself. A single superbeta transistor, perhaps with its own collector bootstrapped, is usually a better choice when base current shot noise is a limiting factor. Homemade Darlingtons can be better—you can make a good one from two BFG25A/Xs. At low currents you can make the bootstrap a fast JFET (e.g., a BF862). That gets rid of the base current problem, at the expense of lower transconductance and higher voltage noise.

Resistor R_E sources its full Johnson noise current into the summing junction, so that the net Johnson noise is that of the parallel combination of R_f and R_E . For convenience, we could simply put a 75 k Ω metal film resistor in parallel with D_1 (assuming that $-V_{\text{bias}} = -15 \text{ V}$). The improvement is enough to meet our design bandwidth, but the noise is degraded by 3 dB, and we have to raise the positive supply enough that the voltage drop across R_f doesn't saturate the op amp. This is a viable solution if we can make $-V_{\text{bias}}$ bigger, perhaps -45 V , so that R_E can grow and its Johnson noise current thereby shrink (as a by-product C_d will shrink, which is a great help). A simple charge pump followed by a capacitance multiplier will get you a nice quiet -24 V , which is usually lots.

The calculated transimpedance gain and CNR of the cascoded transimpedance amplifier appear in Figure 18.10, with and without an additional 30 μA I_{Eq} . Using a higher bias current makes it worse rather than better; there's a big improvement in bandwidth and mid-frequency SNR, but even in this best case, the 1 MHz SNR is down by 6 dB due to the bias current noise, so we have to go a bit further still.

Aside: DC Offset. A minor drawback to the externally biased cascode circuit is that the DC level at the output of the transimpedance amplifier is no longer zero at zero photocurrent. This offset can be trimmed out, but it will drift somewhat with temperature, so that more circuit hackery is necessary if a highly stable DC level is needed. Most of the time it isn't, especially since other drift sources such as etalon fringes are normally much more serious. If we can raise $+V_{\text{bias}}$ enough, a resistor from there to the summing junction can get rid of the DC offset without adding too much Johnson noise. A diode-connected transistor in series with this resistor will provide first-order temperature compensation.




Figure 18.10. Calculated response and CNR of the cascode transimpedance amplifier of Figure 18.8 at $I_d = 2 \mu\text{A}$, with and without a $30 \mu\text{A } I_{\text{Eq}}$.

18.4.8 Bootstrapping the Cascode

If even small DC shifts are obnoxious, or the required value of I_{Eq} is so large that base current shot noise is a limitation, another technique is superior: bootstrapping. As shown in Figure 18.11, driving the cold end of D_1 with a follower Q_2 forces the drop across C_d to be constant, at least at frequencies where X_{C2} is small and $X_{Cd} \gg r_{E2}$.

In order for this to be any use, the bootstrap has to have much lower impedance than the cascode, so make $I_{C2} \gg I_{C1}$. The bootstrap circuit is a bit more complicated to analyze for noise, but the results are nearly the same as for a biased cascode with the




Figure 18.11. Bootstrapping the unbiased cascode circuit reduces the effects of r_E and has performance similar to that of the biased cascode, without the offset current due to R_E .

same collector current. Assuming $I_{C2} \gg I_{C1}$, the noise current from Q_2 flowing to the emitter of Q_1 via C_d is

$$i_{N\text{bootstrap}} \approx \sqrt{\frac{I_d}{I_{C2}}} \sqrt{2eI_d \omega C_d r_{E1}} \tag{18.20}$$

to leading order in ω . This is approximately $(I_{C2}/I_d)^{1/2}$ times smaller than in the unbiased case. It grows linearly with ω , so although the bandwidth is increased by I_{C2}/I_d , the SNR is down 3 dB at about $\omega = (I_{C2}/I_d)^{1/2}/(r_{E2}C_d)$, just as in the biased cascode case.

Bootstrapping basically replaces the r_{E1} of cascode device Q_1 with the r_{E2} of follower Q_2 , which gives an improvement of I_{C2}/I_{C1} times in bandwidth. By essentially eliminating the capacitive loading on Q_1 , it also eliminates the effects of Q_1 's voltage noise.

This trick is reminiscent of the old joke about the cowboy who, after he fell into a well, "pulled himself up by his bootstraps." Bootstrapping suffers the same multiplication of the voltage noise of the follower that we saw in the transimpedance amplifier. However, here the RC product is not $R_f C_d$ but $r_{E1} C_d$, a factor of 8 smaller, and the follower's v_N is usually smaller as well, so this is not nearly as great a problem as it is with the transimpedance amplifier.

We wouldn't be doing this if current errors weren't important, so we'll use a superbeta MPSA18 with $I_{C2} = 290 \mu\text{A}$. The largish C_{cb} of this device appears in parallel with C_d , so it hardly matters; the C_{cb} forms a voltage divider with C_d , but since it's 50 times smaller, it doesn't matter much either. The $7 \mu\text{A}$ flowing through R_{bias} makes the cascode a bit faster, and the offset voltage and drift are canceled by the matching resistor and the V_{BE} of Q_3 . All this together improves the CNR to 1 dB above the shot noise limit in the flatband, degrading to 3 dB above shot noise at 1 MHz, and gets us a 3 dB bandwidth of 2 MHz. The final circuit is shown in Figure 18.12, its calculated performance in Figure 18.13(a). Figure 18.13(b) shows the prototype's measured performance, which is somewhat better than the worst-case calculation. The measured shot noise/dark

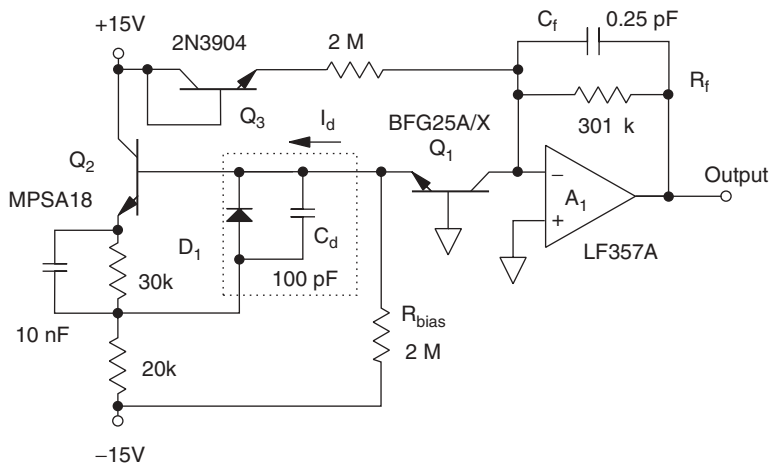




Figure 18.12. The final circuit: cascode Q_1 plus bootstrap Q_2 cope with the obese 100 pF diode, and diode-connected Q_3 cancels the V_{BE} drift of Q_1 .



(b)



(a)

Figure 18.13. Performance of the final circuit. (a) Calculated CNR is down only 3.3 dB at 1 MHz. (b) Output of the prototype: bottom—dark noise; top—noise with 2 μ A photocurrent, showing better than calculated SNR. The setup has an overall voltage gain of 2.7.

noise ratio is 9.5 dB at low frequency, dropping to 4.5 dB at 1 MHz.[†] These numbers correspond to total noise 0.5 dB over shot noise at low frequency, rising to 1 dB over shot noise at 1 MHz. There are no unpleasant surprises, which helps us to be confident that we finally understand the circuit.

18.4.9 Circuit Considerations

Although the cascoding and bootstrapping tricks seem like a free lunch, nevertheless, like all circuit hacks, they have their limitations. The linearity of the transistor's gain may not be as good as that of the photodiode. Normally a small amount of bias (20 μA or so) will linearize the transistor well enough. For the most critical applications, use a highly linear transistor, a Darlington, or (if driven to it) FET. FETs have no base current nonlinearity, but their transconductance is low and their noise high, so they don't usually work as well. The BF862 is sometimes an exception.

18.4.10 One Small Problem . . . Obsolete Parts

Most inconveniently, National Semiconductor discontinued the 75-cent LF357 in 2004. There's no single replacement with its particular combination of virtues, not even among the fancy \$50 parts (See Table 18.4). For the current design, we need $R_f \geq 300 \text{ k}\Omega$ to keep the SNR drop to 0.5 dB at low frequency and 1 dB at 1 MHz with a minimum I_d of 2 μA . If the maximum I_d is less than 100 μA , we can get away with a $\pm 5 \text{ V}$ amplifier, in which case things are easier: we can drop in a \$5 OPA656, and we're done—its input capacitance is about the same as the 357's and it has lower voltage noise. We can leave C_f the same, since the 656 is unity gain stable and we don't need all of its bandwidth. Performance will be nearly identical to the LF357 circuit.

If the photocurrent can go much higher than that (e.g. 500 μA), we really need the extra 10 dB dynamic range we get from $\pm 15 \text{ V}$ supplies. In that case we can either use a compound amplifier, for example, that BF862 differential pair we looked at earlier running into a $\pm 15 \text{ V}$ op amp, or an external fixed-gain buffer inside the feedback loop. The external JFET pair is more easily frequency compensated (e.g., with a lead-lag network around the op amp), but the fixed-gain booster doesn't mess up the nice input accuracy of the OPA656.

Alternatively, we can give up about 1 dB of SNR and use the $\pm 15 \text{ V}$ OPA627. Its other specs are excellent for the purpose, but its input capacitance is 8 pF, so it will oscillate with a $300 \text{ k}\Omega R_f$, and when we increase C_f to compensate, it becomes too slow. Thus we have to use $R_f \approx 80 \text{ k}\Omega$, which costs us 0.7 dB in SNR at all frequencies. A bipolar op amp such as an OP37 will have lower e_N and C_{in} , but its 1 Hz i_N is 0.6 pA, which is comparable to the 0.8 pA shot noise and so costs us 2 dB SNR in the present instance, though it'll be superior at higher I_d . Note that these problems have nothing to do with photodiode capacitance, which has already been fixed by the bootstrapped cascode—the problem here is amplifier C_{in} , that is, the op amp tripping over its own big feet. In general, technological change has made TIA design easier in some ways and harder in others.

[†]Remember that the distance between the two curves is total noise/Johnson noise, not shot noise/Johnson noise.

TABLE 18.4. Suggested Transistors for Cascode Transimpedance Amp and Bootstrap Service^a

Device	Manu- facturer ^b	f_T (MHz)	@ I_C (mA)	β	@ I_C (mA)	$R_{E'}$ (Ω)	C_{ob} (pF@V)	Remarks
NPN								
BFG25A/X	P	5000t	1	50	0.5		0.2t @ 1	Excellent device
BFT25A	P	5000t	1	50	0.5		0.3t @ 1	Easier to get
BFG505X	P	9000t	5	60	5		0.2 @ 6	Higher power, good β linearity
BF240	P+	600t	1	65	1		1t @ 1	
MAT-04	AD	300t	1	175	.01–1	0.6	17t @ 0	Quad, good β linearity, low $R_{E'}$
MPSA14/ 2N6426	Many	125	10	10k	10	0.3t	14t @ 0	Good Darlington, MMBTA14 SMT
MPSA18	M+	160t	1	1000t	0.5–10		3t @ 1	Super- β , good for bootstraps
MPSH20	M	400	4	25	4	1.2t @ 0		
UPA103	NEC	9000		40	5t			Quint, good β linearity, poor $R_{E'}$
2N2484	M	60	0.5	200	0.5		6 @ 5	Very well specified
2N3904	All	300	10	100	1		5t @ 0.05	Ubiquitous; manufacturers differ; well spec'd
PNP								
BFT92	P	5000t	14	20	14		0.7t	
MAT-03	AD	40t	1	90	0.1	0.75	30t @ 5	Dual; accurate, slow
MPSA-64	Many	125	10	10k	10	0.6t	15t @ 0	Darlington
MSC2404	M	450	1	65	1		1 @ 6	
JFET								
BF862	P	715t	10	35	10	10	1.9	0.8 nV/Hz ^{1/2} typical
2SK369	T	50	10	40	10	50	80	0.7 nV
HJFET								
NE3509	N	18G	10	80	10	0.04t	0.3t	$T_N = 35\text{ K @ }2\text{ GHz}$, $R_{DSon} \sim 6\ \Omega$; $C_d = 0.4\text{ pF}$

^aThese are mostly through-hole devices for easier prototyping, but surface mount equivalents exist. The JFETs are most useful as bootstraps and outboard differential pairs to reduce the noise of another op amp.

^bManufacturer codes: AD, Analog Devices; M, ON Semiconductor; N, NEC; P, NXP Semiconductor; T, Toshiba.

Aside: Input Capacitance Specs. Many newer op amps have two input capacitance specifications, for common mode and differential signals; for instance, the OPA656's are 0.7 pF common mode and 2.8 pF differential. This has to do with the way their input structures work—because the sources of the differential pair are connected together, differential signals see the two C_{gs} of the two input devices in series, but common-mode signals don't.

Unfortunately none of the manufactures specifies how these capacitances are measured, so the safe procedure is to use the larger one, or perhaps $C_{diff} + c_{cm}/2$. The OPA657, a

decompensated OPA656, has 4.5 pF of differential C_{in} , surprisingly high for its 1.6 GHz GBW.

Gotcha: SPICE Macromodels. If you're doing your front end design with SPICE or another circuit simulator, good luck to you—with care you can do a reasonable job, but be suspicious: many op amp models contain inaccurate noise models, or none at all. An appalling number also omit the input capacitance of the op amp, which as we've seen is a vital parameter. If you use simulation as a substitute for thought, you'll get the performance you deserve; sometimes the SNR from a correctly formulated SPICE simulation using the manufacturer's op amp models can be optimistic by 20 dB or even more. (Sometimes the C_{in} value is in the model but not in the data sheet, which is another issue.)

Always calculate the noise analytically (it isn't especially difficult) and compare with the SPICE model and with the prototype. Linear Technology has a very well-regarded free SPICE program, LTSpice, that you can download. Generally, when an op amp macromodel simulation does something uninituitive, such as driving its output beyond the supplies, it's very likely to be wrong. (Also note that SPICE models will have the "typical" data sheet characteristics, which isn't enough to base a design on.)

18.4.11 Power Supply Noise

All through this chapter, we've been doing our noise calculations assuming that the photodiode bias voltages have been noiseless. While this is quite doable, it won't happen by accident. The author *always* uses capacitance multipliers to make these bias supplies, and usually runs the front end amplifiers from them too, unless there's a good reason not to. Any jumping around of the supplies will be transferred directly into the photocurrent, via the photodiode capacitance—the noise current will be

$$I_{Nsup} = V_{Nsup}\omega C_d, \quad (18.21)$$

making it just as important an effect as amplifier noise—100 μV of wideband supply noise is just as serious as 100 μV of amplifier noise. Since this is a purely AC effect, capacitance multipliers are a better match than voltage regulators here. Have a look at Example 14.1 for more—none of the circuits in this chapter will work properly without quiet bias supplies.

Even more insidiously, noise can come in via the power supply leads of your op amps. Op amps have power supply rejection (PSR) ratios of 60 dB or more near DC, but it rolls off at higher frequencies. Linear voltage regulators can exhibit nasty noise peaks—their outputs look like small value inductors in series with very small resistors, so with a big bypass cap, you can produce huge noise peaks at the resulting resonant frequency. Putting a few ohms' resistance in series with the regulator's output pin (before the first bypass) will kill the Q of the resonance and make the supply noise much better behaved, at the price of slightly degraded DC regulation. If your front end has a noise peak in the 1–100 kHz range that you can't understand, try this trick.

18.4.12 Beyond Transimpedance Amps: Cascode + Noninverting Buffer

If the C_{in} of your op amp is still a serious inconvenience, you can eliminate the transimpedance amplifier in favor of a simple load resistor following the cascode transistor,

with a low capacitance buffer following. Second-stage noise need not be a limitation, even though the buffer has a gain of 1; because the buffer's output impedance is low, the next stage can be a low v_N bipolar amplifier such as an LT1028. A good voltage follower (e.g., a bootstrapped emitter follower, see Section 18.4.8) can have an input capacitance of less than 0.25 pF along with a 1 Hz noise of $5 \text{ nV/Hz}^{1/2}$, so C_{in} is not inescapable.

Another insidious problem shows up when we let the collector of Q_1 swing: the Early effect. A transistor has a collector current that depends somewhat on its collector-emitter voltage V_{CE} , so that its output impedance has a large but finite value. For small variations of V_{CE} , this effect is approximately linear; increasing V_{CE} increases the collector current. If we plot I_C versus V_{CE} , and extrapolate linearly to the point where $I_C = 0$, the intercept is the *Early voltage* V_{Early} . This voltage is normally in the thousands of volts for general purpose transistors, so it is of little concern. For RF devices and those with very high β , V_{Early} is much smaller (as low as 40 V), and so the Early effect is a significant source of gain error and nonlinearity in common-emitter amplifiers. It is less troublesome in the common-base configuration, but do look carefully for nonlinearity at large signal swings, and take the transistor's collector impedance into account.

Example 18.2: Current-Mode Amplifier. One of the biggest problems we've run into in TIA design is that the resistors are so very noisy compared with the active devices. It's worth trying to build a front end without resistive feedback, by basically stuffing the photocurrent into the base of a BJT, with sub-Poissonian current feedback. This is more or less what the bootstrap and cascode do by connecting (at AC) the PD between the base and emitter. In the bootstrapped cascode, the cascode protects us from the $e_N C_d$ noise peak, and the bootstrap reduces the load impedance seen by the photocurrent. If we can combine the two functions in one device, we might be able to reduce the noise by 3 dB. One possible way to do this is shown in Figure 18.14. (Don't try building it as shown—there's a lot of stuff missing.)

Input transistor Q_1 's 1 Hz input current noise is the shot noise of the base current, $\sqrt{2eI_C/\beta}$, and (neglecting $R_{B'}$ noise) its voltage noise is $(2eI_C)^{1/2}/g_m = kT(2/eI_C)^{1/2}$.




Figure 18.14. Current-mode photodiode amplifier has slightly better performance than the bootstrapped cascode TIA, at the expense of extra engineering. (This is a conceptual schematic only—the analysis considers only the first stage, that is, Q_1 .)

Since the base and collector shot noise are essentially independent, v_N and i_N are uncorrelated as usual. The input resistance is that of the BE junction, $r_{in} = kT/(eI_B)$ (assuming it obeys the diode equation).

The open-loop bandwidth and input-referred noise are approximately those of the input transistor. We'll ignore the Miller effect, since in a real design we'd get rid of it with a cascode, so the current gain is set by β and the parallel combination of C_d and r_{in} :

$$A_{IOL} = \frac{\beta}{1 + j2\pi\beta kTC_d/(eI_C)}. \quad (18.22)$$

Similarly, the total input-referred noise current is

$$i_{Ntot}^2 = 2e(I_C/\beta + I_d) + 2/(eI_C)[kT(2\pi fC_d)]^2 \quad (18.23)$$

(we've ignored the distinction between β and $\beta + 1$ that circuit instructors so often insist on, and have assumed that the small- and large-signal betas are the same, that is, $\partial I_C/\partial I_B = I_C/I_B$). At low frequency, the 1 Hz SNR is $1 + I_C/(\beta I_d)$ times worse than the shot noise, so we need to limit $I_C < \epsilon\beta I_d$ for some ϵ , which we do by pulling $(1 - \epsilon)$ times I_d with a current sink in a feedback loop, making ϵ a convenient optimization parameter. To stay within 1 dB of the shot noise, $\epsilon < 0.25$. The noise doesn't start to rise until the two terms in (18.23) become comparable, that is, at

$$f_N = \frac{eI_d\sqrt{\beta\epsilon(1+\epsilon)}}{2\pi kTC_d}. \quad (18.24)$$

For $\epsilon = 0.13$ (0.5 dB above shot noise), $\beta = 1000$, $I_d = 2 \mu\text{A}$, and $C_d = 100 \text{ pF}$, $f_N = 1.49 \text{ MHz}$, which is slightly better than the bootstrapped cascode result, as expected. This method takes a fair amount more engineering than the cascode TIA, because there's the whole current feedback loop to design, including getting its transient response, temperature compensation, output drive, and supply rejection right, but there's another quarter-turn of the crank to be had if you really need it.

18.4.13 Choosing Transistors

The cascode and its variants can provide a huge performance gain, but only if the devices and operating parameters are appropriately chosen. Fortunately, there are a few rules of thumb to help with this.

The main one is to always start with an NXP BFG25A/X or BFT25A as the cascode device, and use something else only when driven to it.[†] These parts are the same small geometry NPN RF transistor die in different packages, and are about as near to magic as you can get in an SOT-143 surface mount package (or an SOT-23 for the BFT25A). Lest you think that this is mere infatuation, here are the highlights: $I_{C\max} = 6.5 \text{ mA}$ (best below 2 mA), $f_T \approx 5 \text{ GHz}$ at $I_C = 1 \text{ mA}$, highly linear β of about 100 (very good for an RF device), and $\gtrsim 50$ for I_C down to the nanoamps, $C_{ob} \approx 0.2 \text{ pF}$ at 1 V collector

[†]The author favors the G over the T for historical reasons—it has the same pinout as the late lamented MRF9331—and because its feedback capacitance is a bit lower.

to base, $V_{\text{Early}} \approx 50 \text{ V}$, price \$0.30. If these specifications don't excite you, you haven't spent enough Saturday nights designing front end amplifiers.†

This device is a near-universal choice for the cascode transistor in an unbiased configuration, since a minimum β of 30 means that the base current is 30 times less than the collector current, and hence its shot noise power is also 30 times less. The collector current comes from a photodiode, and hence has full shot noise, so the noise power goes up by about 3% (0.12 dB) in the worst case, which is a small price to pay for an $8\times$ bandwidth increase. Anyway, we can recover more than that by jacking up R_f (and so reducing its Johnson noise current), since $R_f C_d$ no longer sets the bandwidth.

Life gets a bit more complicated in the biased case. Here the base current is not 3% of I_d , but 3% of I_{Eq} , which can easily be comparable to I_d . Since it still has full shot noise, this can represent a significant noise increase. In this case, you can use a superbeta or Darlington for the cascode device, as we did in the bootstrap example. You can build the Darlington from a pair of BFG25A/Xs since I_{C3} is still small. Bias the driver transistor so that its I_B is about 0.1–0.25 times I_d , to keep its voltage noise down without contributing large amounts of shot noise.

When a fast transistor running at high collector current has a capacitive load in its emitter circuit, the input impedance has a tendency to look like a negative resistance at high frequency, leading sometimes to UHF oscillations if the driving impedance is too low. These may be very difficult to see on an oscilloscope, but will make the amplifier act very mysteriously. A 100Ω resistor in the base of Q_1 solves this problem in most cases.

In case you have a real embarrassment of riches, and your photocurrent is too large to allow you to use the BFG25A/X, you can use another device of the same general type; a BFG505, for example, or a small-signal Darlington such as an MPSA14. The problem with general-purpose devices such as the 2N3904 is that their f_T s roll off so badly at low collector currents; a 2N3904 running at 10 mA I_C is a 350 MHz transistor, but at 0.1 mA, it's about a 35 MHz transistor, and it gets correspondingly slower as the collector current declines. Remember that it's the AC value of β that matters for passing signal, so unless your transistor has $f_T > 200 \text{ MHz}$, that nice β of 200 at DC won't be there at 1 MHz. Table 18.4 is a comparison chart of several transistors that are good for these jobs.

The current gain at the operating frequency needs to be at least 20 for the unbiased case, and correspondingly more with bias. For single devices, whose betas go as $1/f$, f_T needs to be about $20\times$ the operating frequency (Darlingtons can hold up longer). This is often a problem for run-of-the-mill small-signal transistors. The trouble is that they are relatively large-geometry devices, often able to handle currents of 200 mA or more. In photodiode front ends, we are running them way down on the low current end of their operating ranges, and they are not optimal there. The BFG25A/X's virtues stem mainly from its small die size. Most transistor data sheets don't guarantee f_T values except at a single operating condition: 10 V collector–emitter, and a few milliamps I_C . A rule of thumb is that well below the f_T peak, f_T goes as I_C , but that within a factor of 10 or so below the peak, the dependence is more like $\sqrt{I_C}$. These rules allow extrapolation of published curves to very small collector currents.

†There are much faster small BJTs available, for example, the 25 GHz BFG424F, but they don't have the BFG25A/X's beta linearity or highish Early voltage, and anyway it's hard to keep something that fast from oscillating when you hang a capacitance on its emitter.

With most devices, it is impossible to rely on worst-case specifications way down in the mud, because there aren't any—frequently not even typical ones. Use the typical specs, but build in a safety factor of at least $3\times$ on f_T and β .

In the biased cascode arrangement, positive or negative photocurrent is equally acceptable, provided that the bias is larger than the largest expected photocurrent. This is convenient, since no PNP transistor comparable to the BFG25A/X is available.

In an unbiased cascode (bootstrapped or not), a positive photocurrent requires a PNP transistor or P-channel FET. These are not as good as their NPN or N-channel relatives. The MMBR521 is a good, fast PNP (5 GHz) for high currents, but is not as much use for low ones as its beta falls off badly. Probably the best thing to use is a small-signal Darlington such as an MPSA64, which has been found to work well. Darlingtons ideally should have a $1/f^2$ behavior, because the driver transistor returns the output device's base current to the collector circuit until its own f_T is approached. The deficiencies of the cascode device are partly hidden by the bootstrap, if used.

Some cutting and trying is necessary to get a good result. Databooks and SPICE models tend to get out of date, since many of the transistors we'd like to use are old designs. As the old production processes are closed down, and the old part numbers reimplemented on newer processes, the parameters change, but the databook specifications stay the same. It is very inconvenient to have an unannounced device change break your design. For a situation like this, where the parts are cheap but the consequences of a change can be painful, the smart plan is to buy as many devices as you anticipate ever needing, and stick them in a safe somewhere (see Section 14.7.3). At the very least, keep enough known-good devices on hand to last you for however long it may take to change the design.

18.5 HOW TO GO FASTER

We had a struggle to get to 1 MHz with a $2\ \mu\text{A}$ photocurrent while staying in the shot noise limit. Is there any hope that we can do shot noise limited measurements at higher speed? Well, yes there is. In our example, we purposely chose a moderately high capacitance photodiode and a low light level. We saw that the RC corner frequency f_{RC} from a diode capacitance C_d and a photocurrent I_d was $I_d/(2\pi \cdot 0.2\ \text{V} \cdot C_d)$ if we were to be within 1 dB of the shot noise. If we use a smaller, fast photodiode with a capacitance of 10 pF, and run it at a photocurrent of $100\ \mu\text{A}$ with a $2\ \text{k}\Omega$ R_L , that corner is not at 16 kHz, but at 8 MHz. This is not the limit either; using a biased cascode with a bootstrapped follower (0.25 pF each) will get us to beyond 200 MHz.[†] VHF design is actually quite a bit harder than this, since all our analyses have been based on RC circuits alone. Above 200 MHz, *everything* has enough inductance to worry about, and stray capacitance is generally the limiting factor.

Aside: Optical Communications. In optical communications, and especially in the emerging area of short-range optical interconnection (on-chip, chip-to-chip, and module-to-module), it is frequently necessary to go a great deal faster than this—20 GHz or faster. This is generally done with extremely small photodiodes made from compound semiconductors such as InP or InGaAs, closely integrated with IC

[†]One design meeting this description has gone faster than 200 MHz, shot noise limited above $50\ \mu\text{A}$ I_d , with a 5 pF photodiode.

preamplifiers designed for the specific application. Instrument builders aren't that rich, usually, but if you can piggyback on this technology, consider doing so. (Of course, due to widespread reliance on erbium-doped fiber amplifiers (EDFAs), most telecom detector modules aren't all that quiet.)

There are situations where it is frankly impossible to reach the shot noise in the required bandwidth with the available light and ordinary photodiodes. Long distance fiber optic communication is a good example. In situations like that, you may be forced to use avalanche photodiodes or photomultipliers, which we discussed in Section 3.6, or in a fiber receiver, an optical preamplifier such as an EDFA. These devices have serious drawbacks and should not be used frivolously. With an APD running at a gain M , you can reduce the load resistor by a factor of M^2 without reducing the SNR compared with $M = 1$ (see Section 3.6.3).

The general rule that more photocurrent allows smaller resistors, and smaller photodiodes run at higher bias have lower capacitance, gets you most of the way there, most of the time. Nonetheless, there is one specialized VHF/UHF technique that is worth mentioning, because it is easily understood and implemented: LC networks (see Section 14.3.10).

18.5.1 Series Peaking

The simplest case of such a network is *series peaking*, which is nothing more than putting an inductor between the photodiode and the load resistor, as shown in Figure 18.15.

The peaking coil L provides positive reactance X_L at high frequencies, which partially cancels the negative (capacitive) reactance of C_d . The cancellation is far from perfect, because the magnitude of X_L rises with frequency while X_C 's falls. Nonetheless, a network like this can provide a useful bandwidth increase without an SNR penalty worth worrying about. The ideal photocurrent sees a load impedance (including C_d) of

$$Z_L = \frac{R + j\omega L}{\left(1 - \frac{\omega^2}{\omega_0^2}\right) + j\frac{\omega}{\omega_0 Q}} \tag{18.25}$$

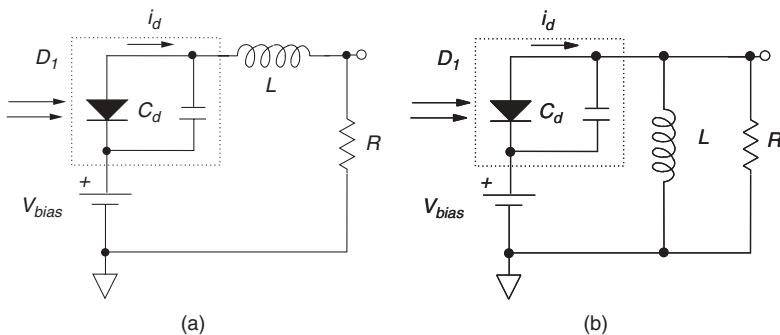


Figure 18.15. Adding an inductor to a photodiode. (a) Series peaking with coil L increases the RC_d bandwidth by $1.4\times$ in a baseband system, or $2\times$ in a narrowband AC system. (b) Shunt peaking keeps the bandwidth at $1/(2\pi RC_d)$ but moves the passband up to $f_c = 1/(2\pi\sqrt{LC_d})$. More complicated networks can do better.

where ω is $2\pi f$,

$$Q = \sqrt{\frac{L}{R^2 C_d}} = \frac{\omega_0 L}{R} = \frac{1}{\omega_0 R C_d}, \quad (18.26)$$

and $\omega_0 = (LC_d)^{-1/2}$ is the resonant frequency of L and C_d alone.

The load impedance Z_{in} is $R_{in} + jX_{in}$. In the absence of losses in L , all the power dissipated by I_d in the real (resistive) part R_{in} gets transferred to R . The total power P dissipated by I_d is $I_d^2 R_{in}$, where R_{in} is

$$R_{in} = \frac{R}{(1 - \omega^2 LC_d)^2 + \omega^2 R^2 C_d^2}. \quad (18.27)$$

Computing the signal-to-noise ratio is a bit more subtle here. The rms noise voltage at the output is the Johnson noise current i_{Nth} of R , times the magnitude $|Z_{out}|$ of the output impedance of the whole network, which is

$$|Z_{out}| = \frac{R(x - 1)}{\sqrt{(x - 1)^2 + x/Q^2}}, \quad (18.28)$$

where $x = \omega^2/\omega_0^2$.

At frequency ω_0 , the series combination of L and C_d is a dead short; thus the Johnson noise current i_{Nth} from R generates no noise voltage whatsoever across R . Nevertheless, a signal power of $I_d^2 R_{in}$ is delivered to the load. This one-way transfer seems a bit odd, not to say impossible.

The reason for it is that we've assumed that the photodiode has no dissipation of its own—that it is a perfect current source in parallel with an ideal capacitor. Such a device is unphysical, since there is no limit to the signal power that it can deliver to a sufficiently high impedance load. There will be some resistance R_s in series with C_d , and even if there weren't, loss in L would prevent the output impedance from going to 0. Nevertheless, in principle, this is a very lucky break: it means we can potentially open a shot noise limited window in the middle of a wall of Johnson noise.

The 1 Hz Johnson noise power is $i_{Nth}^2 |Z_{out}|$, and the 1 Hz shot noise power is $2eI_d R_{in}$, so the 1 Hz CNR is

$$\text{CNR}_{1\text{Hz}} = \frac{i_d^2 R_{in}}{2ei_{dc} R_{in} + 4kT(x - 1)/\sqrt{(x - 1)^2 + x/Q^2}}. \quad (18.29)$$

Of course, we haven't done anything about the amplifier's intrinsic noise, except to short out its input. Some amplifiers work well with shorted inputs, but some don't. None of them has zero noise with a shorted input, unfortunately. You have to ask your amplifier supplier, or do your own simulations and measurements if you're designing your own. Nonetheless, the resonant enhancement in R_{in} is often enough to get you to the shot noise.

If R is not a real resistor, but instead the input resistance of an RF amplifier, the appropriate value of T to use is not the ambient temperature, but the noise temperature of the amplifier. It may not be immediately obvious from (18.29) that things have improved, but they have. For narrowband applications that require high frequency operation (e.g.,

heterodyne systems using acousto-optic modulators), choosing L to resonate C_d and making R small will make the denominator quadratically small. R is increased by a factor of approximately Q^2 .

Reasonable values of Q to use depend on the frequencies encountered, but will seldom be above 10 and never above 25. One thing to remember is that the impedance transformation is accomplished by a large current circulating around the loop. This current is Q times larger than the AC photocurrent, and leads to a large AC voltage across C_d . It may seem strange to be increasing the swing across C_d now, when we worked so hard to reduce it before. The difference is that in a pure RC circuit, all the current going into C_d was lost, and here it isn't. There are limits to this, of course, since at sufficiently high frequency the photodiode stops looking like a pure capacitor, and its intrinsic losses become important.

Example 18.3: Narrowband 160 MHz Heterodyne System. If we have a heterodyne system using two passes through a typical 80 MHz acousto-optic modulator, the photocurrent component we care about is in a narrow band centered on 160 MHz. Using a photodiode with $C_d = 10$ pF and $I_d = 30$ μ A, we could in principle use a huge R such as 10 k Ω , and get a shot noise limited CNR of 140 dB in 1 Hz, as in Figure 18.2. Of course, the RC time constant is 100 ns, so that the voltage would have rolled off by 100 times at 160 MHz, which makes it a bit impractical to use. If we choose $R = 100$ Ω to control the rolloff, then we drop only 3 mV across it, and we're firmly in the Johnson noise limit, with a CNR of 126 dB in 1 Hz. On the other hand, if we put in an inductor of 99 nH, and work straight into a 50 Ω load, then from (18.27), $R_{in} = 197$ Ω . That's 3 dB better, but still far from shot noise limited. Decreasing R to 12 Ω , perhaps by using a 2:1 RF transformer (4:1 in impedance, see Section 14.3.14) between the amplifier and inductor, improves R_{in} to 825 Ω , which would notionally drop 25 mV. (The impedance at DC is only 12 Ω , but for noise purposes it's the DC value of I_d times the AC value of R_{in} that matters.) The Q of this network is 8.3, which is reasonable. The FWHM of R is ω_0/Q , or about 19 MHz, which is equal to that of the equivalent 825 Ω ·10 pF lowpass, as we expect (see Chapter 15).

If the amplifier has a noise temperature of 75 K, then its noise power is only a quarter that of a room temperature resistor; thus $I_d R_{in}$ only needs to be 13 mV for the shot noise to begin to dominate. Thus such an amplifier plus a simple series inductor and a 2:1 transformer will get us to the edge of the shot noise limit. A slightly more complicated network can do better, for example, a π -network or a tapped tank circuit,[†] but this is a good place to begin.

18.5.2 Broader Band Networks

There are two other common uses for reactive elements in photodiode amplifiers: extending the bandwidth of a baseband detector, and a wideband application (say, an octave) well away from DC. A slightly more complicated network (e.g., a π - or T -network) can match a 50 Ω RF amplifier input to weird source impedances. Resonating away the capacitance of the photodiode works well, but only at one frequency, whereas for these applications we need a decent bandwidth as well as a high operating frequency.

[†]See, for example, Terman or *The Radio Amateur's Handbook*.

Just putting a judiciously chosen inductance L in parallel with D_1 (with appropriate biasing and DC blocking, so as not to short it out) moves the low frequency RC bandwidth to the resonant frequency of L and C_d , $f_0 = 1/(2\pi\sqrt{LC_d})$. This happens approximately symmetrically; for example, a 40 MHz lowpass network becomes a bandpass of ~ 40 MHz full width.[†] This assumes that the load resistance remains in parallel with D_1 as well, and that the Q is large enough that the low frequency response has rolled off a long way before hitting DC. If you don't mind building your own RF amplifiers, this can be a good technique; a dual-gate GaAs FET follower can do a good job of buffering such a circuit without loading it down unduly.

The circuit of Figure 18.15 has response all the way to DC. Although we used it in a relatively high- Q application earlier, it is also useful at Q s of around 1 for improving baseband networks. It could in principle be used with transimpedance amps as well. (You may want to try it out. Don't expect good performance from inductors larger than $50\ \mu\text{H}$.)

From (18.27), we can calculate the points at which R_{in} has fallen to α times its peak value:

$$\frac{\omega_\alpha^2}{\omega_0^2} = 1 - \frac{1}{2Q^2} \pm \frac{1}{Q} \sqrt{\left(\frac{1}{\alpha} - 1\right) \left(1 - \frac{1}{4Q^2}\right)}, \quad Q \geq \sqrt{1/2}. \quad (18.30)$$

This simple exact form is valid only for $Q \geq 1/\sqrt{2}$.

Example 18.4: Peaking a Baseband Network. In the previous example, we used a network with $R = 12\ \Omega$, $L = 100\ \text{nH}$, $C_d = 10\ \text{pF}$. This resulted in a peak R_{in} of about $825\ \Omega$. What if we needed to go from DC to 160 MHz? Such a network will obviously not have high Q , and so its input resistance will be of the same order as R rather than being multiplied by a high value of Q^2 .

In a pure RC circuit, a bandwidth of 160 MHz allows a maximum of $100\ \Omega$ for R . With a maximally flat ($Q = 0.414$) network, the 3 dB corner is $\sqrt{2}$ times higher than that set by the RC , so that we can use a $140\ \Omega$ load, and get a 1.5 dB improvement in the SNR for the same bandwidth (the improvement will be greater near f_c because of the resonance, as above). This value of Q gives the maximum bandwidth improvement for a fixed C and R . This does not give us the tremendous bandwidth improvements we saw in the transimpedance amplifier section, but then that was low frequency, and this is VHF. Remember that it's hard to get decent low frequency, high value inductors, so that you can forget peaking a high impedance, low frequency network.

18.5.3 Matching Networks and Bode's Theorem

People working in the 100 MHz to several GHz range often find themselves limited by the capacitance even of an InGaAs APD, which is usually quite small. There are lots of differences between that regime and baseband which make front-end design a challenge. If you're working in a bandwidth of less than an octave, you can do some reactive matching tricks that help a great deal. There's an inescapable trade-off between mismatch and bandwidth. For the common case of a parallel RC circuit, there is a

[†]Just off resonance, X_L and X_C have equal and opposite slopes, making the total reactance change twice as rapidly with frequency as in the lowpass prototype. This reduces the 3 dB width by a factor of 2, making the total bandwidth about the same as the lowpass prototype.

theorem of Bode that states[†]

$$\int_0^\infty \ln \frac{1}{|\Gamma|^2} d\omega \leq \frac{2\pi}{RC}. \quad (18.31)$$

(Darlington and Fano later published more general versions for complex impedances, but Bode's is the most useful.) Thus if you don't mind a return loss of 6 dB (75% efficiency), set $\Gamma = 0.5$, and you can get a BW of $2\pi/\ln(4)$ or $4.5\times$ the RC bandwidth. The 25% average passband loss is -1.24 dB compared to a perfectly matched resistance. Considering that the average passband loss of the unaltered RC rolloff is $10 \log(\arctan(1)) = -1.04$ dB, we get a $4.5\times$ bandwidth improvement for an additional signal loss of 0.2 dB, which is a pretty good payoff. (Note that Γ must be close to 1 almost everywhere in order that the integral have a finite value.)

The basic rule of thumb is that if you use a three-element tee network (sometimes all inductors), you can get within 0.5 dB of the absolute physical limit for bandwidth with a given R_{in} , so that it isn't worth doing anything more complicated. This works by transforming the 50Ω input impedance of your amplifier into some much larger effective load impedance on the diode. Figure 18.16 shows a 10 pF detector matched over a 110–220 MHz band with an effective load of 1.2 k Ω using this trick, which makes it shot noise limited from $50 \mu A$ up with a quiet amplifier. You may not need as large a load impedance as you think, because good RF amplifiers have a noise temperature much below 300 K—some are lower than 50 K. The Miteq catalog has some with NFs below 0.5 dB at 1 GHz, which is pretty impressive—a noise temperature of 36 K. The nice thing about this is that a lossless matching network transforms this low noise resistance into the equivalent of a cryogenically cooled resistive load, so you can be shot noise limited at much lower IR_L values; $2kT_N/e$ for this amplifier is not 50 mV but 6 mV, which is good for an $8\times$ bandwidth improvement over a 300 K load, at the same SNR.

18.5.4 T-Coils

One excellent place to go to learn about building amazingly fast baseband networks is a Tektronix oscilloscope service manual from the 1960s or 1970s, when discrete circuitry




Figure 18.16. Wideband matching network.

[†]Hendrik W. Bode, *Network Analysis and Feedback Amplifier Design*. Van Nostrand, New York, 1945, Section 16.3. Bode's book is well worth reading if you can find a copy.

dominated. The constant-resistance T-coil of Figure 18.17 is a gem you'll find there (and also in Jim Williams's books). The amazing thing about it is that the diode sees a constant load resistance of R_L , and the 10–90% rise time is exactly the same as if only the diode capacitance were loading it—no current is wasted in the resistor while charging the capacitor— $2.8\times$ faster than the RC alone. For a pure capacitance at C_d , the design is symmetrical: $L_1 = L_2 = L$. The design equations are[†]

$$L_T = R_L^2 C_d, \quad C_b = \frac{1-k}{4(1+k)} C_d, \tag{18.32}$$

$$M = kL, \quad \delta = \frac{1}{2} \sqrt{\frac{1+k}{1-k}},$$

where M is the mutual inductance, $L_T = 2L + 2M$ is the end-to-end inductance, and δ is the damping factor, $\delta = 1/(2Q)$. (Don't confuse this with an ordinary T-network—the mutual inductance is key to its operation.)

Example 18.5: Constant-Resistance T-Coil. Getting 30 MHz of bandwidth with a 10 pF photodiode requires a 530 Ω load resistor. Using a T-coil, we can run 1.5 k Ω with a 10–90% rise time of 11 ns and a 3 dB bandwidth of DC–30 MHz, with no overshoot ($Q = 0.707$). The component values are $L = 8.44 \mu\text{H}$, $M = 2.81 \mu\text{H}$, $C_b = 1.25 \text{ pF}$. This represents a 10 dB signal power increase, and since the Johnson noise power is independent of R , a 10 dB SNR increase in a Johnson noise limited system. One problem with the T-coil is that R_L and the output are different nodes. Using an active device such as a transistor with voltage feedback instead of a barefoot resistor will get you the noise temperature of the transistor instead of the resistor, while keeping the noise resistance constant.

Aside: Refrigerators. In case you're still worried about how a 300 K amp can have a 35 K noise temperature, sit down with a cold drink and consider the ice cubes in your glass—they were made in a 300 K ambient too.




Figure 18.17. The constant resistance T-coil gives a $2.8\times$ bandwidth improvement over a plain RC , with constant load resistance.

[†]Carl Battjes, “Who Wakes the Bugler?” in Jim Williams, ed., *The Art and Science of Analog Circuit Design*, Butterworth-Heinemann, Woburn, MA, 1995.

18.6 ADVANCED PHOTODIODE FRONT ENDS

18.6.1 Linear Combinations

Optical measurements are frequently based on sums and differences of the photocurrents in two or more detectors. Examples are position-sensitive detectors such as quadrant cells, as well as autofocusing in CD players, phase detection by Schlieren or Nomarski techniques, and polarimetry.

Doing this at DC is easy; almost any way you think of will work, even digital techniques (Section 17.2.5). When the measurement must be done at some speed, however, the effects of circuit strays become large enough to cause serious problems.

For example, consider trying to measure an extinction of 10^{-4} on top of a rapidly fluctuating background. A typical example is a current-tuned diode laser, whose output power varies rapidly with tuning. A common way to do this measurement is to send a fraction of the laser beam into one detector, and the rest through the sample chamber to a second detector. If the frequency band of interest is DC–1 MHz, then to maintain an accuracy of 10^{-4} in the face of an order-unity fluctuation due to scanning requires that the circuitry following the two detectors be matched to 0.01% in amplitude and 10^{-4} radian in phase at 1 MHz. These requirements push the state of the art if separate amplifiers are used, especially because you can get 10^{-4} radian of phase shift across a 10 k Ω feedback resistor by having an extra 0.0016 pF of stray capacitance on one versus the other at 1 MHz.

This book being what it is, of course, there is a circuit hack for it: just wire the photodiodes in series. With the outer ends of the diodes bypassed solidly to ground, the diodes are actually in parallel for AC and noise purposes. There is no opportunity for the strays to differ: you have one amplifier, one summing node, one ground, and one cable. Differences in diode capacitance are of no consequence, because the two capacitances are in parallel, and so both diodes see both capacitances. This trick works well with discrete devices, but unfortunately split detectors usually come with all the anodes or all the cathodes wired together, so that this is not possible. For high frequency split cell applications, transformer coupling with the DC bias applied to a center tap is a good solution.

It is possible to use the series connection with cascoding: either use a biased cascode, so that the net DC photocurrent can be positive or negative without reverse-biasing the transistor, or use a separate cascode for each diode (one will be NPN and the other PNP), with their collectors connected together and with a big capacitor between their emitters, so that the diodes are connected together at AC.

There are two difficulties with this basic approach. One is that there are slight differences between diodes that do matter. Besides shunt resistance, photodiodes have a small *series* resistance (often 50–100 Ω for fast devices, much more for lateral effect cells), which forms an *RC* transmission line with the shunt capacitance C_d . If the two diodes have slightly different series resistances, there will be a slight phase shift between the currents they produce, given identical illumination. Unlike 1 femtofarad circuit strays, this is easily trimmed out, and will stay trimmed. Figure 18.18 shows how to wire the detectors, plus one version of the R_s balancing tweak, using a loaded pot (Section 14.2.4). It could use a 10 Ω pot, but these are unreliable. The circuit also has a conveniently nonlinear adjustment, which allows 6 \times finer control in the middle of the range. Use a one-turn cermet pot and metal film resistors.

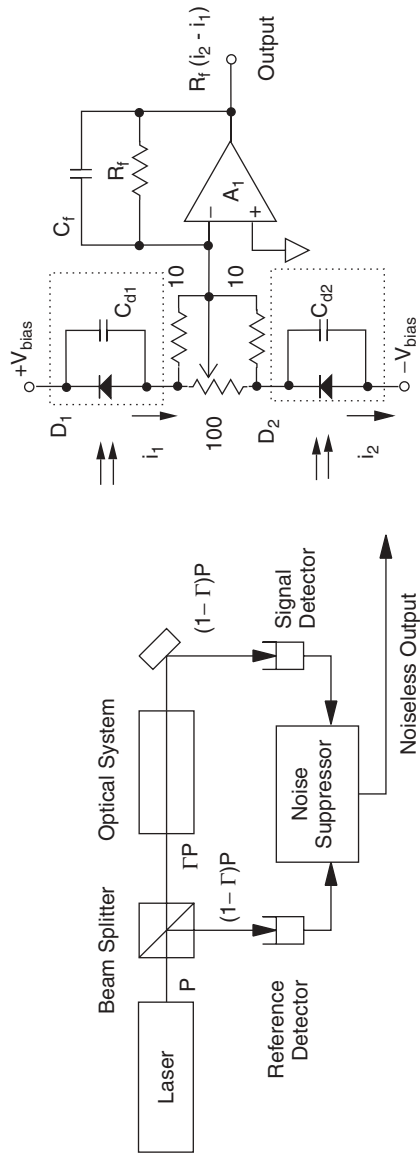


Figure 18.18. Two-beam noise suppression methods: at left, a generic two-beam system that can use subtraction or division; at right, a circuit for higher accuracy subtraction, including a tweak to correct for slightly different series resistances of the diodes.

Circuit problems can be kept at bay, but the second difficulty is more fundamental: how to keep the two photocurrents exactly equal. Can we really make the two beams identical to 1 part in 10^4 over all scanning conditions, with different sample absorptions, etalon fringes, dust, and so on? This is a nontrivial requirement. It would be convenient to have an automatic way of maintaining the adjustment.

18.6.2 Analog Dividers

One possibility is to divide one current by the other, rather than subtracting them. You do this by converting each photocurrent separately to a voltage, and then applying the two voltages to an analog divider IC, which returns a voltage indicating the ratio between the two applied voltages. Due to circuit strays and divider errors, this idea is not adequate for the demanding application above, but may be useful in lower performance situations. Its main charm is that the two photocurrents no longer have to be identical. A less obvious one is that since everything is proportional to the laser power, the signal gets intermodulated with the noise, which can be extremely obnoxious; dividers ideally fix this as well. Dividers unfortunately are too noisy and slow for most uses, and their accuracy is very seldom better than 0.5%.

18.6.3 Noise Cancelers

In principle, differential laser measurements should be totally insensitive to additive noise due to source fluctuations, because of three perfect properties:

1. With lasers, it is possible to make sure that the two detectors see *exactly* the same beam; this requires some care, for example, putting an efficient polarizer at the laser so that spontaneous emission in the polarization opposite to the laser beam does not get converted to uncorrelated amplitude noise in the two beams (VCSELs are especially bad).
2. Optical systems are very wideband (0.01 nm bandwidth in the visible is 10 GHz temporal bandwidth).
3. Optical systems and photodiodes are very linear as well.
4. Therefore, *Given two beams from the same laser hitting two photodiodes, the instantaneous excess noise current is exactly proportional to the DC photocurrent.* This is a very powerful fact, as we'll soon see.

Imagine taking a laser beam, splitting it into two carefully, without vignetting it or introducing etalon fringes, and sending one of the resulting beams (the signal beam) through your optical system into one photodiode, and the second (the comparison beam) directly to a second photodiode. Since everything is very wideband and linear, the fluctuations in the original beam split exactly as the carrier does. (The shot noise of the two beams is of course uncorrelated.) This means that if you adjust the beam intensities so that the DC photocurrents cancel exactly, the excess noise (above shot noise) cancels identically at all frequencies of interest, even far outside the control bandwidth. Twiddling an attenuator to keep this exactly true requires a graduate student to be shipped with each instrument, of course, which may reduce its practicality, but at least he doesn't have to adjust it very fast.

We're rescued by another remarkable fact:

1. A bipolar transistor differential pair is an extremely linear, voltage-controlled current splitter.

Take two BJTs, with their emitters connected together. Ground the base of one, and put some fixed voltage ΔV_{BE} , $-60 \text{ mV} \lesssim \Delta V_{BE} \lesssim 60 \text{ mV}$, on the other. Now inject some current I_{in} into the emitter node. For a fixed value of ΔV_{BE} , the ratio of the collector currents is constant over a range of several decades in emitter current, and the ratio can be adjusted over a very wide range. As a consequence, any fluctuations in I_{in} split in just the same ratio as the DC does.

Putting these five facts together with a garden-variety cascoded transimpedance amp, we can make an electronically balanced subtracter, so the grad student can adjust a pot controlling ΔV_{BE} instead of an optical attenuator. We can go a bit further, too, by noticing that the student can be replaced by an op amp, since the criterion for perfect adjustment is so simple: zero volts means zero excess noise.

We've arrived at the *laser noise canceler*, a version of which is shown in Figure 18.19. It has two outputs. The normal transimpedance output has its DC value nulled, of course, so it puts out a highpass filtered version of the signal photocurrent minus its noise. The servo signal from A_2 is a lowpass filtered ratiometric output, which depends only on the ratio of the signal and comparison photocurrents, minus both the background noise and the noise intermodulation.[†] From the Ebers–Moll equation, it's easy to show that

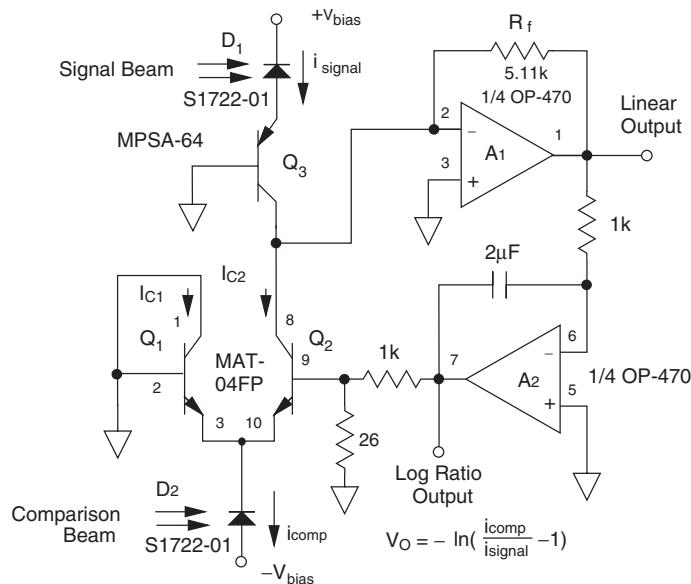


Figure 18.19. Laser noise canceler, single-ended version. BJT pair Q_1/Q_2 splits the reference photocurrent so as to null the total DC. Noise and signal are treated the same way, so the noise cancels at all frequencies. (From Hobbs, 1997.)

[†]Canceling the noise intermodulation means that (e.g., in spectroscopy) both the baseline and the peak heights are independent of laser intensity and intensity noise.

ΔV_{BE} is

$$\Delta V_{BE} = \frac{kT}{e} \ln \left(\frac{I_{\text{comp}}}{I_{\text{sig}}} - 1 \right). \quad (18.33)$$

BJT differential pairs are unique among active devices in that ΔV_{BE} depends only on the ratio of the collector currents, not on their magnitudes. This relation holds over several decades of collector current, and is why the fluctuations split exactly as the DC, which (as we saw) is the key BJT property for cancellation to work. Thus measuring ΔV_{BE} allows us to make measurements of relative attenuation even with order-unity fluctuations of laser power, a key virtue for spectroscopy, for instance.

A cardinal fact here is that the cancellation itself comes from circuit balance, not from the feedback—the feedback just establishes the conditions for the cancellation to be exact. Thus the cancellation operates at all frequencies, *completely independent of the feedback bandwidth*. What that means is that we can make shot noise limited measurements of optical power *at baseband*, even with very noisy lasers. This has very beneficial consequences for measurements, because it makes the bright field quieter than the dark field.

18.6.4 Using Noise Cancelers

Noise cancelers are simple to use, as we saw in Section 10.8.6; you take a sample of your laser beam with some etalon-fringe-free beamsplitter like a Wollaston, shove the more powerful of the two (the *comparison beam*) into the lower photodiode, and run the other *signal beam* through your optical system to the signal photodiode. It is a very good idea to put a good-quality polarizer right at the laser, because otherwise the spontaneous emission contribution doesn't split the same way as the laser light, and that disturbs the cancellation.

The exact ratio is usually not critical; a rule of thumb is to make the comparison beam $1.1 \times$ to $2 \times$ as strong as the signal beam. Choosing $2 \times$ costs more laser power but sets the operating point at $\Delta V_{BE} = 0$, where the temperature drift of the baseline is zero.

The linear output is very convenient to use, because with the comparison beam blocked, it turns into an ordinary transimpedance amp, which makes setup easy—you can adjust the aiming by maximizing the DC, and a spectrum analyzer will tell you how much cancellation you're getting. (The log ratio output, of course, rails when either beam is blocked—you get $\log(0)$ or $\log(\infty)$.) There are several useful variations of the basic circuit, including the differential model of Figure 18.20 and the fast ratio-only model, which extends the log bandwidth to several megahertz.

18.6.5 Noise Canceler Performance

This is a surprisingly powerful technique. In a measurement whose sensitivity is limited by laser residual intensity noise (RIN), the noise canceler can improve the SNR by as much as 70 decibels at low frequencies, and by 40 dB up to 8–10 MHz or so, as shown in Figure 18.21.

It will reliably reach the shot noise even with very noisy lasers; Figure 18.21 shows it getting to within 0.2 dB of the shot noise[†] with a total of 13 mW of 532 nm DPY laser

[†] Q_1 and Q_2 were matched Motorola MRF904 RF transistors, D_1 and D_2 were Hamamatsu S-1722 photodiodes. The signal beam was 5.6 mW ($I_{\text{sig}} = 1.77$ mA), and the comparison beam was 7.2 mW ($I_{\text{comp}} = 2.3$ mA).

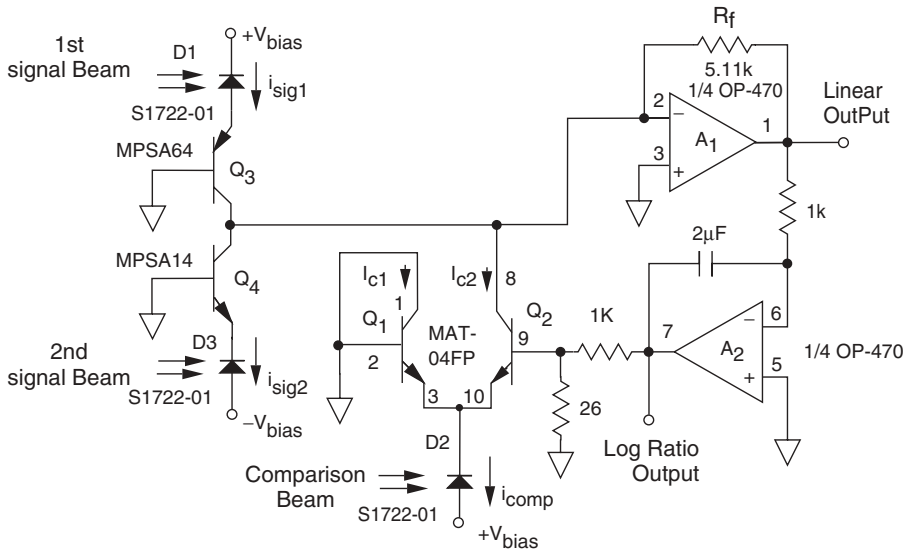


Figure 18.20. The differential noise canceler adds a second cascoded photodiode. Since most of the negative photocurrent bypasses the differential pair, the nonideal behavior of the BJTs at high currents is eliminated. This circuit can achieve 1 Hz measurement SNRs of 160 dB or more, even with lasers 70 dB noisier than that.

light, which is a tough test—154 dB dynamic range in 1 Hz. (The differential model of Figure 18.20 can do 160 dB.) The noise at the log ratio output is given by

$$eN_{\log}(I_{\text{sig}}, V_{\log}) = \frac{2kT}{\gamma\sqrt{eI_{\text{sig}}}} \left[1 + \exp\left(\frac{e\gamma V_{\log}}{kT}\right) \right], \quad (18.34)$$

which is just the total photocurrent shot noise times $\partial V_{\log}/\partial I_{\text{sig}}$ —the 1 Hz SNR at the linear and log ratio outputs is ideally the same. Figure 18.22 shows the dependence of the log ratio output noise on I_{sig} , compared to the shot noise (solid line) and the shot noise corrected for the Johnson noise of a 40 Ω base resistance $r_{B'}$. (The base resistance contribution can be reduced by paralleling transistors if necessary.) The log output’s noise is also flat with frequency; Figure 18.22 also shows the noise PSD of the highest photocurrent data point, where $I_{\text{sig}} = 931 \mu\text{A}$ —it’s flat way down into the low audio, and its <10 Hz behavior is dominated by temperature swings in Q_1 and Q_2 .

18.6.6 Multiplicative Noise Rejection

Because of the ratiometric property of ΔV_{BE} , the noise canceler’s log ratio output is the best thing going for multiplicative noise. It can cancel both additive noise and the noise intermodulation down to the shot noise level most of the time. Figure 18.23 shows at least 70 dB suppression of noise intermodulation, which is much better than any competing




Figure 18.21. Basic noise canceler ultimate noise floor at the linear output. (a) Top trace, comparison beam blocked; bottom trace, both beams on, showing a noise floor 0.15 dB above shot noise. (b) DC [–]10 MHz performance; top and bottom traces as before.







Figure 18.22. Noise of the log ratio output of the fast log version of the laser noise canceler. Allowing for a 40Ω base resistance, the noise fits the theory very well. (a) A 3–10 kHz noise versus signal photocurrent; (b) noise versus frequency at $I_{sig} = 931 \mu\text{A}$, corresponding to the bottom right data point in (a). The noise is flat down to the very low baseband. (The circuit diagram and many additional details are in Hobbs, 1997, <http://electrooptical.net/www/canceller/noisecan.pdf>.)



(a)



(b)

Figure 18.23. Two useful noise canceler variations. (a) Noise cancellation performance of the differential version, showing >70 dB cancellation of additive noise with $I_{sig1} = 1.481$ mA and $I_{sig2} = 1.36$ mA. (b) Noise intermodulation performance of the ratio-only noise canceler: upper trace, comparison beam replaced by flashlight producing the same photocurrent; lower trace, comparison beam unblocked (normal operation). Multiplicative noise is suppressed by 70 dB or more, a result superior to any previous technique. (From Hobbs, 1997.)

technique (if you look closely, you may be able to see that the sideband noise has even been returned to the carrier, as we'd hope).[†]

18.6.7 Applications

Just about any laser-based measurement that's limited by laser intensity noise can benefit. Anyone building a laser-based measurement system would do well to investigate, because the simplification of the optical and signal processing systems is usually enormous—life really is easier at baseband, and you can have that convenience along with shot noise limited performance, even in unattended measurements. It saves a lot of AO cells and wasted photons in heterodyne and FM measurements.

The author and many others have used this device to greatly simplify a number of ultrasensitive measurements, including transient extinction, tunable diode laser spectroscopy,[‡] and coherent lidar. Spectroscopy and extinction experiments use the optical system of Figure 10.10, and the coherent lidar just adds an interferometer, as in Example 1.12.

Noise cancelers are easily constructed and are commercially available.[§] *Note: This circuit is much easier to mess up than to improve.* If you're building your own for the first time, build the basic model exactly as shown (with capacitance multipliers and a few bypasses on the supplies, and perhaps different photodiodes, e.g., BPW34s, which have $C_d \approx 10$ pF @ 25 V) and see how it works before changing it. Seemingly small changes, such as switching to CMOS op amps, can make a profound difference. Especially resist the temptation to put the photodiodes on cables—1-inch leads at most.

18.6.8 Limitations

The most serious limitation of the canceler is the deviations of the transistors from ideal behavior (principally the parasitic series resistances of the emitter and base, $R_{E'}$ and $R_{B'}$, respectively). This can be got round by using the differential model, in which only a small fraction of the photocurrent has to go through the differential pair.

From an optical point of view, the noise canceler's biggest liability is its own strength—after it cancels the big ugly correlated noise, it shows up all the second-order warts on your beams. Cancellation is hindered by anything that decorrelates the noise: vignetting, etalon fringes, and spontaneous emission in the polarization orthogonal to the laser light.

These can usually be fixed easily enough, but finding them does require some care and thought. There isn't space here to go into all of its ins and outs, but if laser intensity noise is a problem for you, check the referenced articles.

Because of all the fine points a noise canceler will show you, it takes a little while to get up to speed with it—the average seems to be about 2 weeks—but the investment pays off over and over. Learning to spot those etalon fringes, vignetted beams, and coherence fluctuations will make all your systems work better, whether they have noise cancelers

[†]P. C. D. Hobbs, Reaching the shot noise limit for \$10. *Optics and Photonics News*, April 1991; and Ultrasensitive laser measurements without tears, *Appl. Opt.* **36**(4), 903–920 (February 1, 1997).

[‡]K. L. Haller and P. C. D. Hobbs, Tunable diode laser spectroscopy with a novel all-electronic noise canceller. *SPIE Proc.* **1435** (1991). Available at <http://electrooptical.net/www/canceller/iodine.pdf>.

[§]Available as the Nirvana detector from New Focus, Inc.

or not. One key piece of advice: there are lots of things that appear to be common mode but aren't, for example, etalon fringes in two different polarizations.

18.7 OTHER TYPES OF FRONT END

18.7.1 Really Low Level Photodiode Amplifiers

We've spent almost all of our time in this chapter fighting to stay at the shot noise limit, but sometimes that just isn't possible. For example, if our $2\ \mu\text{A}$ photocurrent were 50 pA instead, we'd need a $1\ \text{G}\Omega$ load resistor to get within 3 dB of the shot noise, and only physicists use resistors that large.[†] What's more, in a DC measurement, any significant reverse bias will corrupt the data with leakage current. What we'd like to do then is go to an AC measurement, but that may involve choppers and so forth, which are bulky, expensive, and unreliable.

If we have to do a DC measurement, with no opportunity to measure the leakage current independently, we're stuck with operating at exactly zero bias to make the leakage zero. As we saw in Section 14.6.1, the small-signal resistance of a zero-biased photodiode can be quite low—most are nowhere near $1\ \text{G}\Omega$ even at room temperature, and all drop by half (or even a bit further) every $10\ ^\circ\text{C}$. This resistance appears in shunt with the photodiode. It contributes Johnson noise current, of course, but what's more, it increases the noise gain of the stage, in much the same way that the photodiode capacitance does, except that being a resistor it does it at all frequencies. Keep the photodiode small, the load resistor no larger than necessary to override the amplifier noise, and consider cooling the diode.

Aside: Heroic Efforts. Some specially selected diodes can reach 1 or even $50\ \text{G}\Omega$ at $20\ ^\circ\text{C}$, and really careful work can get these down to a few hundred electrons/s worth of noise in millihertz bandwidths, if you can wait long enough.[‡]

18.7.2 Pyroelectric Front Ends

Pyroelectric detectors are difficult devices to interface to, since they convert a temperature change into a charge, rather than a current as quantum detectors and bolometers do. That means that at low frequencies, the current available is proportional to the time derivative of the sensor temperature, which is inconvenient. Example 17.1 shows one way to solve that problem; here we're concerned with keeping as much SNR as we can, which means high stability and femtoamp leakage.

From the front end's point of view, the trouble with pyroelectrics is their low signal level and very high impedance. The most familiar pyroelectrics, namely, porch light sensors, make their AC signal by using a segmented Fresnel lens that casts a dozen or so images on a split detector. The two are wired in opposing parallel, so when you walk up to the door, twelve of you in a row cross from the + half to the - half, generating a nice AC signal. The two pyros are connected between the gate of a discrete MOSFET and ground, with a $10\ \text{M}\Omega$ leak resistor to keep the DC level constant. The FET's drain

[†]Well, electrical engineers use them once in a while, but it takes a physicist to put one on a cable.

[‡]G. Eppeldauer and J. E. Hardis, Fourteen decade photocurrent measurements with large-area silicon photodiodes at room temperature. *Appl. Optics* **30** (22), 3091–3099 (1991).

is AC coupled to the thresholding circuit. It's a simple and elegant idea, which is why they sell tens of millions of them. The problem with this for our purposes is that it's very noisy, far too noisy for an imaging sensor. The leak resistor reduces the signal level and adds a lot of Johnson noise, the discrete MOSFET isn't too quiet at low frequencies, and the thermal drift is bad enough to make your neighbor's porch light come on whenever there's a gust of wind.

The good news is that the pyroelectric pixel itself is quite a good capacitor, so we can use a charge-dispensing readout reminiscent of a CCD. The reason this is a good idea is that you can let the pixel integrate itself for a whole frame time, then dump all the collected charge in one pulse, right when you want to measure it. This has exactly the same nice SNR consequences as the pulsed measurements of Section 13.8.10. Integrated pyroelectrics (e.g., those from Irisys) usually stack the pyro on top of a CMOS readout chip, which makes all the decisions for you. A built-up circuit has to do its own charge dispensing.

The Footprints sensor's multiplexer uses diode switches made from ordinary display LEDs.[†] Ordinary display LEDs have extraordinarily low leakage. One snag is that (being differentiators) pyroelectrics produce a bipolar current, and diodes conduct only in one direction. The basic idea is to put the switch LEDs under an opaque white cover and illuminate them all with a processor-throttled LED, so as to produce a well-behaved bias current of a picoamp or two. Figure 18.24 shows the multiplexer design. Each time one of the column strobe logic lines goes low, six pixels are read out at once, and digitized in succession, which takes about $300\ \mu\text{s}$ out of a frame time of 200 ms. The RC time

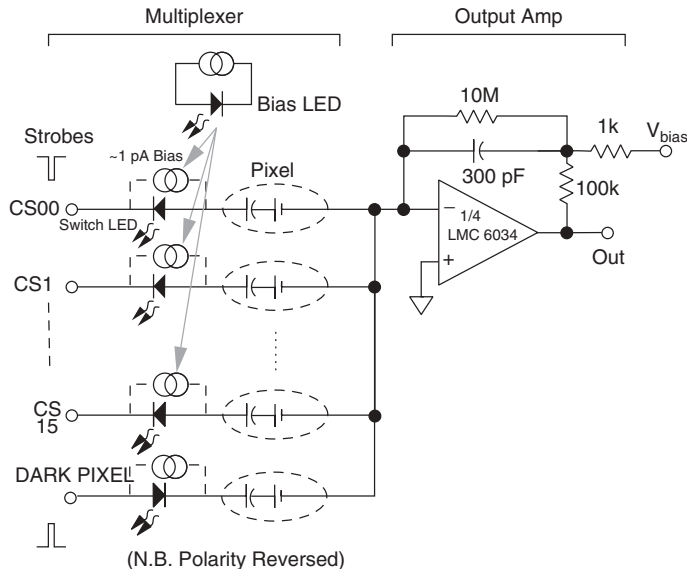


Figure 18.24. Current-dispensing multiplexer for 16 pixels of a 96 pixel pyroelectric sensor: each of the column strobes at left dumps one pixel into each of six charge-sensitive amplifiers with resistive resets. The duty cycle of the pulsed charge measurement is less than 0.2%.

[†]For many more details, see Section 13.11.16 and the papers referenced there, Example 17.1, and Section 14.6.1.

constant of the amplifier is 5 ms. Because we don't get to dispense every single electron as a CCD does, both the dispensing and reset operations have kTC noise, so correlated double sampling doesn't actually help much here.

Example 17.1 explains some signal processing tricks needed to fix up the transfer function, but when it's done, a 96 pixel sensor costing \$10 (including the lens) can give quite competitive sensitivity: 0.13 K NE Δ T, with room for probably 10 \times further improvement.

18.7.3 IR Photodiode Front Ends

Near-infrared photodiodes (InGaAs and germanium) work more or less the same way as silicon ones, because their shunt impedance is high, so that they are current sources to a reasonable approximation. The main addition is that they have significant amounts of dark current, which exhibits full shot noise. Provided that the photocurrent is large enough to dominate the dark current, this is not a limitation. The dark current is a strong function of the bias voltage, so with dim light, it may be necessary to run these devices at much lower reverse bias. This means higher capacitance.

Mid- and far-IR photodiodes are a considerably more difficult problem. The Judson InAs detector discussed in Example 3.2 had a shunt resistance of 100 Ω even at -40°C , dropping to 10 Ω at room temperature—you can't reverse bias that by very much. Detectors with such low shunt resistances are limited by their own Johnson noise, except at extremely high illumination levels. The task of the front end designer is to make the best of this, because improving the detector is usually expensive, impractical, or impossible. Cryogenically cooled far-IR detectors are frequently limited by the shot noise of the thermal background photons, which is also not susceptible to circuit improvements.

Generally, it is difficult to make amplifiers whose noise is significantly (say, 15 dB) below the Johnson noise of such a low value resistor. The problem is usually voltage noise rather than current noise, because the Johnson noise current is so large at these resistances. There are two techniques that work reasonably well: transformer coupling and very quiet bipolar current amplifiers. At high frequency, reactive matching networks are a third.

18.7.4 Transformer Coupling

If the source impedance is a poor match for any available amplifier, why not change it? Of course, the source impedance could be increased by wiring a resistor in series with it, but that would be a strange way to improve the noise performance. Instead, an impedance transforming network is used. For high frequency, narrowband (1 octave or less), an LC matching network is usually best. The same idea can be applied at low frequency or wide bandwidth as well, using a transformer.

A good transformer has strong mutual coupling between its windings and low losses to ohmic heating in the copper and hysteresis or eddy currents in the magnetic core material (usually powdered iron or ferrite, sometimes permalloy). This means that nearly all of the available power from the primary is available at the secondary; furthermore, by the fluctuation–dissipation theorem, low losses mean low added thermal noise.

The reason this is useful is that if we wind a transformer with an N -turn primary and M -turn secondary winding, the voltage at the secondary is $K = M/N$ times the primary

voltage, and the current is $1/K$ times. Thus the impedance has been transformed by a factor of K^2 . An amplifier whose 1 Hz voltage and current noises are 1 nV and 4 pA will be quite good (total noise 1 dB above the detector noise) with a 250 Ω source, but poor (8.5 dB over detector noise) at 10 Ω . Since the increased Johnson noise makes the 10 Ω detector 14 dB noisier to begin with, this is really adding insult to injury. A transformer can make the 10 Ω detector look like 250 Ω to the amplifier, eliminating the additional 7.5 dB SNR loss (though we're still stuck with the 14 dB).

Another advantage of transformer coupling with low shunt resistance detectors is that the DC voltage across the detector is held at zero, because there is a wire connected all the way from one side to the other. There are two main disadvantages: you can't tell what the DC photocurrent is, because there is no DC connection between the amplifier and the detector, and there is no simple way to reduce the intrinsic RC time constant of the detector except by reducing the load resistance, which seriously degrades the noise performance. The first you can fix with some circuit hacks, but the second you're stuck with. Good transformers are available from EG&G PARC, Jensen Transformer, and Mini-Circuits Labs.

18.8 HINTS

These maxims will help keep you out of the worst potholes in front end design. If you ignore any of these, make sure you know why you're doing it.

One dB Matters. A loss of 1 dB in the SNR requires 26% more signal power to overcome it. In a photon-limited system, this can add that 40% to the cost of the optics, or stretch the measurement time by 26%. These factors of 1.26 multiply, so that if the loss is more than 1 dB, life gets a lot worse, fast. This is an absolute, inescapable, information-theoretic limit and cannot be got round by any postprocessing whatever. Put lots of effort into getting your detector subsystem really right; you'll be grateful later, when the measurement is fast and the data are good. Even if you *are* building a spy satellite or solar telescope, where photons are not the problem, make the detector subsystem right anyway. It's good for the soul, builds your expertise, and anyway you're liable to reuse it another time.

Dynamic Range Is Precious. Many measurements must operate over a wide range of optical powers. It is obnoxious to be forced to choose between railing your amplifier on the peaks, or having the troughs disappear into the noise. Don't use 3 V or 5 V supplies in high dynamic range applications. You're throwing away as much as 20 dB of dynamic range, compared with a ± 15 V system. After the dynamic range has been reduced, for example, by filtering out the DC background, this is usually much less of a problem, so the amount of circuitry requiring ± 15 V is usually small. This makes it feasible to power the front end of a mostly 5 V system with a small DC-to-DC converter. You can use charge-pump voltage converters such as the ICL7660 and its descendants, or a small switching regulator. Use fully shielded inductors, and don't omit to filter the output of these devices with a three-terminal regulator or, better, a capacitance multiplier. Bypass capacitors won't do it. Watch out for inductive pickup from switching regulators, and for the fuzz that any of these sorts of devices always puts on its input supply.

Always Plot the SNR. It is depressing how many people ignore how the SNR changes with frequency. In this chapter, we've seen that there are lots of counterintuitive things SNRs can do, so don't omit to calculate what SNR you expect. Sometimes a slower front end with a peaking filter in a subsequent stage to compensate for its rolloff can work just as well as a gold-plated ultrafast front end.

Always Measure the Noise Floor. In Section 1.7, we talked about making sure that the photon budget was met, and not being satisfied with less than full theoretical performance. The noise floor of the front end amplifier is one place that people never seem to expect good results, and often don't even measure, even though it's trivially easy—a flashlight will produce a photocurrent with exactly full shot noise; find out what photocurrent gives a 3 dB noise increase, and you know the input-referred noise. (This works independently of gain, measurement bandwidth, and so on, but don't try to do it on a scope by eye—use a spectrum analyzer or a filter plus an AC voltmeter, see Sections 2.5.4 and 13.6.5.) *Be exhorted:* you really can predict the noise floor accurately—to accept a noisy front end is one of the stupidest and most expensive mistakes you can make in designing sensitive optical instruments. Measure it, and make sure you can explain every half decibel.

Don't Use 50 Ω Unless You're Driven to It. Amplifiers with 50 Ω inputs are all over the place, but they shouldn't be in your front end—unless there's a reactive matching network in front of them, or your photocurrent is at least 1 mA. Long haul fiber optic communications people use a lot of 50 Ω amplifiers, but they struggle for every fraction of a decibel, so that lets them off the hook.

Provide a DC Output. It is very useful to provide a DC output from a detector, for setup, alignment, and troubleshooting. If there's too much gain to allow straight DC coupling without railing an amplifier somewhere, make the DC gain lower or send the DC to an auxiliary output—just don't get rid of it altogether, or you'll wish you hadn't.[†]

Use Feedback Tee Networks. We're accustomed to ignoring the noise of the second and subsequent stages of an amplifier chain, and this is fine as long as the front end has high enough gain. A transimpedance amplifier has a noise gain of 1 (for noise other than $e_{N_{amp}}$), and capacitance limits how big we can make R_f , so the second stage noise can easily dominate if we're not careful.

Use a quiet amplifier for the second stage, or put a tee network in the feedback loop of the transimpedance amplifier, as shown in Figure 18.25. This network increases Z_m and A_{VCL} by reducing H_{fb} , without having to increase R_f and so suffer extra phase shift. Of course, the bandwidth will be reduced by the voltage divider ratio of the tee network as well, so a faster amplifier will be needed. Some people like to put two amplifiers inside the same high gain feedback loop, to get extra bandwidth and eliminate the second-stage noise and input errors. If you do this, the booster stage needs its own local feedback to ensure it runs at a fixed AC gain, and must be fast enough not to mess up the overall loop stability.

[†]For differential detectors, it is nice but not essential to bring out both ends as well as the difference (perhaps using current mirrors to bring the voltages down near ground).

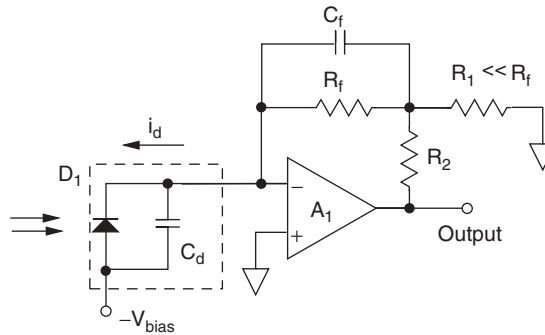


Figure 18.25. A tee network in the feedback loop of a transimpedance amp provides extra voltage gain at the expense of loop bandwidth. The increased signal gain reduces the effects of second-stage noise. Don't reduce R_F .

The value of C_f is not changed by the addition of the tee network. The parallel combination $R_{\text{div}} = R_1 || R_2$ must be small enough so that $1/(2\pi R_{\text{div}} C_f) \gg f_{3\text{dB}}$.

Even without the resistive divider, a tee network made up of small capacitors can allow the use of a somewhat larger C_f if the calculated value is inconveniently small. Don't get carried away, though: since this forces one end of C_f to be essentially ground, C_f then loads the summing junction. Once it gets up to 2 pF or so, don't go any further, or you'll make the instability worse rather than better (consider what would happen if you replaced a 1 pF C_f with a 1:10⁴ capacitive divider and a 10 nF C_f).

Don't Put Photodiodes on Cables. Optical systems are often large and operate under stiff constraints on weight, cost, and complexity. It is therefore tempting to allow the light to come out anywhere it wants, and put the photodiode there. This is reasonable as far as it goes; you can put the front end amplifier there too. Unfortunately, people often just hang the bare photodiode on an RG-58 cable and connect an amplifier (50 Ω , you guessed it) to the other end. This is a ticket to perdition. That cable will pick up signals from everywhere, including ground, FM radio, lightning, you name it. When unmatched, it will exhibit huge capacitances (100 pF/m) at low frequencies, and poorly controlled transmission resonances and phase delays at higher frequencies. If there's a DC voltage on it (as there usually will be with photodiodes), cable vibrations produce capacitance changes that show up as signal. The list goes on and on. Especially when you're trying to do differential measurements, and *especially* with noise cancelers, keep the amplifier and the photodiode together.

Put Capacitance Multipliers on the Supplies. We talked about the virtues of capacitance multipliers in Example 14.1; they have poor regulation near DC where that's OK, and unsurpassed regulation at AC, where it really counts, because the supply rejection of your amplifiers is poor and your switching power supplies very noisy. Front ends are an excellent place for a capacitance multiplier.

Always Build a Prototype and Bang on It. It is not possible to build a first class front end with nothing but SPICE and a PC board layout package. This subsystem absolutely must be prototyped, and the prototype's characteristics measured to within a

gnat's eyebrow to make sure that you understand where all the noise is coming from. If its noise performance at your expected minimum photocurrent is not within a couple of tenths of a decibel of what you expected, stop and find out why. A certain healthy paranoia is indicated here.

The other reason is that circuit strays are very important. The transimpedance amp design we wound up with used an LF357 with a 300 k Ω feedback resistor and a 0.8 pF feedback capacitor. Without the capacitor, its phase margin was negative—it would have oscillated at about 1 MHz, depending slightly on where the second pole fell. Increasing the capacitor will seriously degrade its bandwidth. Ordinary metal film $\frac{1}{8}$ W axial-lead resistors have a capacitance of about 0.25 pF, and surface mount ones less than that, so such a small feedback capacitance is possible. In fact, it is often possible to build this capacitance right into the board layout, for example, by putting a ring of copper around the inverting input, connected to the output pin (it may need to be AC coupled to avoid leakage). SPICE won't be much help in making the board layout right, even if you have a trustworthy model of how your cascode transistor behaves at 5 μ A of collector current, which you probably haven't.

Make sure you follow Pease's Principle[†]: bang on it. Stick a square wave through a big resistor into the summing junction, then into the + input, looking for overshoot (you have to put a small resistor in series with the + input first, of course). If the overshoot is more than 20% of the step height, C_f is probably too small. Finally, bang on the output with a square wave through a low value resistor. Do this at various frequencies, too—sometimes it looks different.

Center Your Design. Component variations are one of the major causes of manufacturing yield problems in analog electronic systems. You can't possibly build enough prototypes to take in the whole range of all components, so use simulation. Most flavors of SPICE can do Monte Carlo sampling of the normal variation in each component, or you can write your own code to do it, with a compiler, a spreadsheet, or a scratchpad program such as MathCad, GNU Octave, or Matlab.

Pick component values that lead to acceptable performance over all the cases. Every last component in the circuit has limits on each of its parameters, beyond which the circuit will not function well enough. In a landscape full of highly multidimensional cliffs, we're almost bound to be near one of them. Simulation will help find it, and tell us how far to move in what direction to be equidistant between cliffs. This is called *centering*, and it will save you lots of headaches. Beware, though, that there are cliffs lurking in the simulation itself: models and model parameters are all lies. Some of them are just more useful than others. Make sure that you check the centering experimentally, by changing the values and seeing where trouble develops.

RF Amplifiers' Noise Figures Depend on Source Reactance. Every RF device has an optimum source impedance, where its noise figure is best. This is generally *not* the matched condition. Amplifiers therefore have noise performance that depends on the impedance mismatch at their inputs, which is a matter of critical concern in high frequency front ends. Make sure that your amplifier is a type that works well with horribly reactive input impedances, and that it is cannot oscillate for any value of source impedance (i.e., it must be *unconditionally stable*).

[†]Robert A. Pease, *Troubleshooting Analog Circuits*. Butterworth-Heinemann, Woburn, MA, 1991.

Bringing Up the System

...it is very easy to be blinded to the essential uselessness of them by the sense of achievement you get from getting them to work at all. In other words—and this is the rock-solid foundation on which the [Sirius Cybernetics] Corporation's galaxy-wide success is based—their fundamental design flaws are completely hidden by their superficial design flaws.

—Douglas Adams, *So Long, and Thanks for All the Fish*[†]

19.1 INTRODUCTION

The design work is finished, two-thirds of the electronics seem to run, and though the optical prototype has quite a bit of hot glue on it and lots of scars from engineering changes, the system is over the hump. It's about time to make sure that all the bits connect together properly and actually work in concert. Since the technical risk has been taken care of in the early prototypes, the rest should be fairly routine. The development schedule hasn't slipped much, so from now on it should be 40 hour weeks and lots of time with the kids. System integration and debug is just a matter of turning the crank and fixing small oversights, right?

Usually not. System integration is the time when the really weird bugs show up, the ones that are not reproducible nor susceptible to simulations. It is not that there are evil spirits infesting the subsystems, or that the basic physics is not understood—it's just that it is remarkably difficult to simulate an effect you haven't thought about. Here's a war story.

19.1.1 The Particle Counter That Wouldn't

A group of the author's friends transferred the design for a liquidborne particle counter[‡] to a small vendor. The sensor was to be used to detect submicron foreign particles in fluids (water and aqueous HF primarily). The vendor had previously made dark-field instruments that weren't too complicated. This one was a bright-field Nomarski-type interferometric system with a lot of electronics and software.

[†]Douglas Adams, *So Long, and Thanks for All the Fish*. Harmony Books, New York, 1985.

[‡]J. S. Batchelder and M. A. Taubenblatt, Interferometric detection of forward scattered light from small particles. *Appl. Phys. Lett.* **55**(3), 215–217 (July 1989).

The system was complex because it was specially designed to reject spurious counts due to bubbles, the bane of other techniques. Our group had built a few demonstration systems, which successfully operated in the semiconductor manufacturing line, but which were the size of a projection TV, weren't easy to align, and had too many expensive commercial mounts—not a very saleable package. (Still, having a working prototype to start from is pretty comforting.) The vendor had a good optical guy, who undertook to cost-reduce the design for production. Besides designing a package of manageable size and substituting a diode for the HeNe laser, he changed the interferometer by replacing the Nomarski wedges at the pupils of the objectives with Savart plates in front of them.[†] Because of the violently corrosive fluids, the flow cell windows had to be made of sapphire, and to aid servicing, the (very thin) cell had to be removable—a stiff requirement for a moderate-NA system. In addition, they redesigned all the electronics.

Optical particle counters don't really measure the size of particles, but rather their scattering or absorption cross sections. These are converted into fictitious particle diameters by reference to a curve of cross section versus diameter for polystyrene latex spheres, which are a good, commonly available standard. In order to get this cross section from a light pulse, we have to know the intensity of the light incident on the particle. Since optical beams usually have intensity profiles that vary smoothly between the center and the wings, that means we need to know the particle's radial position.

As shown in Figure 19.1, the Nomarski-type interferometer (see Example 10.2) produces two overlapping, orthogonally polarized focused spots, whose phase is adjusted to near 90° in quiescent conditions, so that there is ideally no background signal. Due to the phase shift from the particle, the interferometer generates a positive voltage when the particle is in one of the two spots, and a negative voltage when it is in the other. The particle counter took advantage of this to measure the particle trajectory. By tilting the axis joining the two spots, a particle that crossed closer to the positive beam was made to produce a pulse with a big positive lobe and a weaker negative lobe (and conversely). The signal processing strategy involved measuring the asymmetry of the S-shaped signal pulse, in order to determine where in the beam the particle had passed, so that the extinction and phase shift measurements could be normalized correctly. The details of the pulse shape depended sensitively on the optical system, the laser beam quality, the aberrations in the flow cell window and Savart plates, the state of focus, the impulse response of the amplifiers and filters, and the design of the track/holds in the digitizer, among other things. It was not easy to test these subsystems realistically, especially because the data available from the early prototype was inapplicable due to the large number of engineering changes. In consequence, a great deal was left until the system integration phase.

The first problem was laser mode hopping due to optical feedback. In a dark-field system, we don't usually mind mode hopping much, since the noise is multiplicative—if there's no signal, there's no noise, and a 1% laser fluctuation can only produce a 1% signal level error. In a bright-field system, on the other hand, a 1% fluctuation may completely mask all signals of interest (see Chapter 10 for other ways around this). Although the differential system is nominally zero background due to the subtraction of the two signals, enough of the mode hopping got through to cause serious trouble. This was eventually solved by destroying the temporal coherence of the laser by large

[†]A Savart plate is a symmetrical walkoff plate that shears the two linear polarizations with respect to each other without angular deviation (see Section 6.6.2).




Figure 19.1. Nomarski-type interferometric liquid particle sensor. The Wollaston is oriented at 45° to the Nomarski prisms, so that the two beams are made to interfere losslessly while keeping the background zero.

injection current modulation at UHF (see Section 2.14.3). The coherence length dropped to a few hundred microns, and the mode hopping stopped completely.

If the system had still used Nomarski prisms instead of Savart plates, this would have been a complete solution to the optical part of the problem. Nomarski prisms have very nearly constant phase delay with angle, and at the pupil there is only one plane wave component to worry about. Unfortunately, the delay of a Savart plate is a strong function of field angle and polarization. Besides aberrating the beams, this property interacted with the short coherence length of the UHF modulated diode laser to produce a disaster. We saw in Section 2.5.3 that the relative phase fluctuations of the beams become comparable to a radian at a path difference equal to the coherence length; at 1% of the coherence length, they amount to 0.01 radian. The low frequency fluctuations of the relative phase now became the dominant noise contribution. This was far from simple to solve, eventually requiring a special nonuniform wave plate to correct the relative delay (by this stage, it was no longer trivial to go back to Nomarski prisms).

The flow cell windows were made of *c*-axis normal sapphire (see Section 4.6.1), which looks isotropic near normal incidence. Because the beam was converging, though, this wasn't good enough—window birefringence changed the polarization and state of focus of the marginal rays. This led to a peculiar pulse shape.

The electronics presented other knotty problems, for example, severe thermal drift in a nonlinear amplifier which was implemented with diodes in the feedback loop of an op amp. They were hip-deep in problems. All of them were eventually solved, and the system went on to win a prestigious industry award and get lots of good press, but only

after an entire year's schedule slippage and at very serious human cost to the participants. The product line was shortly sold off, and it went out of production soon after.

The force of this tale is that none of their individual design decisions was silly at all—the engineers were a talented bunch of guys, who did most of it right. What killed them was underestimating the difficulty of the system integration and fighting on too many fronts at once.

19.2 AVOIDING CATASTROPHE

19.2.1 Incremental Development

Hindsight is of course an exact science, and in the particle counter case it shows that the best way to proceed would have been to make a much closer copy of the system that was known to work, and to try out each change individually. The system would have incrementally converged on a workable configuration, and the devastating interaction of the optical design changes would have been avoided. Bringing to market a somewhat more expensive but less risky design would have generated revenue, which would have helped to fund the cost reduction steps and so reduce the time pressure.

Horror stories like this can be avoided, once we know what to look for, but most of us have to learn by making the mistake. All practicing electro-optical instrument designers should have the following military maxim taped to their bathroom mirror: “A pint of sweat saves a gallon of blood.” Inadequately tested design changes are extremely risky, and it is very easy to underestimate the difficulty when making a change to technologies at the limit of our competence. Very often, such a change seems an attractive way out of our current difficulties; fiber optics and diode lasers especially seem to exert a fatal attraction on designers unfamiliar with their less attractive features (see Section 2.12 and most of Chapter 8). The author once spent a surprisingly long while simply changing from a diode-pumped green Nd:YAG laser to a high power, single longitudinal mode diode laser in an ultrasensitive interferometric system. It took almost a year, off and on, during which time he learned a great deal more about diode lasers than he cared to. It was fortunate that the rest of the system was already working; if the diode had been in the system from the beginning, the design would in all probability have failed altogether. It was only the incremental character of the change that saved it.

19.2.2 Greedy Optimization

Complicated systems that work are invariably found to have developed from simpler systems that worked.

—Attributed to John von Neumann

Concurrent engineering was a 1990s management fad, but one with some basis: it would be the perfect development strategy if systems engineers were clairvoyant. It extols the virtues of getting a whole lot of people working on a project right from the beginning, with everything being developed at the same time, in order to get a rapid time to market. Since this may involve documenting nonexistent features or designing signal processing for a front end that is misconceived, it is somewhat fraught and requires a durable sense of humor.

This apparently chaotic process is actually very well suited to electro-optical instruments, because it tends to prevent the sorts of disasters recounted earlier. Properly done, engineering of electro-optical systems starts with a conceptual design (see Section 1.1), whose technical risks are carefully assessed and then reduced by breadboarding the hard parts. It proceeds to a detailed system definition and from there, immediately to a brassboard prototype.[†] This prototype is not pretty to look at, but it does allow every major design decision to be sanity-checked at a very early stage. As subsequent versions of the major subsystems become available, they are immediately tested in concert with the rest of the system, so that integration proceeds in easier steps along the way, and the ability of unconsidered effects to cause catastrophes is greatly reduced. Building prototype hardware also gets everybody asking the right questions—practical ones—and highlights the hard problems. The crucial discipline is, *don't break the prototype*. Ever. Each version change must leave the system working.

Numerical analysts call this strategy *greedy optimization*. We start with a guess and continue to accept each step that makes things better. Numerical greedy algorithms work well with a good enough initial guess, but with a poor one, they can get caught in local minima, or get stuck on a slope that never ends. The incremental engineering model is similar, in that if the initial brassboard is too far from the final system configuration, convergence on an acceptable system may not occur in the time available. Thus it is very important to build a prototype that includes all the fundamental features of the final system. If the final system is to be miniaturized, for example, a breadboard covering a 4 by 10 foot optical table is too far off to be safe.

19.2.3 Specifying the Interfaces

In order to improve our chances of convergence, the confusion must be minimized. Specify all the interfaces in the system in as much detail as possible, in advance, and keep these specifications up to date as the project progresses. As far as possible, subsystems should be independent of each other, so that as long as the interfaces are the same, internal changes to one subsystem do not affect others. Modularity is easiest to achieve in software, followed by electronic hardware, and is most difficult in the optics. This difficulty makes it critically important that a realistic brassboard of the optics be used, starting as early as possible.

19.2.4 Talking to Each Other

Even with incremental development and the use of a brassboard, there is always a very lively possibility that some unnoticed flaw will render the entire system a failure. Reducing that risk is one reason for making sure that the designers and systems engineers talk together a lot, individually and in groups. Get everybody thinking about it by making sure each designer always has an up-to-date copy of the detailed draft system specification, so that oversights can be corrected as early as possible, and that each revision highlights any changes so they're hard to miss.

[†]Prototypes are named differently, depending on how closely they approximate the final system. In order of increasing sophistication, it goes: mockup/demo, breadboard, brassboard, preproduction.

19.2.5 Rigorous Subsystem Tests

Localize each problem as closely as possible, as early as possible. It is easy to design a system that is the electro-optical equivalent of spaghetti code, or (even worse) of series-string Christmas lights: when one goes out, they all go out. The way to avoid this is by hierarchical design, dividing the function into independently testable subsystems, each of which has its own independently testable blocks. Really work hard at devising good subsystem tests.

19.2.6 Plan the Integration Phase Early

These interface specifications should lead naturally to a plan for system integration; group tightly coupled subsystems together, and test these combinations by themselves, then assemble those few groups into the complete system. This is the same strategy of divide and conquer that is used in troubleshooting: localize the problem as quickly as possible.

A planned system integration is not a substitute for incremental development. If it is used that way, it is nothing but a way of heading for trouble neatly. However, eventually the last prototype version of a system must be left behind, and the jump made to an early production version. Problems must be expected to surface here as well, and this is where the integration plan comes in. Bring up the lowest level of function first—power supplies, computers, light sources. Pay particular attention to the ways in which your subsystem test conditions differ from the actual system environment (e.g., using separate lab power supplies for each electronic board and running the optical system with the covers off). The key to (relatively) painless system integration is to allow lots of time for it, and really watch any places where this system differs significantly from others you've built in the past—removing on-card voltage regulators, switching to diode lasers, changing from dark to bright field (or vice versa), or just building a more complicated system than you're comfortable with can lead to integration problems.

19.2.7 Don't Ship It Till It's Ready

With all we hear nowadays about how time to market is the key to profitability, it is easy to get buffaloed into shipping a system that isn't ready. Missing ship dates is embarrassing; customers and managers get upset; we feel we've let them down by being too slow. Why not ship a couple just to reduce the heat, so we can finish our jobs? They can always be brought up to the final revision level later, after all.

Don't fall for this one. The first customer shipment is a Rubicon that cannot be uncrossed. If you ship an inadequately debugged system, your previously irate customer will now be apoplectic, and so your management will be scared. Scared management does stupid things, such as pouring all available engineering time into customer support, starving the debugging effort, which drags out. Where's your rapid time to market now? Meanwhile, your system is rapidly acquiring an evil reputation among its prospective customers. It is possible to recover from such a disaster, but it is most uncommon (see Figure 19.2).

The one exception to this rule is omitting features. If the basic unit is really solid, but there are one or two features that do not work properly just yet, it is quite permissible to leave them out for the first production run. This applies to the particle counter example, where the omitted features would have been some of the cost-reduction steps. Provided




Figure 19.2. Profit versus time to market: hype and reality.

the customer knows just what he will be getting, this can be a valuable source of flexibility in development.

19.3 DEBUGGING AND TROUBLESHOOTING

Troubleshooting and debugging are related but not identical. Troubleshooting is when we're finding out what's wrong with a single device that once worked, or which is supposedly identical to another 23 devices that worked (for arbitrary values of 23). Debugging is getting it to work the first time, or chasing a problem that shows up intermittently in production units.

They can be a blessing or a curse. The blessing comes when the chasing makes us think about the design in a new way, leading to improvements or even to a really new idea. It is also a wonderful intellectual discipline, which designers often need. The curse is when we're unprepared or haven't budgeted time enough.

It's really worthwhile getting good at troubleshooting. Jim Williams, a famous analog circuit designer at Linear Technology, has written an important chapter on fixing things, which says in part:

The inside of a broken, but well-designed piece of test equipment is an extraordinarily effective classroom. The age or purpose of the instrument is a minor concern. . . . Good design is independent of technology and basically timeless. The clever, elegant, and often interdisciplinary approaches found in many instruments are eye-opening, and frequently directly applicable to your own design work. . . . The specific circuit tricks you see are certainly adaptable and useful, but not nearly as valuable as studying the thought process that produced them.[†]

Design work has no fixed cutoff point; it's done when the designer says it's done, which is an opportunity for leaving loose ends and unexamined assumptions. There is no

[†]Jim Williams, ed., *The Art And Science of Analog Circuit Design*. Butterworth-Heinemann, Woburn, MA, 1995, p. 5.

such sloppiness in troubleshooting. Troubleshooting is sometimes thought to be an art, but really it's an applied science. The trouble exists out there somewhere to be found, rather than being created for the occasion; once the bug is successfully found and fixed, everyone can agree that it's gone. (By this criterion design is an art, and debugging has some qualities of both.) All sciences have their own characteristic mindset, and the mindset of a good troubleshooter is a suspicious and very observant one. If you notice anything funny, never, never ignore it just because it doesn't fit with anything else you're working at. Write it down prominently in your notebook (you do keep a notebook, right?) and think about it later.

Debugging is like troubleshooting, only working without a net. This thing has never yet worked right; it's just your confidence that it ought to, and will eventually, that you have to go on. This means that there's no single right answer, and that you can't assume that the design is basically correct. Debugging often uncovers design errors and usually results in design or process changes.

In fact, debugging uncovers many more design mistakes than it does broken pieces, so your initial guess should be that you goofed: not that the seemingly misbehaving Wollaston prism is defective, but that you've got a stray reflection you didn't notice. Most of the construction flaws we find are stupid; for example, the circuit may not be wired the way the schematic says, or a lens may be in backwards. Purely analog electronic debugging is complicated enough that books have been written on it,[†] and when optics, mechanics, digital electronics, and software get mixed in, the result is a real witch's brew.

Frequently, there may be many possible causes for a given problem; for example, say the bias level at the output of the photodiode front end amplifier is drifting. Is this because of circuit board leakage, an oscillating transimpedance stage, etalon fringes, laser instability, or temperature drift? Who knows? You'll just have to make a list of possibilities, then test for them in some intelligently chosen sequence, and with good test gear. If all you have to work with is your own buggy pre-alpha-test software, you'll get turned into a newt for sure.

19.4 GETTING READY

Guiding Principle: Problems never “just go away by themselves.”

Insist on rational behavior. Circuits are highly interdependent, especially DC-coupled circuits; when one part of the circuit misbehaves, *all* the signal voltages go nuts. In this situation, we have to use what we know about circuits—the component that isn't behaving rationally is the problem. For example, if an op amp's noninverting input is 50 mV above its inverting one, the output had better be railed high. That transistor with the 150 mV V_{BE} had better not be pulling any collector current. Signal voltages shouldn't depend on where along the trace you measure them. Being suspicious and logically scrupulous really pays off.

[†]Robert A. Pease, *Troubleshooting Analog Circuits*. Butterworth-Heinemann, Woburn, MA, 1991.

Divide and Rule. Okay, you have a gizmo that isn't producing the right output. Somewhere inside, something is wrong, but where? The quickest way to find out is to figure out what the major functional blocks in the instrument are, and test somewhere in the middle—misapplying a political analogy, we can call this the *divide-and-rule principle*. If the middle looks right, the problem is in the back half, and if it doesn't, it's in the front half. If the system has a single, well-defined signal path with no ancillary loops such as AGC or APC (automatic phase control, i.e., a PLL), you can do a nice binary search and be sure of finding it in $\log_2 N$ iterations, rather than N if you start from one end and work toward the other. Most systems aren't like that, but generally they're close enough that divide and rule is the right strategy. It's certainly the right approach to traversing our lists of possible problems.

It is frequently useful to cut in at the points where the major transformations occur: optical to electrical, analog to digital, and hardware to software, that is, the detectors, the digitizers, and the bytes read by the computer from the hardware interface. It is easy to look at the signals there, and to inject synthetic "good data."

Examine Your Assumptions. As the saw has it, "logic is an organized way of going wrong with confidence." It's not logic that's at fault, of course, it's that most of us aren't sufficiently rigorous about it. In debugging, we often find that one of our basic assumptions is wrong. This is surprisingly difficult to spot, because of the conceptual blinkers we're wearing, the ones that led us to make the mistake in the first place. The author once spent a long time contemplating an apparently perfectly aligned interferometer that wouldn't work, watching the fringes blink on and off uniformly across the field, before realizing that the beam wasn't well enough collimated, and that since the blinking fringes indicated only the *relative* phase of the two wavefronts, they were insensitive to the collimation error.

If you've eliminated the first thing you thought of, but after a day of head-scratching you can't find what else it could be, go back to the first thing and examine the evidence *very carefully*, preferably explaining it to someone else. Often the problem is hiding there, and your reasoning is what is at fault.

Learn What's Normal. During system integration, a lot of time is spent staring at perfectly functional equipment, trying to figure out what it is doing. Just getting a feel for the normal operating behavior of the system is time consuming and should not be skimmed on, so budget time for finding out.

All the design and simulation in the world won't prepare you for what you'll see on the scope. For example, CMOS analog multiplexers exhibit large switching transients when run into high resistance loads, such as a series RC filter. This looks really bad, but in fact the charge injected is always a lot lower than the spec, because the manufacturer's test circuit has a shunt capacitor connected to the switch (as in a T/H). An optical example is learning what defocus or misalignment does to your detected signal, and how to tell them apart.

Believe Your Noise Budget. It is an elementary scientific principle that when theory and experiment disagree irreconcilably, experiment wins. On the other hand, the strength of the evidence needed to overthrow a theory depends on how well established the theory is. In building apparatus, we use well-established theory to tell us how to change the system to get the performance we want, that is, when there's disagreement, *the apparatus*

loses and the theory wins. Don't discard your noise budget when the system isn't as good as you expected—the purpose of that budget is precisely to help you recognize the problem and fix it. Go through the math again, and think about what physics you might have missed, but don't chicken out and toss it. Electrical engineers seem to be temperamentally prone to this error.

Use Signature Analysis. The advantage a troubleshooter has over a debugger is that he knows that the system can work—it isn't misconceived, no design error will prevent its ever functioning properly. He also has knowledge of how it should work. What the beam quality, bias voltages, and signal levels should be throughout the system have been written down on the drawings, and from familiarity, the troubleshooter learns where to poke the patient to provoke a response.

This knowledge can be systematized into *signature analysis*, which checks the response of the system to carefully chosen stimuli. For example, if you're building a system for detecting Cerenkov light pulses from cosmic ray showers, you don't want to have to wait for a 10^{20} eV proton to arrive in order to check the response of every last PMT amplifier. What you want is a reproducible dim optical stimulus with good timing characteristics, available on demand. This need not be elaborate; in this case, you could use an auto spark plug with a strategically placed series resistor to provide a 5 V pulse for synchronization, plus a few neutral density filters to attenuate the light.

Back ends can be tested by using a digital scope or A/D card to capture a few representative waveforms from the front end, and a D/A card or arbitrary waveform generator (AWG) to replay them into the back end under test, to check for appropriate responses. Make sure that you play them back faithfully enough, or you're just fooling yourself: that means using enough bits, filtering out the clock ripple without group delay distortion of the signal, and making sure the records are long enough for the system to settle down before the event takes place. Test the tests by playing them back into the system they were measured on, and check that the response is identical to what the real signal provoked. Watch out for DC offsets, and remember that fast scopes are good to about 6 bits' resolution.

For each test, a few oscilloscope screen shots pasted into a notebook will enable a good technician to zero in on the location of the problem very rapidly. This is a huge time saver, which is not a negligible consideration when you're building highly complex systems that only the designer may understand completely.

Testing the response to signals in noise may require a white noise generator. Good calibrated ones are available from NoiseCom, and these are especially good for testing thresholding and statistical signal processing systems, for which prerecorded noise is not adequate—you can't measure the detection statistics when the test signals and noise are held in a fixed relationship, as in a recording.

On the other hand, testing statistics-based systems for adequate performance, for example, the very low false alarm rate in the ISICL sensor of Example 1.12, may be very time consuming. For those cases, it's worth having a library of near misses and squeakers-under-the-wire in your recordings, because that way you can verify that the device under test can distinguish them as well as it ought to.

Keep a Good One on Hand. Breadboards of each critical subsystem should be on hand so that suspected bad components can be tested. This works better for electronics than optics. Similarly, a few known good components should be kept around, so that if a bad batch arrives and causes problems, you can pin them down.

19.5 INDISPENSABLE EQUIPMENT

Since debugging and troubleshooting require close reasoning from ambiguous clues, they require first-class data; the reasoning and data have to work together like muscles and bones.

19.5.1 Oscilloscopes

The first thing you absolutely must have if you're going to work on instrument back ends is a good, fast analog scope. Its bandwidth should be at least 350 MHz, and 1 GHz is not too much. Remember that in order to gauge the shape of a waveform at all, you need to be able to see at least the first three harmonics, and to see it accurately, you need the first 10. Even if you're not working at VHF intentionally, many of the components you're using can oscillate up there, and not infrequently they will. SPICE will often miss this, because it depends a lot on layout strays.[†] If your scope misses it too, finding it will be that much harder.

A sufficiently good digital scope (1 Gs/s) is an adequate substitute, *provided* that it has some sort of envelope or spatial-histogram mode. In those modes, the digitizer works at full speed all the time, and each time bucket on the scope has its upper and lower limits plotted, so your signal is pretty well guaranteed to be somewhere inside that stripe. Live in that mode, and especially avoid “high resolution” modes, which use fast sampling speeds and some simple digital filter to get the illusion of more bits. This is useful occasionally but is seriously misleading in general. Some scopes make you use high res mode for measuring rise times and so on, unfortunately.

Get a scope from a manufacturer that makes really good vertical amplifiers. It is surprising how many scope makers will tout a 1 Gs/s scope that has a 300 MHz vertical bandwidth (for a while, one maker was selling a 2 Gs/s unit with a 350 MHz bandwidth). When called on this one, they will often claim that it is necessary for anti-aliasing protection, but that's just smoke and mirrors—real anti-aliasing would require that the vertical amplifier roll off by at least 40 dB by the folding frequency, and that they do not. Pay special attention to the scope's user interface; a lot of the low end models of the fancy brands are actually rebadged imports, some of which have user interfaces as bad as a \$35 DVD player's.

Regardless of your budget, don't buy a superfast scope (>2 GHz bandwidth) unless you really need it. You lose the option of high input impedance, and you trade off an enormous amount in accuracy and edge fidelity to get the extra speed. If you need the bandwidth, sampling scopes such as the Tektronix 11801C (widely available used) have much cleaner edges and are much faster, if you have repetitive signals. (They're also an order of magnitude cheaper.)

19.5.2 Spectrum Analyzers

For any high frequency work, a spectrum analyzer is also a near-necessity. You don't have to have your own private one, but for any sort of serious analog development, you'll

[†]If you're doing time-domain analysis of a circuit with a high frequency instability, the results will depend sensitively on the time step chosen and will be uniformly completely wrong. Gear-type integrators will mask the high frequency instability and are less sensitive to step size, but their results will be equally wrong.

at least need to have one for every three designers, and be ready to rent some more during system debug. Measurements of close-in modulation and noise on sine wave signals is only possible with an analyzer with a frequency-synthesized LO, so look for that. An analyzer with no frequency synthesizer has to phase lock to your signal in order to work at high resolution, and the phase locking masks frequency drift and close-in phase noise in your signal—very puzzling if you haven't tumbled to it.

Remember that weird problems are often traceable to VHF or even gigahertz oscillations, so get an analyzer that works up to at least 2 GHz. Tracking generators are nice for aligning filters and measuring frequency responses, and with an external directional coupler, they let you use your spectrum analyzer as a scalar network analyzer. Avoid the Swiss Army Knife combination spectrum analyzer/network analyzers, though. The best ones work about as well as amphibious cars, and the worst more like a combination shoe horn and pistol.

19.5.3 Probes

An often-overlooked necessity is good probes. Get several $\times 10$ probes (old time circuit design columnist Tom Kneitel used to say that they were marked 10:1 “because they ten to one away when you're not looking”). All those little grounding doohickeys are important, too. Don't throw them away, and make sure they stay together with the probe. A good, short probe ground is one of the keys to good measurements. Whatever you do, resist the temptation to use long floppy grounds. In the context of a 50 MHz measurement, a 1 inch ground lead or coax pigtail is long and floppy (it's about 20 nH, or 6 Ω). If you like using the 3 inch grounding clip leads that come with the scope probe, reduce the loop area—and hence the ground inductance—by twisting the wire back on itself (like a tangled phone cord, *not* coiled like a spring).

A typical $\times 10$ probe has an input impedance of 10 M Ω in parallel with 10 pF. That sounds great, until you figure out that at 20 MHz, 10 pF has an impedance of 800 Ω , and even less at higher frequencies. At VHF (30–300 MHz) and above, a better probe is often a 450 Ω chip resistor soldered in series with the center conductor of a piece of double-shielded 50 Ω coax, right at the node to be tested, together with a 50 Ω scope input. For a reusable test point, put it in series with an SMB jack (with a good ground).

Sometimes you need a fast probe, but 500 Ω isn't enough. For that case, you can get active probes, which have fast FET amplifiers similar to those used in the scope front end, but mounted right down near the probe tip. Tektronix and Agilent (Hewlett-Packard) make nice ones. You can make your own if gain flatness isn't much of a worry, for example, in a narrowband measurement. These probes are somewhat delicate with respect to overvoltage and electrostatic discharge, but they have 50 Ω outputs, which means you can use them with spectrum analyzers as well as oscilloscopes.

19.6 ANALOG ELECTRONIC TROUBLESHOOTING

One of the places that people get hung up fairly often is in debugging the analog signal processing system. Pease has written a good book on the troubleshooting part of this, and we've discussed ground problems in some detail in Section 16.5.2, so we'll touch on some of the other major debugging issues: pickup, supply coupling, oscillations, and assorted head scratchers such as analog multiplexers that won't switch.

Look for Trouble in the Right Domain. Oscilloscopes are famous for being bad at detecting frequency-domain problems, such as spurious signals or mild distortion, which any spectrum analyzer will pick up right away. On the other hand, frequency-domain instruments often miss things like major-carry glitches in DDSs and DACs—they look fine on a spectrum analyzer, because the glitch energy spreads out over a wide band—it isn't always in the same place in the cycle, so it's not just harmonics. Make sure you look in both domains.

Check that your noise floor is where you expect it to be. This is especially important with sampled data, where aperture jitter causes wideband phase noise.

Don't Expect Inputs to Be Just Inputs. Instruments are designed to tell us about our circuit without adding artifacts of their own. Sometimes, though, they aren't quite as innocent as they look. It is not unusual for a DVM, A/D board, or FFT spectrum analyzer to put small amounts of noise out on its inputs. Because this isn't what we expect to find, it's hard to spot the first few times, and will drive you nuts looking for your mistake, when it's really someone else's. Whenever you get a new instrument, connect its input to a scope and then to a spectrum analyzer, just to see what's coming out.

Don't Use Zillions of Cables. Physicists especially are prone to build apparatus with more coax patch cords than have ever been seen in one place before. This is a guaranteed way to make a measurement flaky. Lab patch cords have short and difficult lives—if you have 20 of them in one place, at least one is broken and two out of the 40 connectors are dirty, bent, or have had their center pin driven out of position. Even new cables talk to one another—don't expect perfect isolation from RG-58. Try to build your prototype signal processing in big enough chunks that you don't need so many cables, and use barrel connectors to join connectorized mixers, filters, and so on.

Get Rid of Pickup. Chasing spurious signals due to stray capacitance and inductance and supply coupling is a nearly inescapable debugging task.

Example 19.1: Unusually Severe Pickup Problems. One time, the author was debugging a card for an inexpensive optical sensor, which used three solar cells and a shadow mask to sense the position of an object in three dimensions (see Figure 19.3). Illumination was provided by IR LEDs IRED1–IRED5, which are chopped at about 100 kHz. This one turned out to be a real onion problem. The new surface mount card was exhibiting severe 100 kHz pickup problems, which nothing seemed to improve.

Onion Layer 1. After a bit of puzzling, it emerged that the 100 kHz was entering from two sources, capacitive pickup and power supply fuzz, which were roughly equal in amplitude.

Because of the femtofarad-level coupling capacitance, the pickup had a relative phase of $+90^\circ$, whereas because of the supply bypass capacitors, the ripple had a phase of -90° ; the two were thus 180° out of phase with each other, and because they were nearly the same size, they substantially canceled one another. This was a tricky debugging problem, because almost any change (even beneficial ones) disturbed the balance and made the symptoms worse.

Ultimately, the sources were largely separated by grounding the LED string, so that it was on continuously. This of course made the symptoms worse, but because it cleaned

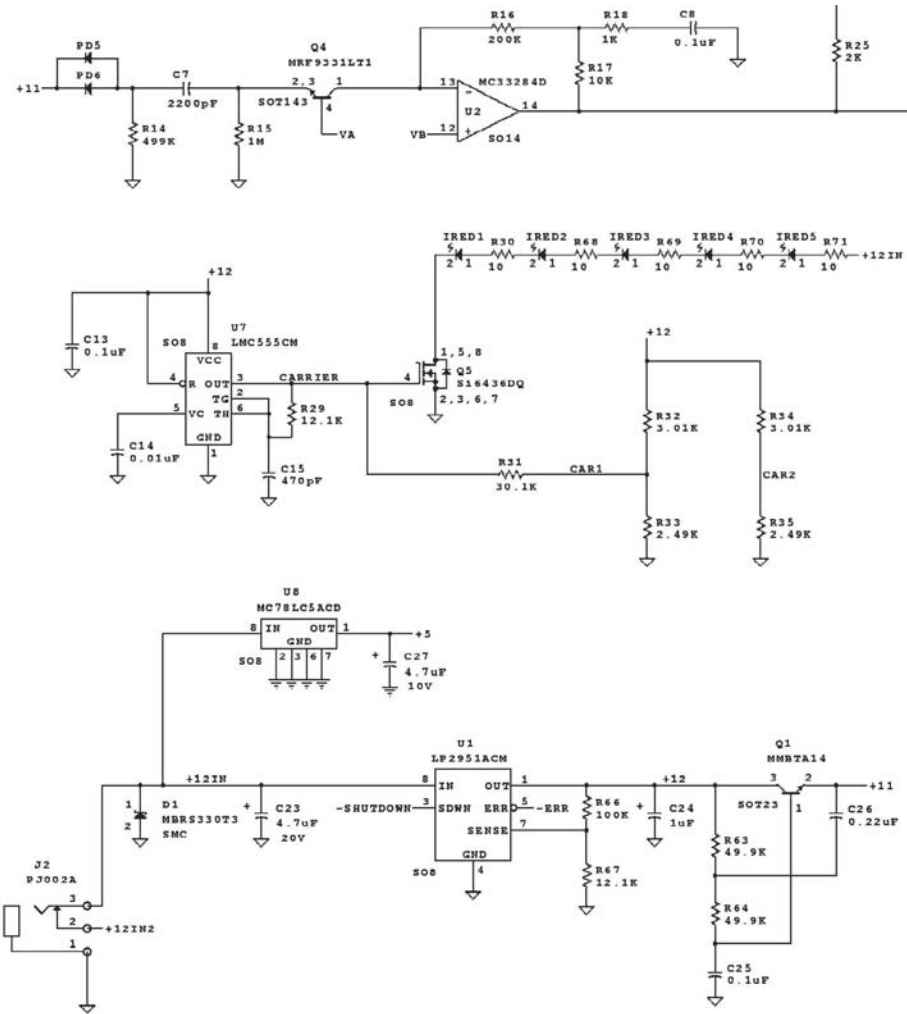


Figure 19.3. Noncontact head tracker schematic fragments: a tiny layout problem caused serious pickup troubles.

up the 250 mV ripple on the 12 V supply somewhat, it seemed like a worthwhile thing to stay with for a while. After cutting traces and hanging wires on two cards to the point where they were too flaky to use, the author applied foil shields to the top and bottom of the cards. This reduced the pickup signal by about 20 dB, which eliminated the cancellation that had been muddying the water. When the LEDs were allowed to switch again, the ripple problem surfaced, but that was cured by putting a 100 μ F capacitor from ballast resistor R71 (between the top LED and the supply) to the (grounded) source of the driver FET Q5.

Onion Layer 2. These measures largely cured the problem, except in one channel, where the pickup remained severe enough to rail the digitizer. Since the channels were identical

except for layout (all the amplifiers were on the same chip), this looked like a board layout problem. The layout had been done carefully, with ground current loops kept small, analog and digital devices kept separate, and the sensitive analog circuitry laid out one side of the card, over top of the analog ground plane. The power plane was on the far side of the ground plane, split between the ugly +12 V supply under the digital stuff and the +11.2 V analog supply under the analog stuff. Upon closely inspecting the layout once more, the only salient difference was that one single PC pad, 0.8 mm square with a via hole, connected to the sensitive input node of that channel, sat over the singly regulated 11.2 V supply.

This supply was isolated from the ugly unregulated 12 V supply by a low-dropout linear voltage regulator IC (U1), and was used for the less sensitive analog functions. The CMOS 555 timer (U7) that drove the LED driver transistor Q5 was connected to this supply and made it bounce up and down by about 50 mV. This presented no problem for the other two channels, because the front end circuitry was isolated from it by a two-pole Sallen and Key active filter (a fancy capacitance multiplier, see Example 14.1). Since no sensitive traces crossed the 11.2 V plane, it seemed that it should cause no problem to the third channel either. Nonetheless, touching a 1000 μF capacitor from there to ground instantly cured the problem, and permanently moving the V_{cc} pin of the 555 to the 10 Ω –100 μF bypass point completed the cure.

How on earth did this happen? The I-V converter in each channel had a 200 k Ω feedback resistor and had an AC-coupled $\times 11$ voltage divider (see Section 18.8), so that the total first-stage AC transimpedance was 2.2 M Ω . These were followed by phase-sensitive stages (Gilbert cell multipliers), which had a gain of around 400 for in-phase signals (the phase was not particularly carefully tuned—these were primarily frequency converters to allow bandwidth narrowing as in a lock-in amplifier). The size of the detected signal was about 200 mV DC, which at a frequency of 100 kHz means that the required stray capacitance was

$$C_{\text{stray}} = \frac{200 \text{ mV}}{2\pi f V_{\text{pk}} R_f A_V \cos \phi} \approx 16 \text{ fF}, \quad (19.1)$$

assuming that the phase shift ϕ is zero (reasonable, because a 90° lag from the filter capacitor is followed by a 90° lead from the capacitive pickup).

Considering that FR-4 circuit board has a dielectric constant of about 4.5, and that the pad was about 0.020 inch (0.8 mm) from the power plane, the area required to get 16 femtofarads of capacitance is only

$$A_{\text{min}} \approx \frac{4\pi d C_{\text{stray}} (\text{pF})}{1.12 \epsilon_R} \approx 0.003 \text{ cm}^2, \quad (19.2)$$

where fringing has been neglected, all spatial dimensions are in centimeters, and C is in picofarads (the Gaussian unit of capacitance is the centimeter, which is equal to 1.12 pF).

Fringing decreases this number, and the relief of the ground plane under the pad increases it, so this is not an accurate estimate. Nevertheless, it's clear that no very large area is needed—the 0.8 mm pad is about 0.0064 cm², so this is an easily believable source of the stray coupling—but only in hindsight. How can this sort of problem be prevented?

In this case, the problem was overconfidence. A dead bug prototype had worked well, and the board layout looked sensible, so the step of estimating the sensitivity numerically

was omitted. Once that 16 femtofarad number surfaced, it became obvious that very, very small layout problems could be responsible. In a high gain, high impedance circuit, it is absolutely vital to estimate just how large the stray coupling is allowed to be.

Moral: Greedy optimization isn't always the right approach. For pickup, fix anything you find, regardless of what it appears to do to the symptoms; watch, though, that it is really fixed.

Eliminate Supply Coupling. This example is somewhat unusual in that the coupling was indirect; most of the junk on the supplies had been removed by regulators and active filters. The best way of eliminating supply coupling in general is to make sure the local decoupling of all high level signals and high gain amplifiers is adequate. Note that this does not necessarily mean sprinkling 0.1 μF capacitors everywhere. The key to decoupling is to make the undesired AC path through the supply much higher impedance than the desired one (through the bypass capacitor to the other end of the AC load). This requires locally raising the impedance of the supply with series resistors or low- Q inductors (e.g., ferrite beads), and reducing the impedance of the bypass by choosing the right value and the right capacitor type, or alternatively by using active supply filters, as here.

Sprinkle Active Supply Filters Strategically. In the previous example, after the first onion layer was peeled off (pickup and supply ripple), the 555 was still putting large (50 mV p-p) 100 kHz transients on the 11.2 V supply, which is where the 10 V photodiode bias came from. Even in the face of that 50 mV ripple, and with photodiode capacitances of about 200 pF, it was not necessary to put any additional decoupling on the +10 V supply from the active supply filter. From Eq. (19.1), and noting that the total offset voltage in the well-behaved channels was less than 10 mV, we can estimate that the in-phase spurious signal on the 10 V supply had to be less than 0.2 μV . The calculated rejection of the Sallen and Key Butterworth filter is approximately 140 dB at 100 kHz, so although it is probably operating far less well than this, essentially no ripple at the operating frequency is getting through. This is rather remarkable, because of the apparently direct connection of the quiet 10 V rail to the troublesome +11.2 V; that 5-cent MMBTA14 does a terrific job.

19.7 OSCILLATIONS

Many of the more vexing flaky problems in analog electronics come from high frequency oscillations. Somewhere in our signal processing chain, an emitter follower has a source impedance that is too low, together with a capacitive load, and it starts shrieking at umpty-ump megahertz. Bandwidth limitations in subsequent stages filter out the VHF energy, leaving only the gross bias shifts, intermodulation, phase funnies, and extreme flakiness behind.

Troubleshooting oscillations is easier if you have good test equipment: a 1 GHz oscilloscope, a good spectrum analyzer, and some FET probes. A big, rich lab will have these things, if you can get time on them. In most labs, someone else is usually using them, and you'll have to wait till after hours. Or maybe you don't have access to them at all. In any case, you need some practical skills in troubleshooting them.

19.7.1 My Op Amp Rings at 1 MHz When I Put This Cable on It

We looked at op amp instabilities in Section 15.4.1; they can be caused by too much capacitance on the summing junction or on the output, both of which add excess phase shift to the loop gain, and so reduce the phase margin. Unterminated coaxial cables look like capacitors of 100 pF per meter of length, for frequencies $f \ll 1/(2\pi t_{\text{transit}})$. This capacitance combines with the open-loop output impedance of the amplifier, forming an RC lowpass network in series with the feedback path. The phase shift associated with this network reduces the phase margin of the amplifier and produces ringing or even oscillation. It can be corrected by putting a resistor of about the same size or slightly larger than the op amp's open-loop output resistance (look for it in the data sheet or device schematic, though you may have to dig a bit) in series with the output. You can close the DC feedback to the cable's input, as long as you put a small capacitor from the op amp's output to the summing junction. (Note: amps with rail-to-rail outputs are the worst for this.)

19.7.2 When I Wave at It, It Waves Back

If the circuit's behavior depends on where you wave your hands, something is probably oscillating, or your bypassing and shielding is inadequate. It's pretty understandable: we're so often putting a wideband front end amplifier on some device (e.g., a PMT or packaged photodiode) that looks as though its maker intended it only for electric-eye doorbells. We wind up having to ground the cases of photodiodes with copper tape and solder in order to prevent them from picking up FM stations, or oscillating due to input-output coupling. If sensitive high frequency amplifiers are involved, give lots of thought on how to stitch all the grounds together adequately—usually two ground planes with lots of plated-through holes, and perhaps a metal clamp for the diode package.

19.7.3 My Circuit Works Until I Let Go of It

If the circuit is poorly isolated from outside influences, those influences behave like circuit elements. Sometimes these show up as performance variations that depend on where you hold the box, or whether there's a metal surface nearby. These kinds of problems are generically called *hand capacity*. There is a very old and probably apocryphal IBM story of a then-new mainframe that wouldn't work unless there was a scope probe hung from one particular circuit node—the first unit was allegedly shipped like that, and quietly fixed during a preventive maintenance call.

19.7.4 My Transistor Amplifier Oscillates at 100 MHz

Or 1.5 GHz, or someplace that's invisible on your scope. Transistors can usually find something around to use for a resonator, so you'll have to de- Q the oscillation by putting a damping resistor in somewhere. A 10 Ω resistor in series with the base will usually do it. In a differential amp, check if it's a differential or common-mode oscillation, because the medicine is different; a differential oscillation has to be fixed with emitter resistors, whereas a common-mode one usually goes away with a series RC (small R , really small C) from the emitters to ground. To check which it is, you can use two probes, or put a 1:1 resistive divider between the collectors and look at the midpoint to see if it moves.

19.7.5 Another Kind of Digital Troubleshooting

The first thing to do with a circuit that is oscillating or has hand capacity problems is to start poking your fingers into it. Not figuratively, either. Touch the input and output of each stage with your finger, while looking to see if the strange behavior changes. (As long as there are no hazardous voltages present—don't try this on an APD preamp!)

Putting your finger on the summing junction (– input) of a fast op amp will cause oscillations rather than cure them, but for discrete stages or noninverting inputs, this test is both easy and sensitive. At RF, human fingers look like a lossy capacitor, which is an excellent temporary fix for most oscillations.

If the finger trick doesn't locate the problem, start looking for impossible bias conditions: things like an emitter follower that is obviously conducting but whose DC bias is impossible (e.g., $V_{BE} = 0.11$ V). That sort of thing can happen during large-signal (Class C) oscillations, in which the transistor is biased off for most of the cycle. Gigahertz oscillations can be really funny that way, for example, an oscillating common-base stage whose apparent V_{BE} changed from 0.15 to 0.4 V when the scope probe was moved a few millimeters along the lead—it looked like a nice DC level on a 500 MHz scope. (It was a biased cascode for a photodiode—see Section 18.4.6.) Remember that (as Pease quotes), “if you notice something funny, Write Down Amount of Funny.”

Warning: If you're using a DVM for this, put a balanced decoupling network consisting of a 10 k Ω resistor in series with each lead, and a 10 nF capacitor connecting them on the meter side. Large amounts of RF will drive most DVMs *crazy*, not to mention what those nice antenna-like leads will do to the circuit—if it wasn't oscillating before, it probably will when you hang a meter of probe wire directly on it. Once you've found the stage that is most sensitive, you can try the following tricks.

Add 10–100 Ω in the base lead.

Add extra supply decoupling—1–20 Ω in the supply lead.

Put ferrite beads on base and emitter.

Put 10 Ω in series with the load. Normally, putting 10–100 Ω in any two leads of a BJT will stop UHF oscillations.

Put an RC snubber on the output.

Use constant-resistance diplexers. This eliminates the reactive load.

Reduce the collector current. This reduces f_T and hence the high frequency gain.

Use a slower device.

19.8 OTHER COMMON PROBLEMS

In debugging and integration, remember that both are onion problems: once the worst problem has been fixed, the next worst will surface, and so on until we get exhausted or the system is perfect. If something we do affects a problem, that usually means that we're getting close to its cause; often, being able to make the problem worse is nearly as valuable as being able to improve it.

Hmm. Must Be a Dead Chip. . . . This is a strong possibility in troubleshooting, but not in debugging, unless it was working before, and you just destroyed it, for example, by putting it in its socket backwards and applying power (if you did that, don't even bother

checking—chuck it). It is amazingly rare to get a bad part from a reputable supplier these days. If the circuit hasn't worked yet, it's almost certainly a design or wiring error. Check all the supply voltages on all supply pins of all ICs. Check that you know why any unconnected pin is unconnected—usually because it's an unused output or an offset adjust. Anything else needs to go somewhere, especially resets, chip enables, three-state enables, and so on.

This Precision Op Amp Has an Offset of 100 mV. Assuming the chip is good (which we always do at first, right?) and that you haven't done anything perverse with the offset adjust pins, the main possibilities here are exceeding the input common-mode range, out-of-spec supply voltages, nonlinearity due to slew limiting, saturation, or oscillation, or a DC error due to the input bias current times a huge source resistance. You see this symptom often with high- Q active filters and the early stages of a signal processing system—subsequent filters mask the clipping, leaving only the low frequency nonlinearity, offsets, and splatter—signals showing up frequencies where they have no business being.

Sometimes I Have to Cycle Power to Get It to Work. Some circuits work fine once they're going, but have trouble starting up properly when you power them on. Oscillators are sometimes like this, when the large-signal and small-signal behavior of the transistors are very different. Startup problems are especially likely when you have cascaded NPN and PNP differential pairs, or op amps whose output swing is bigger than their input common-mode voltage range, because in those situations there are two stable operating points: besides normal bias, you might have the output wake up at one rail, which will completely turn off the input stage of the following amplifier. Depending on the design, this may be harmless and self-correcting, or it might destroy the whole system (as in a motor control loop). Make sure you consider what happens under all imaginable startup conditions, with one supply coming up before the other or both coming up slowly, or one pausing and going down again for a short time, and so on. Startup behavior is usually different at different temperatures and supply voltages, so make sure you cover the whole range. Startup circuits and power on resets really are surprisingly subtle.

It Can't Handle the Input. Overdriving IC inputs can lead to very strange behavior, for example, in CMOS, you get charge migrating all around the isolation tubs and changing all the threshold voltages, and some op amps reverse their output polarity when their inputs are driven too close to the positive supply, which can lead to the destruction of your motorized or temperature-controlled device. Watch out for lockup conditions, and always test short circuit all inputs and outputs to their supplies and grounds (make sure you design things to survive this). Some op amps can drive their own inputs nuts—especially watch CMOS ones with rail-to-rail outputs but not rail-to-rail inputs. Instruments can be saturated by strong out-of-band signals—watch out for oscilloscopes, lock-in amplifiers, and FFT analyzers here.

It Can't Handle the Load. Overloading the output of an amplifier will usually cause worse problems than simply reduced output swing. Symptoms include sudden death (e.g., the function generator with a “50 Ω ” output that blew up when 50 Ω was attached, because its output transistors wouldn't handle the current); clipping due to current limiting; frequency response degradation; distortion; or oscillation due to protection circuitry or to capacitive loads.

My Low Noise Amplifier Has a Weird Gain Peak. This is a common problem with low voltage noise op amps when used with high source and feedback impedances (“high” can be 1 k Ω sometimes). There is a big peak in the hundreds of kilohertz to low megahertz, which is hard to explain; it’s actually a mild case of the same disease that causes loop oscillations—loading due to input or load capacitance. There’s more in Section 18.4.2.

My Simulation Is Perfect, Why Doesn’t It Work? Usually because of simulation inaccuracies, circuit strays, and supply coupling. Try putting your finger on the supply trace and see what happens. Improve the grounding and decoupling by massive overkill. Then look again—it’ll probably be a lot closer, provided you did a really realistic simulation (i.e., included the actual package and PC board strays, parametric variations of components, transistor-level models of your ICs rather than macromodels, and so forth). Theory works, but circuits are hard to model exactly, especially in the time domain.

I Don’t Understand My Oscilloscope Trace; It Doesn’t Make Any Sense. . . . If you’re using a digital scope, run it up to its fastest horizontal sweep rate and work down. You’re probably seeing aliasing. The 20 MHz and 100 MHz input filters (usually on the vertical menu) are often helpful in sorting this out. You ought to be using the HISTOGRAM or ENVELOPE function, which runs the digitizer at full speed and displays a color histogram (or at least the high and low values) in every pixel width. Use the “Sample” mode only for automated measurements (e.g., rise time) and absolutely avoid the “High Res” mode.

My Entire Analog Multiplexer Circuit Wiggled Out. A seldom-discussed pathology of CMOS circuits is their tendency to connect all their pins together when even one is overdriven. This weird behavior is caused by the input protection circuitry, among other things, and happens when any pin is driven beyond the supply or ground rail. If a circuit containing a multiplexer or analog switch IC is acting strangely, failing to switch properly, and all the inputs seem to be loaded down in unaccountable fashion, look carefully for somebody being driven beyond the supply.

I’m Dying of Capacitance. A common complaint, unfortunately. Low level, high impedance signals have to be shielded to protect them from picking up TV signals and so on, but coaxial cable has a capacitance of 100 pF per meter, which rolls the bandwidth off really badly. There are three basic techniques for this. The best is to put the first-stage amplifier right at the transducer (photodiode, microphone, what have you). This is usually possible with a bit of creativity; you can put a hermetically sealed (ceramic or metal can) op amp inside a vacuum chamber, if you heat-sink it to the chamber walls (you’ll probably have to drill a hole in it for UHV use, and watch out for soft solder). Next best is reducing the input impedance with a really quiet common-base input stage or transimpedance amp, as we did in Chapter 18. (Remember what that did to the input voltage noise of the amplifier, though.) The neatest-looking but least satisfactory method is guarding, that is, driving the cable shield from the output of a voltage follower, to bootstrap the cable capacitance. In Section 18.4.8 we saw that bootstrapping can work well but has the same noise-multiplying tendency as transimpedance amps; your frequency response will improve, but your SNR won’t.

My Network Won't Tune. In Section 16.11, we did our tuning by a greedy algorithm—tune knob #1 to a return loss maximum, switch to knob #2, and continue round robin fashion—and discussed what to do in noisy situations where the indicators are unreliable: tune in big enough steps that you can clearly see where things are getting worse, and head for the middle.

My Network Won't Tune, Part 2. There are other reasons for networks failing to tune. In a matching network, the maximum impedance transformation ratio is Q^2 , so if your circuit Q is too low, the matching condition will be inaccessible. In that case, you can increase the Q by making your series inductors and shunt capacitors bigger, and your shunt inductors and series capacitors smaller, then try again. Calculate first, though, because that's a fair amount of work; keeping the resonant frequency of each section constant in the process is usually enough. Maybe the most common problem of all is running out of range. There should be two equivalent positions of each trimmer capacitor—if not, you've run out of range.

My DC Level Is Drifting All Over the Place. Watch out for thermoelectric potentials, especially gradients across active devices. Thermoelectric coefficients of wire thermocouples are usually 20–50 $\mu\text{V}/^\circ\text{C}$, but Cu–Si's is 400 $\mu\text{V}/^\circ\text{C}$, and copper/tin oxide is even bigger. Soldered connections are not usually a problem, but Cd–Sn solder has a bit lower thermoelectric coefficient with copper than Pb–Sn has, and Cu–Ag and Cu–Au are very low.

Narrowing My Bandwidth Isn't Helping. In Section 13.6.1, we saw that drift and $1/f$ noise isn't helped by bandwidth narrowing alone. You'll have to get more of the measurement bandwidth further from DC, with AC modulation or fast scanning and signal averaging. You might then find that you can widen your bandwidth instead.

My Noise Looks Like Data. Some measurements, especially slow-scanning ones like magnetic force microscopy, are prone to generate reasonable-looking wrong answers. This is a difficult test of one's scientific integrity sometimes; it helps to remember that other people (customers or readers of your paper) are going to be ruthless about examining it, so you might as well be too. If you wouldn't believe it in somebody else's work, where you weren't too familiar with the apparatus and had no personal stake, others are unlikely to believe yours either. This is a tough one; some people are their own severest critics, so apply this rule with some sensitivity, and show the data to a friendly colleague.

19.9 DEBUGGING AND TROUBLESHOOTING OPTICAL SUBSYSTEMS

There's a considerable amount of overlap between this topic and Chapter 12, because unlike electronics, just assembling an optical system requires a lot of testing and checking as we go along. We'll concentrate here on properly aligned systems that don't work properly.

Optical debugging is simpler than electronic in one way, because there are normally very many fewer optical elements than electronic ones. In other ways it is more complicated, because each element does much more complicated things, and the available test equipment is impoverished compared with oscilloscopes and spectrum analyzers. It isn't

too bad in the visible, where you can see where the beam got vignetted (say), but in the IR and UV, it isn't so easy.

This difficulty is aggravated by any unclarity in our intuitive idea of just what the system does, and how. It is crucially important in optical debugging to be able to describe the operation of each part of the system in at least two complementary ways, for example, geometric imaging, wave propagation, and photon following. Different things are obvious in different pictures, and thus a flexible understanding is a vital tool.

The most common optical debugging problems are caused by things you didn't think of, for example, that iris-type shutters cause intensity nonuniformity unless they're at the pupil, or that a thin lens illuminated with a laser exhibits very fine Newton's rings. The second most common type comes from having too simple a mental model of how things behave, so you miss their subtle interactions; diode laser feedback and thermal lensing are examples. Table 19.1 has some possible sources of instability to start from.

Here's some advice.

Each Mode Is Its Own Interferometer. Lasers often oscillate in several modes simultaneously, so that their optical spectra look like fine-toothed combs with 2 to 100 teeth. Because the speed of light is so high, the mode spacing in ordinary sized lasers is 100 MHz (in a big gas laser) to >10 GHz in a Fabry–Perot type diode laser. Detection makes all these modes interfere with themselves and each other, producing a strong baseband signal plus a forest of huge ugly spurs near harmonics of the mode spacing, and

TABLE 19.1. Some Sources of Instability in Bulk Optical Measurements

Optical
Etalon fringes
Polarization nonuniformity (especially in VCSELs)
Aberrated beams
Misalignment
Not enough space for misalignment to show up clearly
Schlieren due to air currents or temperature gradients
Vignetting
Stray light
Glints
Laser
Collimation errors
Interferometer path difference $\neq 2N$ cavity lengths
Wiggle noise due to spatial side mode
Mode partition noise
Mode hopping
Interference with delayed reflections
Coherence fluctuations
Mechanical
Order overlap in gratings
Bending due to thermal gradients
Mechanical creep
Hysteresis
Backlash
Stick–slip energy storage
Lubricant film instability

(at least in gas lasers) unstable baseband spurs caused by the anharmonicity of the comb (see Section 2.13.7). Because of the spurs, we confine ourselves to the neighborhood of DC, and so for our purposes, each mode is its own interferometer.

There are two important consequences of this: first, an interferometer using a multimode laser must have a path difference that is close to $2N$ cavity lengths (taking the refractive index into account of course), or else the sensitivity will be reduced and the noise increased, since the different phases of the different modes cause mode partition noise to be transformed into intensity noise (see Section 2.5.3). Second, any high frequency components of the desired signal will be mixed down into the baseband region and made unstable by the frequency instability of the mode spurs.[†] A noisy interferometer can often be traced to one or both of these effects, so dig out a ruler and check that $\Delta z/\ell = 0, 2, 4, \dots$

Mentally Follow Your Photons. Our familiar descriptions of wave propagation in terms of plane waves, rays, and δ -functions are all singular in some way, and hence unphysical. A plane wave has infinite energy, so (although the math works out fine) it is a bit unintuitive sometimes to think of waves of finite extent and finite power being sums of plane waves. If it's hard to get a crystal-clear view of how some optical subsystem should behave, try changing domains. Go back and forth between frequency, space, and wave packets (pulses or Gaussian beams). Thinking about a general beam as a superposition of Gaussian or tiled, square uniform beams at different positions is often helpful too, for example, thinking of all the tiled components of a high-NA focused beam not adding up as scalar addition, due to their different propagation angles (Section 1.3.1). (See Figure 19.4.)


Use Different Descriptions. There are often a few different equivalent ways to describe the physics of your setup. For example, consider a system for inspecting a smooth but not necessarily flat transparent film for defects, as shown in Figure 19.5. The system uses a two-pass geometry with a small, fairly dim source and a low tech imaging system, plus a big sheet of cheap retroreflector to send the light back to the source.

The physics can be described as two passes through a pellicle, or as four independent beams (top, top), (top, bottom), (bottom, top), (top, top). Phase and polarization issues and other delicate effects are often more easily understood by switching back and forth this way. Another well-known example is the plane mirror Fabry–Perot interferometer, which can be solved by a geometric series approach as in Section 1.6.2 or by patching together the plane waves inside and outside the cavity.


Watch for Polarization. Polarization funnies can drive you nuts. If you're using a polarized light source, it's worth following the polarization around with an analyzer to make sure it does what you expect, and watch out for sign errors (e.g., right- and left-circular), because that's where a lot of mistakes lie. For example, beyond Brewster's angle a right-handed elliptical polarization stays right-handed on reflection.

Get a few polarizers of good optical quality and wide angular aperture (film for the visible, Polarcor for the near-IR, Rochon prism or walkoff plate for the UV), to make sure that the polarization does what you expect. Mark their orientation clearly, and look

[†]The modes are remarkably narrow as a fraction of the optical carrier frequency (10^6 Hz/ 10^{15} Hz), but much wider as a fraction of the mode spacing, which is what's relevant in the baseband signal.



(a) Self-Consistent Fields



(b) Sums of Scattered Components

Figure 19.4. Following photons through a thin-film stack: self-consistent fields versus sums of single-path contributions.




Figure 19.5. Two-pass inspection system for surface defects.

hard; polarization effects are sometimes subtle, and small angular errors can have a big effect on measurements.

Stress birefringence and retardation nonuniformity in wave plates and Pockels cells are frequent villains in systems needing high extinction. A less obvious one is frustrated TIR, for example, from optically contacted surfaces that aren't quite in contact, which can cause really major polarization funnies. These difficulties usually change significantly with temperature, so a little work with a hair dryer and crossed polarizers will often tell a clear story.

Know Where Your Stray Light Goes. Front surface reflections are a bit like land mines: they don't disappear on their own, and the only way to be safe from them is to know exactly where they are. Every dielectric surface produces at least two reflections. Make sure you know exactly where each one goes, and what it will hit. This becomes so complicated after a few bounces that you'll probably fall back on measuring it instead, which is OK as long as it's done in a realistic setting. For example, using a HeNe laser to measure the back-reflection from a laser diode collimator ignores the wildly nonlinear nature of the interaction between the back-reflection and the laser cavity.

First-Order Analysis Isn't the Whole Story. Physicists (e.g., the author) are especially prone to get snookered by first-order analysis, for example, "components above f_c get chopped off by the filter," neglecting the fact that this filter might easily have a shape factor of 30, so that the transition band is very important. Similarly, "polarization insensitive" doesn't mean zero polarization sensitivity, and "shot noise limited" doesn't mean that all your outliers couldn't be coming from an amplifier with popcorn noise, even though its rms contribution is way below the shot noise. The wisdom here is to calculate a hard upper limit to the effect you're neglecting before neglecting it, probably making a few measurements along the way.

Etalon Fringes Are Sensitive to Everything. One of the things that makes optical systems mysterious sometimes is that they're hundreds of thousands of wavelengths across, and often sensitive to effects at the 1 ppm level. Thus parts-per-billion effects can easily become a problem. Gentle air currents, too slow to detect on the skin, can cause big enough index gradients in the air to steer beams around a little: the effect isn't much, but the amplifying effects of etalon fringes can turn it into serious low frequency instability in your measurement. You can often spot this by waving your hand around near the optical path. Small temperature gradients can do the same thing; dn/dT for air is comparable to that of glass. The solution is good packaging and, even better, getting rid of the fringes.

19.10 LOCALIZING THE PROBLEM

If the optical problems are unstable with time, it's an etalon fringe, thermal gradient, or source instability. If your SNR is too low, check for major design blunders, for example, a misplaced factor of 10 or assuming that the focal volume of a Gaussian beam is equal to πw^2 times twice the Rayleigh range (that's a factor of 10 right there[†]), and then look

[†]This one bit the author once, leading eventually to an $8\times$ reduction in the system specification—and that was after having promised the original spec. Very embarrassing.

for places you're losing a decibel at a time. If your beam quality is poor, look for laser feedback. If your CCD pictures are ugly, look for bloom and clock level problems, or too low a temperature in a cooled CCD camera; if that isn't it, put an IR blocking filter in front.

One technique that isn't used enough is intelligent manual manipulation. If your baseline is swimming around all over the place, try pushing with your fingers on the optical mounts, flexing the lens barrels, tapping things with the eraser end of a pencil, breathing on things to warm them up a bit and stir up some index gradients. By finding the place that's most sensitive, you've got much closer to the locus of the problem.

19.10.1 Is It Optical or Electronic?

When chasing some unknown spurious junk in your signal, the first question that has to be answered is whether it's optical or electronic in origin. This is a special case of the divide-and-rule principle but is worth putting a box around because the two cases need such different treatment. Often it's as simple as blocking the light beam and seeing if the junk goes away, but that can be seriously misleading sometimes; for example, consider a front end amplifier that oscillates whenever the photocurrent exceeds $100\ \mu\text{A}$, as can easily happen with a common-base input stage without a base resistor. That's an entirely electronic problem that goes away when you block the light beam. Other systems (e.g., power-stabilized lasers) won't operate properly without the beam present.

Thus determining which domain the problem lies in can be a subtle business. The first rule is not to actually turn anything off, because doing that changes both the optical and electronic operating conditions, and hence muddies the water. Try attenuating the light gradually, with a polarizer, an ND filter, or slight vignetting, and see if the junk is proportional to the optical power, which suggests an optical problem, or perhaps cuts off at some nonzero beam brightness, which suggests an electrical problem. Note whether it cuts on again at exactly the same place—hysteresis here is a nearly infallible indication of an electronic problem.

Alternatively, replace the light source with a flashlight to maintain the DC bias conditions. Statically balanced systems like dual-beam spectroscopy or DC Nomarski setups can be fooled by replacing one of the beams with a flashlight; autobalanced systems such as laser noise cancelers can be tested by unbalancing the two arms by adding a bit of quiet light to one arm; if the noise gets much worse without changing significantly in character otherwise, the problem is probably optical. If the problem goes away when you turn off the room lights, it's probably optical too.

If none of these tricks is applicable, you can examine the light beam itself to see if you can see the problem there; put in a beam sampler of some sort, and detect the resulting sample, looking for similar spurious signals in the sample output. The main thing to be aware of is that simple-minded tests can be very misleading here—use a few lines of evidence, and keep the possibility that you're wrong in the back of your mind.

19.10.2 Component Tests

Many optical components are made in small batches, with a lot of hand work and without logistics systems of any sophistication. Thus there is some likelihood of a mixup, such as one lot not being tested, or the wrong part shipped. This is sometimes obvious, but often not—and then it will reliably give you fits. One real-life example was half a dozen beam

separator cubes[†] in which the quarter-wave plate was tuned to the wrong wavelength, so the transmit/receive selectivity was only 15 dB, rather than >30 dB. This didn't show up until the diode laser started mode hopping madly. It's really worthwhile testing critical optical components before using them. Furthermore, swapping in a new one and seeing if it helps is an unreliable test unless the new one is really a known-good old one, or at least from a different manufacturing batch. In the beam separator case, all six were identical in performance, but identically wrong.

19.10.3 Beam Quality Tests

Intensity ripple and scattered light are common problems, too. A video camera and frame grabber with some simple software to display a false colored image of beam intensity versus position is a big help, but you can also just shine the light on the back wall of the lab; intensity ripple has to show up as far-field artifacts eventually. Knife-edge and chopping tests are a poor match for this, because they integrate over one dimension, which usually washes out the variations we're trying to see.

19.10.4 Collimated Beam Problems

A shear plate[‡] inserted into the beam will show up gross collimation errors or aberrations; it needn't be at 45° to work, so you ought to be able to fit it in somewhere. It isn't particularly sensitive, so for more subtle beam quality problems, use a pickoff mirror and a better-grade collimation tester such as a shear plate dithered back and forth in angle (like the late lamented Collimeter, formerly made by Blue Sky Research—see Section 12.6.4), or a measuring interferometer.

19.10.5 Focused Beam Problems

Test focused beams with a knife edge. A piece of a broken mirror makes a really good knife edge for systems that are already assembled, because it can send the test image off at a convenient angle; it isn't easy to make a good reflecting knife edge any other way. The illuminated area is very small, so straightness of the edge and flatness of the mirror hardly matter—aluminized window glass is fine. You can find the focus in the usual way, by looking at which way the shadow moves when you move the blade, and measure the spot diameter by plotting total reflected light versus blade position. Do make sure the motion is perpendicular to the edge, or you'll be off by a factor of $\sec \theta$. You can mount the mirror shard on a piece of bent aluminum sheet hanging from one screw like the blade of a paper cutter, so that it can swing in and out as needed.

If the knife edge test shows a badly aberrated beam, you can often identify the problem by racking the focus in and out and looking at the near-field pattern. Gross astigmatism shows up clearly as well-separated tangential and sagittal foci, for example.

Beam scanning systems can use bar targets, for example, the ubiquitous USAF 1951 three-bar target, which function like a chopper test (the inverse of a knife edge); the

[†]That is, polarizing beamsplitters with $\lambda/4$ plates cemented on one end, for the circular polarized transmit/receive duplexer trick.

[‡]For example, an unwedged optical flat or a specially selected microscope slide that happens to have a sweet spot where the faces are flat and parallel.

beam moves and the blade stands still. The beam usually scans too fast to be examined in detail, but the fact that you're measuring the real signal in a realistic way offsets that disadvantage. That also works much better with invisible light, because it doesn't rely on eyeball estimates. Watch out for asymmetric aberrations here—an astigmatic spot (which is long and thin) shows up beautifully on a properly aligned bar target, but degrades far more rapidly than a symmetrical one as the target is twisted.

19.10.6 Viewing Techniques

In Section 11.8, you were exhorted to put a viewer into the system. That idea can be exploited for debugging and troubleshooting, either by using the built-in viewer and racking focus in and out, or by examining a low-NA aerial focus with a microscope. Astigmatism, coma, and spherical aberration show up pretty clearly this way. Before doing this, calculate the maximum the power that can exit the eyepieces, and after it's installed, measure it to make sure it's safe and comfortable ($1\text{--}10\ \mu\text{W}$ is a good range, nice and bright without being painful, and far too dim for any danger).

19.10.7 Test Techniques for Imaging Systems

Imaging system performance is best characterized by MTF versus position and wavelength, contrast, uniformity, and efficiency. If you're using a focal plane array detector (e.g., a CCD), some simple testing software and a set of sinusoidal test reticles are a more or less complete solution for detailed testing, but during setup we need something quicker, that we can eyeball. One good method is to use a radial bar pattern, and mark it with a sharp-tipped felt pen and a circle template, to indicate the finest lines your system should be able to reproduce. If you're handy with Postscript, you can generate reticles like this on paper, on demand, which is pretty convenient.

19.10.8 Test Techniques for Light Buckets

Systems with little or no spatial resolution (e.g., photometers and differential detectors) usually rely heavily on stability, linearity, and good baffles. Long-term stability is hard to measure really adequately, especially during troubleshooting, but surrogates exist. Intensity instability is usually caused by thermal drift of filters, etalon fringes, and the leakage current of detectors and front end amplifiers. Warming with a hair dryer is a good test for this, but watch out for Schlieren effects from air turbulence—warm the system up and watch it as it cools down undisturbed. Putting in a chopper is another good test in a DC system—detector and amplifier leakage won't get translated to AC, whereas optical instability will; comparing the AC and DC outputs will identify the source of the trouble. The easiest way to check for linearity in a light bucket system is to use an LED with a small audio-frequency sinusoidal modulation on top of a constant DC drive current. Move the LED around or attenuate it with ND filters, while looking at the ratio of the AC and DC signals on a couple of DVMS—in a really linear system, the two outputs will track exactly. Alternatively, two LEDs in an integrating sphere, with different modulation, will generate mixing products in a nonlinear system. You need the sphere to keep the intensity ratio independent of position.

19.10.9 Invisible Light

All these tests get harder in the UV and IR. Inexpensive cameras work between about 250 nm and 1 μ m, but outside that range things get difficult. Fluor-coated CCDs work well into the UV, and an image tube IR viewer with an S-1 photocathode gets to 1.3 μ m for really bright light. Further into the IR, life is really miserable. Lead-salt vidicon cameras get to 2.2 μ m but have such strong memory effects that they are nearly useless for fine work. Mid-IR focal plane array cameras are very expensive, although the new bolometer array units are beginning to change this.

If you do use a camera, get an articulated mount such as a Bogen Magic Arm to hold it. Magic Arms are beautiful—they're about the size of your own arm, and articulate the same way, but can be twisted into any position and orientation and then frozen solid by 1/4 turn of a lever. They have 1/4-20 mounting holes, so they bolt right onto your optical table, and also come with an excellent clamp. Cameras are bad enough without adding pointing problems.

19.10.10 Test Techniques for Fiber Systems

There are fewer things to go wrong with fiber systems, but on the other hand, they are more pervasive and much harder to fix. Trouble from cladding modes can be eliminated by mandrel wrapping or potting the stripped fiber end in wax. Polarization problems you just have to live with, although switching to polarization-maintaining fiber helps a fair amount. Galloping etalon fringes are a sad fact of life, that only Faraday isolators plus angled fiber ends, or very low coherence sources can fix (the latter, at the price of FM-AM conversion). Broken or dirty fibers are pretty obvious, except in bundles; collimation errors you fix as usual. Microbending loss can be due to too-short wavelengths where there is a leaky mode, or to a poor mounting scheme (e.g., metal gooseneck armor).

The major effects of thermal and mechanical instabilities on fiber systems are excess loss in launching light into the fiber, microphonic effects due to wiggling launchers and acoustic microbending and fringes of a few percent P-V going by constantly.

If your launching efficiency is poor, look for defocus, misalignment, dirt, surface damage, or a bad cleave. Cleave problems include hackles (little chunks added or missing where the crack front took a sharp turn) and undesired facet tilts, which misdirect the light. A fiber microscope is an easy way to look for this. In connectorized systems, the fiber might have pulled out or not have been polished well enough; dust on a physical-contact connector can destroy the facet.

19.10.11 Test Techniques for Frequency-Selective Systems

The main things that go wrong with grating systems are misalignment, ghosts, and stray light. Misalignment shows up in broadened or asymmetric lines, low efficiency, and wavelength variations along the exit slit. Use a narrowband source and look for the light intensity peak in the output slit moving around with slight detuning. To look for stray light, use a filtered narrowband source, detune the instrument, and detect with a photomultiplier. The amount of stray light usually depends strongly on the illumination angle; it gets worse when you fall off the grating.

Colored glass filters tend to be fluorescent, and some will drift with time, temperature, and abuse, so try changing their order or rotating them to see if anything changes; the

desired signal should not change at all, because the transmittance of a string of filters should be the product of their individual transmittances.

19.10.12 Source Noise Problems

Laser noise always varies with position and angle. Try putting a black-painted razor blade, a bit of cleaved silicon, or a shard of broken mirror on a translation stage, and gradually vignette the beam while looking at the detected noise level—often it'll be nearly a minimum with an unvignetted beam and no sample, but get big when you cut off half the beam. Try it from a couple of different directions, too. Gas lasers with relatively poor mode selectivity (e.g., big argons) are especially prone to this; a spatial side mode that is close to oscillating will make the beam wobble around very badly. (Pick a knife edge that minimizes the scattered light, to avoid confusion from speckle or mode hopping.)

Side-mode problems in laser diodes are easier to spot, because the mode separation is big enough that you can use a grating. Try coming out of the grating near grazing incidence; the anamorphic reduction in beam width (and attendant angular spreading) makes the modes show up better over small distances.

19.10.13 Pointing Instability Problems

Pointing instability shows up best with a quad cell alignment aid (see Section 12.9.11) and an oscilloscope or data logger.

19.10.14 Source Pulling

Etalon fringes can pull your laser frequency; of course, real Fabry–Perots are the worst for this (especially optical spectrum analyzers), but even stray fringes will do it. Faraday isolators are the best remedy if mild misalignment isn't enough.

19.10.15 Misalignment

We talked about alignment a lot in Section 12.4. One quick test for misalignment is to examine a focused spot on axis; a properly aligned optical system shows no asymmetric aberrations (e.g., coma) on axis and the beam doesn't move sideways when defocused.

19.10.16 Etalon Fringes

If you're suffering from etalon fringes in some component, then tapping that element should provoke a response, whereas components further downstream should be less sensitive (upstream ones may be as sensitive or even more, because moving them will move the fringe pattern). Going round tapping things gently with a screwdriver handle or the eraser end of a pencil is a surprisingly good diagnostic test. Etalon fringes are notoriously temperature sensitive, so gentle heating with your hand, a hair dryer, or a small heat gun (the little plastic kind you use for heat shrink tubing) will reveal a lot too. Etalon fringes forming across an air gap are sensitive to wind, temperature, and especially small amounts of vibration. With experience, you can locate the problem pretty rapidly this way.

Multiple effects cancel sometimes—a change that makes the fringes less vibration sensitive is probably an improvement, even if the fringes look somewhat worse.

19.10.17 Thermal Drift

There are lots of things that can drift with temperature besides etalon fringes; besides the obvious problems of expansion and bending, which we talk about in Section 20.2 (available at <http://electrooptical.net/www/beos2e/thermal2.pdf>), the two most drift-prone optical parameters are interference filter bandpasses and the focal lengths of infrared lenses.

19.10.18 Environmental Stuff

Dr. Erwin Loewen of Richardson Grating Laboratories used to tell a story of a famous ruling engine that would only work in the wintertime. It was a massive cast iron and granite apparatus, on the ground floor, and with a carefully temperature-controlled room.

After a few years of getting only 5 months' service per year, they finally nailed it by asking some very pointed questions about the day that the engine started misbehaving. It turned out that it usually happened several hours after the first big spring rainstorm, which pointed to something expanding due to moisture—but the iron, granite, and concrete seemed like poor candidates. Finally they took up the floor, only to discover that the engine had been mounted on several inches of wood over the concrete. The humidity changes made the wood swell and shrink, which was enough to prevent accurate gratings being made. Replacing the wood with something impervious fixed the problem.

19.10.19 Take It Apart and Put It Together Again





If all else fails, try taking the optical system apart, cleaning and inspecting everything, and putting it back together. Sometimes you'll find some hidden blunder, for example, a mount hanging from one overtightened screw and three loose ones, or a machining burr that didn't get removed. Other times, you won't know what changed, but the problem will vanish (those ones are usually the flaky ones due to etalon fringes or scattered light). Of course, sometimes the problem will get worse, and even if it does go away, you shouldn't ignore the problem—it's going to recur in another unit, count on it.

APPENDIX: GOOD BOOKS


If you're going to build electro-optical systems for a living, you're going to have to change fields a lot. If you're still in school, use the opportunity to broaden your expertise—an optics or experimental physics student should learn signal processing and circuit design, and an electrical engineering student, classical and quantum mechanics. It's well worth devoting an extra semester to, if you can afford it—graduate students usually can. The rest of us need good readable books that don't demand total devotion.

The present author's usual way to start learning a new field is to get a good undergraduate-level textbook and read it like a novel. Pay special attention to getting the very basic concepts down perfectly; as we've seen in earlier chapters, our confidence in working in a new area depends entirely on how sure we are of our tools. For example, Fourier optics is an approximation, but the Fourier decomposition of an optical field is exact; in electronics, metal resistors are linear and have only Johnson noise, and capacitors subtract DC noiselessly. In this class also are the classical techniques of other fields. For instance, the use of UHF modulation to suppress mode hopping in diode lasers is an idea first used around 1930, in the quench circuits of superregenerative receivers. (Superregeneration is an amazingly cute idea—you can amplify the Johnson noise enough to drive headphones with a single low gain stage; see Terman, *Radio Engineering*.) Thanks are due to the assembled expertise of the Usenet groups sci.optics, alt.lasers, and sci.electronics.design for suggestions and criticism. Those are also excellent places to find electro-optical design help. One thing to remember about Usenet technical groups: the crisper the question, the more helpful the answer, in general.

Legend

	Introductory level: Accessible to technical people from outside the field.
	Intermediate: Advanced undergraduate level, some background required.
	Advanced: university-level preparation required.
	Available free online.

Mathematics

George B. Arfken and Hans-Jurgen Weber, *Mathematical Methods for Physicists*, 4th edition, Academic, 1995. Good standard undergraduate text on classical analysis, with the usual physicist's mathematical permissiveness; especially good on complex variables. 

- Carl M. Bender and Steven A. Orszag, *Advanced Mathematical Methods for Scientists and Engineers: Asymptotic Methods and Perturbation Theory*, Springer, 1999 (reprint of the 1978 edition). The best applied math book of all time: light on the rigor, but powerful and full of physical insight (and good graphs). Excellent problems. ✂ ✂
- Ronald N. Bracewell, *The Fourier Transform and Its Applications*, 3rd edition, McGraw-Hill, 1999. Bracewell teaches you how to think in Fourier space, which is hugely valuable for optics and electronics folk. ✂
- I. M. Gelfand and S. V. Fomin, *Calculus of Variations*, Richard A. Silverman, translator, Prentice Hall, 1963. Variational methods are a uniquely powerful way to get solutions to a lot of problems in electromagnetics and mechanics, which lead straight to useful algorithms. Short, accessible and has all the variational horsepower you're likely to need. ✂ ✂
- Roger A. Horn and Charles R. Johnson, *Matrix Analysis*, Cambridge University Press, 1987. Covers advanced linear algebra topics, such as singular value decomposition, condition numbers, and so on. It's a math book, all right, but it's nice and clear, and states results in algorithmically useful forms. ✂ ✂
- Athanasius Papoulis, *Probability, Statistics, and Stochastic Processes*, 3rd edition, McGraw-Hill, 1991. Readable and full of useful examples in statistical signal processing, for example, Wiener and Kalman filters, maximum entropy estimation, maximum likelihood tests, Monte Carlo methods, and detection of signals in noise. ✂ ✂

Mathematical Tables

- Milton Abramowitz and Irene A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, US National Bureau of Standards, 1964. All about practical uses of special functions. Contains low and high accuracy tables, integral representations, asymptotic forms, and polynomial approximations. ✂ 📊
- Izrail S. Gradshteyn and Iosif M. Ryzhik, *Tables of Integrals, Series, and Products*, 4th edition, Academic Press, 1994 (corrected reprint of the 1980 edition). A gigantic, well-organized, and highly accurate set of integral tables. ✂ ✂
- Daniel Zwillinger, *CRC Standard Mathematical Tables and Formulae*, 30th edition, Chemical Rubber Company, 1996. Smaller and more accessible than Gradshteyn and Ryzhik, this one has trigonometric and combinatoric identities and a much smaller table of integrals. ✂

Electromagnetics

- Roger F. Harrington, *Time Harmonic Electromagnetic Fields*, 2nd edition, IEEE, 1996 (reprint of the 1961 edition). Especially good on variational principles and the radiation reaction concept—powerful but less widely known tools for EM problems. ✂ ✂ ✂
- Roger F. Harrington, *Field Computation by Moment Methods*, 2nd edition, IEEE, 1993 (reprint of 1968 MacMillan edition). The method of moments is the standard way to solve complicated antenna and metal waveguide problems. This book gives a unified and easy-to-follow development. ✂ ✂ 📊
- John D. Jackson, *Classical Electrodynamics*, 3rd edition, Wiley, 1998. An excellent upper-level electromagnetics text, with particularly good chapter problems. ✂ ✂

- Lev D. Landau, Evgeny M. Lifshitz, and Lev P. Pitaevskii, *Electrodynamics of Continuous Media*, Pergamon, 1984. Gives another angle on electromagnetics; especially good on crystal optics. ✂ ✂
- Simon Ramo, John R. Whinnery, and Theodore Van Duzer, *Fields and Waves in Communication Electronics*, 3rd edition, Wiley, 1994. An advanced E&M text with unusual coverage that's often very helpful; things like conservation of complex power, calculating inductance and capacitance, antennas, transmission lines, and resonators. ✂ ✂
- Albert Shadowitz, *The Electromagnetic Field*, Dover, 1988. A "div, grad, curl" E&M textbook with an engineering angle and lots of very clear explanations, at an introductory level (Maxwell's equations don't show up till Chapter 11). Inexpensive. ✂
- Allen Taflov and Susan C. Hagness, *Computational Electrodynamics: The Finite Difference Time Domain Method*, 3rd edition, Artech House, 2005. A big book on the state of the FDTD art, with lots of details; explicit equations and algorithms you can code. ✂ ✂

Optics

- Nicolaas Bloembergen, *Nonlinear Optics*, 4th edition, World Scientific, 1996 (reprint of 1965 edition with updated references). The best books are usually written by the pioneers, and Bloembergen is an excellent example: clearly written and physically lucid. ✂ ✂ ✂
- Max Born and Emil Wolf, *Principles of Optics*, 6th edition, Cambridge University Press, reissued 1998. The standard work on optical theory: propagation, diffraction, coherence, aberrations, you name it. ✂ ✂
- V. A. Borovikov and B. Ye. Kinber, *Geometrical Theory of Diffraction*, IEEE, 1994. If you have to do GTD for real, this is the book. Not elementary, but clear and complete. ✂ ✂ ✂
- Robert F. Fischer, *Optical Systems Design*, 2nd edition, McGraw-Hill, 2007. More centered on lenses than Smith. Has a lot of hands-on stuff, from design to glass selection to tolerancing, and a great list of bloopers at the end. ✂
- Grant R. Fowles, *Introduction to Modern Optics*, 2nd Edition, Dover, 1989. Fowles provides a lot of good physical insight into optical phenomena. Inexpensive. ✂
- Jack D. Gaskill, *Linear Systems, Fourier Transforms, and Optics*, Wiley-Interscience, 1978. Develops the whole theory of 1D and 2D linear systems from scratch, then derives Fourier optics and CTF/OTF as consequences. More mathematical than some. ✂ ✂
- A. Gerrard and J. M. Burch, *Introduction to Matrix Methods in Optics*, Dover, 1994. Presents the mathematics of ABCD, Gaussian beam, and polarization matrix calculations in a clear and comprehensive way; Inexpensive. ✂
- Joseph W. Goodman, *Introduction to Fourier Optics*, 2nd edition, McGraw-Hill, 1986. A well-written brief introduction, whose elementary level limits it to the paraxial case. ✂
- Joseph W. Goodman, *Statistical Optics*, Wiley-Interscience, 1985. Very readable and complete: covers speckle, classical and quantum fluctuations, and the complete theory of imaging with partially coherent light, using fairly advanced statistical reasoning. ✂ ✂ ✂
- Eugene Hecht, *Optics*, 4th edition, Addison-Wesley, 2002. An introductory optics text, with lots of physical insight. Especially good on Fourier and physical optics. ✂

- Kenji Kawano and Tsutomu Kitoh, *Introduction to Optical Waveguide Analysis*, Wiley, 2001. All about numerically calculating modes and propagation in optical waveguides and fibers. Lots of details. ✂✂✂
- Rudolf Kingslake, *Lens Design Fundamentals*, Academic Press, 1978; and *Optical Systems Design*, Academic Press, 1983. These books are classics, reflecting the Conrady tradition of hand calculation in optical design. You can't run a numerical optimizer with a desk calculator, so Kingslake talks a lot about heuristics and intuitive methods, and this alone is worth the price of the books. ✂✂
- Miles V. Klein and Thomas E. Furtak, *Optics*, 2nd edition, Wiley, 1986. Compared with Hecht, this book has much more on geometric optics and aberration and less on physical optics. Some mathematical errors. ✂
- Virendra N. Mahajan, *Aberration Theory Made Simple*, SPIE, 1991. An accessible but oddly split treatment of wave and ray aberrations. A bit weaker on the connections and disconnects between them (e.g. the subtle changes in terminology). ✂
- Bahaa E. A. Saleh and Melvin C. Teich, *Fundamentals of Photonics*, 2nd edition, Wiley, 2007. Starts from the basics and covers most devices of interest. ✂
- Warren Smith, *Modern Optical Engineering*, 4th edition, McGraw-Hill, 2007. Good self-teaching book on all aspects of classical optics—rays, aberrations, diffraction, lens design, combining stock lenses. Accessible but not at all simplistic. ✂
- J. J. Stamnes, *Waves in Focal Regions*, Adam Hilger, 1986. A very complete treatment of wave propagation in all cases where diffraction is important: especially good on the fine points of the plane wave spectrum. ✂✂
- Pyotr Ya. Ufimtsev, *Elements of the Physical Theory of Diffraction*, Wiley, 2007. Ufimtsev's theory is the foundation of stealth technology—PTD is GTD taken one perturbation order further. (An old version is available at <http://www.dtic.mil>.) ✂✂✂📖
- US Department of Defense, *Optical Design: Military Standardization Handbook 141 (MIL-HDBK-141)*, US Defense Supply Agency, 1962. Everything from optical glass to atmospheric optics and lens testing. Still remarkably current. (Available at <http://www.optics.arizona.edu/optomech/references/Mil-hdbk-141.mht>.) ✂📖

Other Physics

- Lev D. Landau, Evgeny M. Lifshitz, and Lev P. Pitaevskii, *Statistical Physics*, 3rd edition (2 volumes), Pergamon, 1980. Especially good on black body radiation and the fluctuation–dissipation theorem, from a theoretical physics point of view. ✂✂
- Simon M. Sze, *Physics of Semiconductor Devices*, 2nd edition, Wiley, 1981. (The third edition (2006), isn't as good on the basics.) Gives detailed explanations of how semiconductor devices work; enough equations but few derivations. ✂✂

Circuits

- American Radio Relay League, *The Radio Amateur's Handbook*, ARRL, Annual. Covers modulation, signal processing, and wave propagation with lots of lore and heuristic arguments: accessible, not rigorous. Pre-1990 editions had lots of RF construction lore. ✂
- Harold S. Black, *Modulation Theory*, Van Nostrand, 1953. Lots of detail on AM, FM, time and frequency division multiplexing, Shannon's theorem, and entropy in modulation schemes, in a clear style, by the inventor of negative feedback. ✂✂

- Hendrik W. Bode, *Network Analysis and Feedback Amplifier Design*, Van Nostrand, 1945. Bode more or less invented network analysis, and the combination of clear explanation and conceptual rigor is unbeatable. Especially good on matching network design. ✂ ✂
- Hans Camenzind, *Designing Analog Chips*, <http://www.designinganalogchips.com>. By the designer of the 555 timer chip; an interesting read with a lot of good background on devices and simulation. ✂ 📀
- Ralph S. Carson, *High Frequency Amplifiers*, Wiley, 1982. A lucid explanation of S-parameter design of high frequency amplifiers, based on Smith charts. This is how you really do it. ✂ ✂
- Jiri Dostal, *Operational Amplifiers*, 2nd edition, Butterworth-Heinemann, 1997. Combines arm waving and rigorous mathematical treatment of high performance op amp circuits. Very readable and full of insight, and covers every op amp problem imaginable. Out of print; soft copies are available at various places on the Web but may be bootlegs. ✂ ✂ 📀
- Floyd M. Gardner, *Phaselock Techniques*, 3rd edition, Wiley, 2005. All about analog phase-locked loops, including frequency synthesis, PLL signal detection, telemetry, and Doppler tracking. Just the right amount of math so you can do it yourself. ✂
- Paul A. Gray and Robert G. Meyer, *Analysis and Design of Analog Integrated Circuits*, 2nd edition, Wiley, 1984. Lots of insight into BJT, FET, and IC performance and behavior. If you don't know how a transistor works, this is your book. ✂ ✂
- Paul Horowitz and Winfield Hill, *The Art of Electronics*, 2nd edition, Cambridge University Press, 1989. Indispensable if you design circuits, even if it's just the occasional op amp. A third edition is due in 2009. ✂
- Donald E. Lancaster, *Active Filter Cookbook*, 2nd edition, Newnes, 1996. Only slightly changed from the 1975 edition, but still very useful when you need to hack together a low frequency filter. ✂
- Robert A. Pease, *Troubleshooting Analog Circuits*, Butterworth-Heinemann, 1991. An entertaining and readable look at troubleshooting. Full of war stories, and applicable to a lot of other types of systems besides analog circuits. ✂
- Frederick E. Terman, *Radio Engineering*, 3rd edition, McGraw-Hill, 1947. A big, very good book on all aspects of analog signal processing and network design. ✂ ✂
- Frederick E. Terman, *Radio Engineer's Handbook*, McGraw-Hill, 1943. Full of old-but-useful tables, curves, and classical techniques for when the derivative modern treatments don't have enough detail. ✂ ✂
- Mac E. Van Valkenburg, *Reference Data For Engineers*, 8th edition, Howard W. Sams, 1995. Lots of practical stuff about design of RF, communications, and analog and digital circuits. This edition branches out from the previous radio-only focus, but still has a lot of great RF stuff. ✂
- Arthur Williams and Fred J. Taylor, *Electronic Filter Handbook*, 4th edition, McGraw-Hill, 2006. Perfect for people who occasionally need to design a custom active or LC filter but don't care about the math. Previous editions were better printed. ✂
- Jim Williams, editor, *Analog Circuit Design: Art, Science, and Personalities*, Butterworth-Heinemann, 1991; and *The Art and Science of Analog Circuit Design*, Butterworth-Heinemann, 1995. The lore of analog design, told in a playful and

sometimes whimsical style by masters of the art. The second book is especially good on high impedance, high frequency circuits—very useful in front end design. ✂

Anatol Zverev, *Handbook of Filter Synthesis*, Wiley, 1967. A big book of passive LC filter design tables and plots. Has unique designs (e.g. the equiripple group delay filter). ✂ ✂

Noise and Interference

Aldert van der Ziel, *Noise in Solid State Devices and Circuits*, Wiley, 1986. A brief, accessible, but complete introduction to noise in semiconductors and circuits, which is a surprisingly pretty subject. ✂ ✂

Ralph Morrison, *Grounding and Shielding Techniques*, 4th edition, Wiley, 1998. Morrison gives the physical causes of interference and spurious radiation, then shows how to get rid of it. ✂

Henry W. Ott, *Noise Reduction Techniques in Electronic Systems*, 2nd edition, Wiley, 1988. Solid on grounding, shielding, decoupling, filtering, and how to calculate noise and interference. Essential for low noise circuit designers. ✂

T. H. Wilmshurst, *Signal Recovery from Noise in Electronic Instrumentation*, 2nd edition, Adam Hilger, 1991. The gory details of lock-in detection, signal averaging, and so forth. ✂

Optomechanics

Eugene A. Avallone and Theodore Baumeister III, *Marks' Standard Handbook for Mechanical Engineers*, 10th edition, McGraw-Hill, 1996. Covers a wide field, even incineration and power generation, but the depth of detail and choice of material are excellent. ✂ ✂

Erik V. Oberg and Christopher J. McCauley, *Machinery's Handbook*, 25th edition, Industrial Press, 1996. Full of good stuff about machine shop practice. ✂

Stephen P. Timoshenko, *Theory of Plates and Shells*, 2nd edition, McGraw-Hill, 1959 (recently reprinted). A wonderful mixture of explanation and cookbook formulas on deflection, resonance, buckling, and deformation of structures, especially those that have to be stiff and light. ✂ ✂

Paul R. Yoder, *Opto-mechanical Systems Design*, 3rd edition, CRC Press, 2005. Covers the mechanical design of optical instruments, test jigs, and so on. Yoder has immense design experience, and the book repays close attention. ✂ ✂

Detection and Front Ends

Stephen B. Alexander, *Optical Communications Receiver Design*, Institution of Electrical Engineers, 1997. A good book on how to detect optical signals in the presence of noise, with attention to various modulation schemes. ✂ ✂

Eustace L. Dereniak and Glenn D. Boreman, *Infrared Detectors and Systems*, Wiley, 1996. Has a good discussion of most modern IR detectors, plus sections on noise, MTF, statistics, and general photodetector operation. ✂ ✂

Silvano Donati, *Photodetectors: Devices, Circuits, and Applications*, Prentice-Hall, 1999. A clear presentation of the physics and circuit characteristics of photodetectors. ✂ ✂

- Hamamatsu Photonics, *Photomultiplier Tubes: Basics and Applications*, 3rd edition. Not just product information, but how PMTs work and lots of good lore on the ins and outs of making good PMT measurements. ✂
- James R. Janesick, *Scientific Charge-Coupled Devices*, SPIE Press, 2001. A big book full of CCD basics and arcana: sometimes verbose but very complete. ✂
- Mark Johnson, *Photodetection and Measurement: Maximizing Performance in Optical Systems*, McGraw-Hill, 2003. How to make optical measurements, starting from simple concepts and progressing to the whole system. Especially good hands-on sections. ✂

Measurements and Systems

- Brian Culshaw and John Dakin, editors, *Optical Fiber Sensors Volume 3: Components and Subsystems*, Artech, 1996. A survey of fiber sensors and some of the underlying technology. Uneven in quality but good on fiber gratings and spectral measurements. ✂
- Robert Hanbury Brown, *The Intensity Interferometer: Its Application to Astronomy*, Halsted Press (Wiley), 1974. An epic technical war story of getting a completely novel type of telescope built, and its unique contributions. ✂
- Richard D. Hudson, *Infrared Systems Engineering*, Wiley, 2006 (reprint of the 1969 edition). The essential problems of infrared systems don't change much with time, and Hudson remains the best guide to IR system design. ✂ ✂
- Albert V. Jelalian, *Laser Radar Systems*, Artech House, 1992. Everything about detecting objects by bouncing laser beams off them. Jelalian is readable and clear. ✂ ✂
- David S. Kliger, *Ultrasensitive Laser Spectroscopy*, Academic Press, 1983. Now somewhat dated, but still very good on the techniques and difficulties of ultrasensitive optical measurements. Engaging and full of lore. ✂ ✂
- Marc D. Levenson and Satoru S. Kano, *Introduction to Nonlinear Laser Spectroscopy*, Revised edition, Academic Press, 1988. Good for coming up to speed in laser spectroscopy quickly; includes chapter problems and lab exercises. ✂ ✂
- Merrill I. Skolnik, *Radar Handbook*, 2nd edition, McGraw-Hill, 1990. More advanced signal processing and the problems and trade-offs of instrument design. Applicable far beyond radar. ✂
- Eric Udd, *Fiber Optic Sensors: An Introduction for Engineers and Scientists*, Wiley-Interscience, 1991. A bit long in the tooth now, but contains a fair amount of lore from experts. Nonfiber alternatives are not really given their due. ✂ ✂

Construction

- Daniel Malacara, editor, *Optical Shop Testing*, 3rd edition, Wiley, 2007. Covers all kinds of interferometric testing, as well as screen (Ronchi and Hartmann), star, profilometer, and other tests for optical surfaces. If it isn't here, you probably can't test it. ✂ ✂ ✂
- John H. Moore, Christopher C. Davis, and Michael A. Coplan, *Building Scientific Apparatus*, 2nd edition, Addison-Wesley, 1989. A very readable guide to building apparatus, with lots of stuff about basic optics, mechanics, vacuum, and electronics. ✂
- John Strong, *Procedures in Experimental Physics*, Lindsay Publications, 1986 (reprint of 1938 edition). A great hands-on book on glasswork, molding, vacuum, and so on. ✂

Lasers

- Samuel M. Goldwasser, *Sam's Laser FAQ*, <http://www.repairfaq.org/sam/lasersam.htm>. Everything about lasers, from basic physics to connector pinouts. A labor of love drawing on decades of nitty-gritty experience. ✂️📖
- Walter Koechner, *Solid State Laser Engineering*, 6th edition, Springer, 2006. The standard work on building solid state lasers—theory, materials, pumping and thermal strategy, *Q*-switching, mode locking, and nonlinear elements. ✂️✂️
- Motoichi Ohtsu, *Highly Coherent Semiconductor Lasers*, Artech House, 1991; and *Frequency Control of Semiconductor Lasers*, Wiley, 1996. Design and experimental data on how to stabilize diode lasers, for spectroscopy, frequency standards, and other uses. ✂️✂️
- Anthony E. Siegman, *Lasers*, University Science Books, 1986. Covers all aspects of lasers at a first-year graduate level; accessible and full of good intuitive examples. Due to its age, better on gas and solid state lasers than diodes. ✂️✂️

Digital Signal Processing and Numerical Analysis

- Forman Acton, *Numerical Methods That Work*, Mathematical Association of America, 1990 (reprint of 1970 Harper & Row edition). Salty, opinionated, and very lucid; full of lore, including a lovely section on what *not* to compute. ✂️
- E. Oran Brigham, *The Fast Fourier Transform*, Prentice Hall, 1974. The ins and outs of FFT algorithms. Readable and very worthwhile if you write any signal processing code. ✂️
- Philip J. Davis and Philip Rabinowitz, *Methods of Numerical Integration*, Academic Press, 1975. Covers all sorts of numerical integration, by way of hundreds of short discussions of the algorithmic ingenuity of each approach, which sheds a lot of light on numerical mathematics generally. ✂️✂️
- Dennis C. Ghiglia and Mark D. Pritt, *Two Dimensional Phase Unwrapping*, Wiley, 1998. Two-dimensional unwrapping is a ubiquitous and very thorny problem; Ghiglia and Pritt give the gory details for times when the seat-of-the-pants approach fails. ✂️✂️✂️
- Richard W. Hamming, *Digital Filters*, 3rd edition, Dover, 1997 (reprint of Prentice Hall 1989 edition). Covers nonrecursive digital filters from an intuitive point of view; good for coming up to speed rapidly. Inexpensive. ✂️
- J. F. Hart et al., *Computer Approximations*, R. E. Krieger, 1978 (reprint of 1968 Wiley edition). Computer-generated rational function approximations to special functions. ✂️
- Alan V. Oppenheim and Ronald W. Shafer, *Digital Signal Processing*, Prentice Hall, 1975. Old but good; fairly heavily theoretical. ✂️✂️
- William H. Press, Brian P. Flannery, Saul A. Teukolsky, and Steven Vetterling, *Numerical Recipes in C*, 2nd edition, Cambridge University Press, 1992. A pretty good numerical analysis book that comes with a lot of working code you can use right away. There's a third (C++) edition, but the second (C) edition is available free, and the C routines may be easier to integrate with your code. ✂️✂️📖
- L. R. Rabiner and B. Gold, *Theory and Application of Digital Signal Processing*, Prentice Hall, 1974. The applied follow-on to Oppenheim and Shafer. Not so mathematical, and with a fair amount of lore and good insights on when to use the different algorithms. ✂️✂️

- Anthony Ralston and Philip Rabinowitz, *A First Course in Numerical Analysis*, 2nd edition, McGraw-Hill, 1978. An undergraduate text with lots of examples, good for self-teaching. Ralston and Rabinowitz have no patience with “neat but useless” techniques. ✂
- Steven W. Smith, *The Scientist and Engineer’s Guide to Digital Signal Processing*, <http://www.dspguide.com>. An elementary DSP book with a wide coverage, lots of pictures and discussion, cookbook tables, and very little math. Excellent for getting your feet wet. ✂ 📖

Handbooks Worth Having

- J. S. Accetta and David L. Shumaker, *The Infrared and Electro-Optical Systems Handbook* (8 volumes), Infrared Information, 1993; also available for download from <http://www.dtic.mil>. All kinds of tables and discussion—not as handbooky as most others. ✂ 📖
- C. W. Allen, *Astrophysical Quantities*, 3rd edition, Athlone, 1985. Where else can you find out that the night sky has ~ 0.1 naked-eye visible stars per square degree? ✂ ✂
- Michael Bass et al, editors, *Handbook of Optics* (4 volumes), Optical Society of America and McGraw-Hill, 1995–2000. Volume I: Fundamentals, Techniques, and Design; Volume II: Devices, Measurements, and Properties; Volume III: Classical, Vision, and X-Ray Optics; and Volume IV: Fiber Optics and Nonlinear Optics. Truly gigantic—not the last word on any optical subject, but a good place to start for most. ✂ ✂
- Burle Industries (formerly RCA), *Electro-Optics Handbook*, 1974. The “RCA handbook” that started it all. Still very very useful, and free. ✂ 📖
- Chemical Rubber Company, *CRC Handbook of Chemistry and Physics*, CRC Press, Annual. A huge book full of tables, indispensable for mixed-technology systems like EO systems: the “rubber bible” has lots about materials, including enough optical properties that you may not need Palik and Ghosh. Recent editions are *much* better than the old ones. ✂
- G. W. C. Kaye and T. H. Laby, *Handbook of Physical and Chemical Constants and Some Mathematical Functions*, 15th edition, Longman, 1986. A real instrument-builder’s handbook, very useful when you have to choose materials or calculate heat transfer—it has obscure things like the specific heat and thermal conductivity of tantalum wire. ✂
- Edward D. Palik and Gorachand Ghosh, *Handbook of Optical Constants of Solids*, 2nd edition (4 volumes), Academic, 1997. Very expensive but sometimes invaluable. Lots of detailed critiques of how the constants were measured. ✂ ✂
- William L. Wolfe and George J. Zissis, *The Infrared Handbook*, Revised edition, ERIM, 1985. Lots of good tables and curves make this a useful handbook for designers. ✂ ✂

- Abbe, 319
- Abbe number *V*, 146
- Aberration, 336–340
 - astigmatism, 339, 341
 - calculation, 309
 - chromatic, 342–343
 - coma, 341
 - defocus, 337, 339, 341–342, 349
 - dependence, 340–341, 352
 - effect on SNR, 329, 344
 - field curvature, 336, 341
 - grating, 319, 346
 - manufacturing, 340
 - nomenclature, 337–338
 - oblique, 339–340
 - ray, 336–337, 341
 - scanning, 328–329
 - seidel, 341
 - and SNR, 344
 - spherical, 310, 312, 339, 341, 343n1
 - theory, 337, 344
 - wave, 310, 312, 335, 337, 340–341
- Abramowitz, Milton, 331
- Absorption, 39, 92, 215–216, 515
- Acousto-optic, 381, 383, 561
- Acousto-optic cell, 258–259
 - shear258– 260
- Acousto-optic deflector (AOD), 260, 563, 589
- Acousto-optic modulator, 234n1, 717
- Adams, Douglas, 509, 738, 738n1
- Adder, saturating, 680
- Adhesive, 441–443
 - bond line, 442
- Adiabatic coupling, 268
- Adjustment
 - interacting, 411
 - orthogonal, 431
- Admittance, 512–513, 526, 618, 703
- Ahab, Captain, 354
- Airplane rule, 391
- Airy pattern, 32, 347
- Alasing, 41
- Algorithm
 - optimal, 687
 - testing of, 671, 680
 - unwrapping, 685–686
- Aliasing, 654–655
- Alignment, 421–425
 - beam, 424
 - clamping, 422–423
 - design for, 422
 - detector, 419–420
 - fiber, 431
 - invisible lasers, 429
 - optical, 425–426
 - orthogonal, 425–426
 - spatial filter, 420–421
 - spectrometer, 421
 - too accurate, 419, 433
- AM-PM conversion, 466, 499–500, 547, 579, 592
- Amplifier
 - BJT, 533, 702
 - cascade, 536
 - common-base, 536, 701
 - common-base, 701
 - common-collector, 536
 - common-emitter, 536, 619, 711
 - current noise, 472
 - current, 103, 733
 - differential, 542–543
 - EDFA, 290, 715
 - emitter follower, 535–536, 542
 - fiber, 290
 - front end, 542, 561, 575, 690
 - gain bandwidth, 501, 642, 699
 - gain-bandwidth product (GBW), 472, 693
 - inverting, 472–474, 580
 - isolation, 543–544
 - limiting, 546–548
 - lock-in, 548
 - loop gain, 573
 - low noise, 471, 567, 757
 - noise figure, 737
 - noise, 471
 - nonlinear, 555, 740
 - nonlinearity, 555
 - operational, 542
 - phase margin, 539, 695

- PMT, 109, 120n1
- RF, 543
- settling time, 545–546
- slew rate, 545
- specs, 540
- stability, 544–545
- track-hold (T/H), 550–551
- transimpedance, 693–714
- transistor, 530–539
- tuned, 539
- unity gain stable, 699–700, 708
- voltage noise, 703–704
- Amplitude, instantaneous, 478
- Analog devices, 549, 596, 624n1, 709
- Analogue divider, 95
- Analogue-to-digital converter (ADC), 45n1, 551–554
 - $\Delta \Sigma$, 552–554
 - dynamic range, 555
 - error, 551
 - flash, 551
 - missing codes, 551, 554
 - noise, 553
 - nonlinearity, 554, 555
 - pathology, 532
 - resolution, 553–555
 - successive approximation, 552
 - testing, 555–556
- Analytic signal, 451–453
- Analytical models, 9, 409
- Analyzer, polarization, 88n2, 208n1, 760
- Anamorphic magnification, 235–237, 242, 311
- Angle, critical, 7, 64, 160, 164, 175, 211, 222–223, 287
- Angle tuning, 189–190
- Angular spectrum, 323, 346, 350–351
- Anodizing, 200, 419
- Answer
 - right, existence of, 659, 666
 - wrong, reasonable-looking, 605, 649, 758
- Antenna, 282, 343, 471, 526–527
 - phase efficiency, 343
- Aperture, 18–19, 346–347, 550n3, 551
- Aperture stop, 18–19, 24, 169, 177, 318, 327
- Approximation
 - scalar, 24, 267–268, 325
 - sleazy, 9, 596–597
- ARRL Handbook, 495, 509, 527, 560
- Art, black, 154, 186
- Art of Electronics, 468n3, 509
- Assembly procedure, 425
 - errors, 425
 - height adjustment, 427
 - optical, 425–426
 - sequential, 635
 - temporary, 442
- Assumptions, 746
- Astigmatism, 82, 310, 314, 765
- Asymptotic behavior, 22, 482, 681
- Asymptotic series, 331
- Asymptotic theory, 9, 316
 - weirdness of, 15
- Athermalization, 147, 168
- Attenuator, 254n1, 283, 488, 528
- Autocollimation, 37, 428
- Autocorrelation width, 20–21
 - statistical, 479, 682
 - time, 157, 479, 667, 682
- Automatic gain control (AGC), 486, 599
- Automatic level control (ALC), 86, 589, 599
- Avalanche, 68, 113, 607
- Avalanche photodiode (APD), 113–115
- Average/Averaging
 - ensemble, 35, 60, 478
 - fringe, 36, 157
 - spatial, 62, 100, 683–685
 - time, 209, 500n1
- Back focus, 163, 167, 281
- Background reduction, 46, 354, 376
 - additive, 355–356, 360
 - coherent, 360
 - correction, 645–646
 - diffuse, 237
 - drift, 502
 - fluctuation, 134
 - incoherent, 360, 376
 - light, 134, 137
 - noise, 137, 355–356
 - phase noise, 88, 378
 - reducing, 408
 - reduction, 46, 354, 376
 - resistant, 649
 - sky, 43
 - thermal, 98–99, 134, 134n2, 733
- Backlash, 434–435
- Baffles, 174, 198–203
 - band-limited, 199
 - design, 198–199
 - diffraction, 199
 - flat black, 203
 - ink dot, 200
 - shiny, 203
- Bandwidth, 455
 - 1-Hz, 472, 474
 - 3-dB, 40
 - bode limit, 682
 - cancellation, 725
 - detection, 49
 - electrical, 382
 - filter, 48–49, 409, 489, 497, 503, 507
 - finite, 11, 61–63, 467, 654
 - FM, 462
 - fractional, 241, 494
 - front end, 405
 - IF, 487, 499, 547, 599

- improving, 718
- LED, 142, 304
- limitation, 701, 753
- loop, 372, 375, 501, 564–565, 574, 588, 601, 694
- matching network, 718–719
- matching, 710
- measurement, 477n1, 562–563
- Miller, 701
- narrowing, 502, 752, 758
- noise, 45n1, 455, 457, 480, 494
- Nyquist, 655
- one-sided, 40, 457
- optical, 61, 114
- oscilloscope, 440, 748
- PLL, 499, 587
- RC, 692, 701, 718–719
- vs sensitivity, 293
- vs settling time, 489, 545–546
- Shannon limit, 681–682
- tradeoffs, 673
- transimpedance, 694–696, 701
- transit time, 49–50, 260
- Base plate, 418
- Baseband, 358–359, 455, 718, 725
 - filtering, 498, 596
- Beam dump, 155, 198–203
- Beam separator, 231
- Beam walkoff, 4, 213, 217, 314
 - alignment, 89, 432
 - cleanup, 345–348
 - collimating, 426–427
 - comparison, 47, 723–727, 725n1, 729
 - focused, 37–38, 252, 282, 764–765
 - Gaussian, 12–14
 - monochromatic, 233
 - noncollimated, 156–157, 418
 - ‘nondiffracting’, 179, 314, 402
 - pencil, 15, 311, 314, 318, 323
 - quality, 343, 427, 429
 - signal, 362–363, 723–726, 725n1
 - transmit, 47–48, 250, 344, 360
 - uniform, 179, 345–348, 760
- Beamsplitter, 158–159, 164, 193–195
 - cube, 164
 - pellicle, 159
 - plate, 158–159
 - polarizing, 193–195
- Bear, Edward, 415
- Beat length, 270, 274, 288, 304, 366
- Beating, 75
- Beer check, 406
- Belt and suspenders, 155, 415
 - wearing, 400–401
- Bessel expansion, FM, 305
- Bessel function, 267, 384, 402, 460–461
 - modified, 267
- Bessel null (FM/PM), 461n1
- Binary optics, 244
- Birefringence, 152–153, 211–212, 277–278
 - bend, 277–278
 - negative, 222
 - stress, 152, 196, 221, 260–261, 762
- Black
 - designer, 200
 - flat, 203
 - stick-on, 200
 - wax, 201
- Black body radiation, 56–57
- Black glass, 201
- Black level (video), 126–127, 370, 376
- Black light, 380, 441
- Black paint, 155, 199
- Black plastic, 200
- Black surface, 199
- Black wax, 201
- Block diagram, 483, 485, 493, 544, 564, 566
- Blunders, 659
 - design, 762
- Blur, 369
- Bode plot, 573
- Bode, H. W., 492n2, 719n1
- Bond line, 442
- Bootstrapping, 602, 705–708, 757
- Borescopes, 176
- Born, M., 7n1, 15n1, 31, 60, 211n2, 227, 313
- Boundary condition, 242, 266–267, 320–321, 323,
 - 342, 523, 612
 - electromagnetic, 10n2
- Bracewell, Ronald, N., 19, 22, 449, 650
- Bragg cell, 260, 362, 383, 598
- Bragg diffraction, 243
- Bragg grating, 243–244
 - photorefractive, 271
- Brewster’s angle, 6–7
- Bright field, 350, 359–360
- Bright field rule, 362, 377
- Brown, Sarah, 644
- Buffer (FIFO), 606
- Bug zapper, 561, 567, 579, 585
- Bulb, tungsten, 58–59
- Bulletproofing, 607–611
- Calculation, back-of-the-envelope, 561
- Calculus of variations, 315, 316n1, 426, 495
- Calibration, 590–592
 - detector, 590–591
 - errors, 564
 - limiter, 592
 - phase, 590–591
- Camera, 127
 - CCD, 214, 336, 427, 429, 763
 - chopping up, 381
- Capacitance multiplier, 536–537, 736
- Capacitance, 99
 - detector, 40–41, 99, 690
 - fringing, 616

- Capacitance (*Continued*)
 - input, 709–710
 - load, 757
 - photodiode, 699, 708, 710, 731, 753
 - resonating away, 717
 - stray, 518, 616
- Capacitor, 67n2, 512–522
 - effective series inductance (ESL), 515
 - effective series resistance (ESR), 515
 - offset voltage in, 753
 - soakage, 515–516
 - variable, 494, 516, 637
- Capture effect, 546
 - in rectifying detectors, 548
- Carrier to noise ratio (CNR), 477, 562, 690
- Cascode, 536, 702, 705, 710
- Cauchy's integral formula, 452
- Causality, 489–490
- Cavity dumping, 74
- Cavity ring-down, 294, 384
- Centration, 65, 422–423
- Cepstrum, 645
- Characteristic impedance, 522, 527, 543, 577, 614
- Charge-coupled device (CCD), 120
 - electron multiplying, 124–125
 - fixed pattern, 124, 126, 128
 - full well capacity, 122, 125–126, 130
 - noise, 124–126
 - nonlinearity, 123
 - TDI, 124
- Chartier, Émile, 402
- Chebyshev, 251, 488, 491, 494, 593–594
- Chinese hat function, 330–331
- Chirp, 77–78
 - zero, 257
- Chopping, 381
- Chromophore, 298, 357–358
- Circuit board, 613
 - ground plane, 618–619
 - hook in, 614
 - SMT, 635
 - strays, 612–617
- Circuit model
 - Norton, 468–469, 472
 - Thévenin, 468–469, 472
- Circuit strays, 612–617
- Circulator, 231
- Clairvoyance, 741
- Cleaning gratings, 445
- Cleaning lenses, 444–445
- Cleaning, 443–446
 - Peel, 442, 446
 - when, 444
- Clock, 552–553
 - stopped, accuracy of, 157
- Coating, 180–207
 - angle tuning, 189–190
 - antireflection, 184, 190
 - bandwidth, 189, 190, 192, 195
 - beam splitter, 193–195
 - black, 198, 200
 - dielectric, 184–186
 - drift of, 184, 195, 200
 - gold, 189
 - metal, 184
 - mirror, 182–184, 195
 - multilayer, 187–188
 - narrowband (V), 193–194
 - protected Al, 189–191
 - stagger tuned, 191–193
 - theory, 186–195
 - transmissive, 184–186
- Coaxial cable, 257n1, 285, 488n1, 491, 522, 525–526, 623
- Cockroft-Walton generator, 110
- Coefficient of coupling, 521, 642n1
- Coherence, 59–63
 - area, 31, 60
 - complex, 60
 - fluctuations, 61, 290, 292, 405, 730, 759
 - length, 60, 63–65, 72
 - partial, 61
 - spatial, 62, 64–66, 68, 73
 - temporal, 60, 62, 64, 68
 - time, 60, 62
- Coherent detection, 34–36
- Coherent lidar, 47n1, 47, 134, 730
- Collimation, 426–427
- Collimeter, 427
- Colour correction, 198
- Coma, 341
- Complex variables, 695
- Complexity, 401
- Component, 233–261, 419, 539–548, 763–764
- Compression, 463, 562
- Condensers, 69–71
- Conductance, 448, 468–469, 513
- Confidence level, 38, 48, 362
- Conformal map, 594
- Conjugate points, 166
- Conjugate variables, 19, 336, 346
- Construction, 612–643
- Continuous wave (CW), 453
- Contrast, 352–353
- Control, 504–508
- Converter, digital-to-analogue (DAC), 548–550
- Convolution, 19–20
 - circular, 672–673
 - and time invariance, 450
- Corner reflector, 163–164
- Correction, 173, 645–647, 649–666
- Correlation, 20, 682
- Cost-reduction, 743
- Counter
 - Johnson, 493, 558
 - walking-ring, 499, 558, 590

- Counting statistics, 41–42, 359–360, 362, 689
- Coupled mode theory, 243, 271
- Coupler, 288, 530
- Coupling
 - capacitive, 750
 - evanescent, 7, 288
 - inductive, 520–521
 - space variant, 271
 - transmission line, 617
- Crest factor, 475, 642
- Critical angle, 7
- Cross correlation, 20, 60, 500n1, 683
- Cross-dispersion, 236, 240
- Crosstalk, 77, 115, 263, 525, 544
- Cryogenic cooling, 108, 112, 117, 128, 138, 415
- Current mirror, 358, 538–539, 538n1, 735n1
- Customer, 390–391, 393–395, 743–744

- Damage threshold, 150, 200, 215–217, 222–223, 258
- Dark field, 40, 350, 359, 361, 373, 444
- Data sheet, 123, 539
- Data windowing, 659–664
- Databook, 714
- Datasheet, 540n1, 700
- Dead time, in photon counting, 111
- Debugging, 744–745
 - assumptions, 744, 746
 - blunder, 762, 768
 - electronic, 745
 - onion problem, 416, 755
 - optical or electronic, 763
 - overkill, 405, 624, 757
 - source noise, 131–132, 767
 - system, 749
 - and troubleshooting, 422, 744–745
- Debye approximation, 14–15
- Decibels, 39
- Decimation, 678
- Deconvolution, 675–676
- Decoupling, 518, 558, 753, 755, 757
 - power supply, 636, 755
- Defocus, 332, 437
- Delay discriminator, 274, 296, 303, 586–587
- Delay line, 266, 300, 302, 500, 586–587
- Delta function, 22, 24, 452–453
- Demo Karma, 414
- Demodulation, 281, 357, 453–462
 - I/Q , 457
- Demodulator
 - FM, 462, 476, 501
 - phase, 305, 500, 599
- Depolarization, 302
- Depolarizer, 229–230
 - rotating, 229
- Descanning, 250
- Design parameters, 403
 - for alignment, 410–413
 - blunder, 762
 - changes, 198, 413, 741
 - common-centroid, 533
 - conceptual, 1, 45, 47, 402, 742
 - hack, 413
 - hierarchical, 743
 - instrument, 402–407
 - interface, 568
 - manufacturable, 186
 - simple, 114
 - strategy, 396
 - for test, 627
- Detection, 33
 - amplitude, 591–592
 - coherent, 34–36, 138
 - coincidence, 384
 - differential, 735n1, 765
 - low-level, 486
 - in noise, 41, 467, 479
 - non-optical, 374
 - optical, 91–144
 - photon counting, 374, 482–483
 - pulse, 501–502, 580–581
 - square law, 33, 92–94
 - strategy, 1, 46, 136, 352, 565
 - successive, 584
 - threshold, 362, 482
 - false alarm rate, 481
 - time-gated, 138, 503
- Detector
 - amplitude, 591–592
 - avalanche photodiode (APD), 113–115
 - CCD, 121, 127–128, 367, 370, 376, 415
 - charge coupled (CCD), 120
 - diode, 499, 578, 582, 584, 591, 642
 - emitter, 578–579
 - envelope, 579–580, 590
 - FM, 500, 544, 586
 - HgCdTe, 128
 - indium antimonide (InSb), 108
 - indium arsenide (InAs), 98, 108
 - IR, 101, 108, 118, 294, 471, 733
 - logarithmic (DLVA), 384
 - logarithmic, 583–585
 - mercury cadmium telluride (HgCdTe or MCT), 108
 - nonlinearity, 498
 - phase, 501, 548, 587–592
 - photomultiplier (PMT), 92, 109–110
 - platinum silicide (PtSi), 108, 128
 - PLL, 499
 - position-sensitive, 105, 721
 - pulse, 501, 580–581
 - shadow mask, 106–107, 399
 - signal, 577–587
 - specific detectivity (D^*), 97
 - split, 44, 103–104
 - square law, 33, 92–94

- Detector (*Continued*)
 - synchronous, 476n1, 579
 - thermal, 117–118
- DFT, blunders using, 659
- Diabatic, 197
- Dichroic, 127n1, 216, 216n1
- Dielectric absorption (soakage), 515
- Dielectric, 184, 514
 - isotropic, 180
 - loss tangent, 9
- Differential interference contrast (DIC), 221, 364
- Differential pair, 537
- Diffraction grating, 233, 239, 244, 419n1
 - polarization dependence, 234
- Diffraction rings, 199, 347, 432
- Diffraction, 319–336
 - Bragg, 243
 - geometrical theory (GTD), 316–317
 - physical theory (PTD), 316n3
- Diffuser, 204–207
 - holographic, 65, 207
- Digital signal processing, 4n1, 550n1, 650, 671, 678n1, 689n2
- Digital to analogue converter (DAC), 548–550
 - dynamic range, 557
 - Glitch, 549–550
 - multiplying, 549
 - pathology, 532
- Digitizer, 103, 548–558
 - dynamic range, 557
- Dimensional stability, 249, 423–424
- Diner syndrome, 407
- Diode equation, 101, 530, 712
- Diode laser controller, 595–596
- Diode laser, 75–83
 - collimating, 82
 - collimation, 82
 - control, 89–90
 - distributed Bragg reflector (DBR), 77–78
 - distributed-feedback (DFB), 77–78
 - drift, 64
 - gain-guided, 75, 82
 - index-guided, 82
 - mode hopping, 84–85
 - modulation, 81–82
 - spectrum, 52–54
 - stability, 78–80
 - tuning, 78
 - visible, 77
- Diode
 - superluminescent, 65
 - zero-bias resistance, 101
- Diplexer, 595, 755
- Direction cosines u and v , 12, 323–325, 333–334, 338
- Discouragement, 388
- Discrete Fourier transform (DFT), 63, 657–659
- Discriminator
 - delay, 274, 296, 303, 586–587
 - delay line, 266, 300, 587
- Dispersed aqueous graphite (DAG), 113
- Dispersion, 271–272
- Distortion, 341
 - cross-modulation, 463
 - geometric, 124, 251, 310, 370, 427
 - intermodulation (IMD), 464–465
- Divider, 372
 - analogue, 601, 723
- Doppler broadening, 2, 369, 383
- Doppler shift, 47, 49, 250, 300, 344, 357, 369, 666
 - scattering, 563
- Dostal, Jiri, 472, 571
- Double-conversion, 466
- Drawings, 26, 407, 413, 747
- Drift, 649
 - fringe, 305
 - polarization, 274, 306
- Drinking straw test, 30, 323
- Duck test, 134
- Duplexing, 230
- Duty cycle, 138, 138n1, 251
- Dynamic range, 557
 - compression, 359
- Ear, sow's, 681
- Early effect, 534, 536, 711
- Ebers-Moll model, 533, 537
- Edison, Thomas A., 398, 398n1
- Efficiency, 121, 405
- EG&G PARC, 734
- Eigenmode, 232, 274
- Eigenvalue, 228, 266–267, 269, 269n1, 680
- Eikonal equation, 313–314
- Electro-optic effect, 254–255
 - linear (Pockels), 254
 - quadratic (Kerr), 254
- Electromagnetic compatibility (EMC), 620
- Electromagnetic interference (EMI), 263, 620
- Electron multiplication, 92, 109, 119, 128
 - CCD, 124–125
- Emissivity, 56, 56n1
 - tungsten, 58
- Emitter degeneration, 534–535
- Empire State Building, 354–359
- Endoscope, 176
- Engineering changes, 407, 626n1, 738, 739
- Engineers, 449, 451, 492, 561–564, 731n1, 741–742
- Envelope, 13, 58–59, 111, 113, 324, 456–457, 459–461, 494, 757
- Equation, diode, 101, 530, 712
- Equipartition of energy, 470
- Equivalent width, 20–21
- Erbium doped fiber amplifier (EDFA), 290, 715
- Ergodicity, 478, 667

- Error
 - dynamic, 556
 - propagation of, 371
 - RMS, 170, 505, 645
 - roundoff, 650, 652n1, 671, 674, 680
 - standard, 645
 - systematic, 46–47, 555
- Error bound, 111, 678
- Error budget, 507, 566–568
- Error signal, 365, 375, 575–576, 604
- Etalon fringes, 154–155, 266, 413, 762, 767
- Étendue
 - acousto-optic cell, 258–259
 - birefringent polarizer, 217–218, 221
 - conservation of, 32, 288, 318
 - and mismatch, 32
 - Pockels cell, 762
 - walkoff plate, 217
 - of wave plate, 224
- Evanescent wave coupler, 288
- Exercises, chapter problems, 125, 413
- Experiment
 - loses to theory, 747
 - numerical, 683
 - plan, 415
- Extinction sum rule, 360

- Fabry-Perot, 37, 279, 297, 301–302, 436
- Factor of 2, 156–157, 468n2, 718n1
 - mysterious, 55
- Fads
 - optical, 402
 - signal processing, 401
- Fallback position, 400
- False alarm rate, 48–50, 481, 503, 747
- Faraday cage, 626
- Faraday effect, 215, 298
 - quenching, 298
- Faraday isolator, 85, 89, 231, 280, 291, 766–767
- Faraday rotator, 230–231
- Faraday-rotator mirror (FRM), 232, 264–265, 292, 302
- Fata Morgana, 316
- Feedback
 - AC-coupled, 599
 - bandwidth, 601
 - DC, 599
 - emitter degeneration, 534–535
 - frequency compensation, 571–573
 - loop gain, 573
 - matched, 603
 - negative, 571
 - network, 571–573, 575–577, 589, 597
 - nonlinear, 603
 - nulling, 598
 - stability, 571
 - T-network, 582
 - unstable, 565, 574
- Feedback loop, 532n1, 571–577
 - windup, 575–576
- Feedback stabilization, 410
- Feedforward, 602, 604
- Felgett advantage in FTIR, 368
- Ferrite bead, 518, 624, 753, 755
- Fertilizer, 366
- Feynman, Richard, P., 390, 390n1
- Fiber
 - holey, 299
 - photonic crystal, 299
 - polarization-maintaining (PM), 273–274
- Fiber amplifiers, 290
- Fiber bundle, 274–275, 297
- Fiber connector, 263, 286–287
- Fiber-coupled device, 287
- Fiber device, 287–292
- Fiber illuminator, 70, 353
- Fiber interferometer, 210, 232, 299–300
 - Fabry-Perot, 299–301
 - Mach-Zehnder, 299–300
 - Michelson, 299–300
 - Sagnac, 300–301
- Fiber modes, 8, 282, 285
- Fiber optic, 262–308
- Fiber-optic faceplate, 275
- Fiber sensor, 307–308
 - intensity, 293–295
 - two-mode, 300
- Fiber stretcher, 291, 304–306
- Field angle, 18–19
- Field-programmable gate array (FPGA), 558, 568, 607, 628, 680
- Fields
 - vector addition, 10
- Fienuip, J. R., 336, 687, 687n2
- Figure error, 150, 162, 421
- Filter
 - acousto-optic, 258–259
 - all-pass, 491–492
 - band pass, 593
 - band stop, 593
 - bandwidth, 435
 - banks, 502–503
 - Bessel, 491
 - boxcar, 41
 - Butterworth, 46, 593, 595, 753
 - cascaded, 488
 - Chebyshev, 251, 488, 593–594
 - choosing, 494–495
 - coloured glass, 150, 152, 156, 196–198, 374, 766
 - deconvolution, 674, 676
 - defuzzing, 127, 214
 - delay equalized, 494, 671
 - design, 491, 674
 - digital, 553, 671–674
 - noise gain, 664
 - electrical, 19, 91, 345

- Filter (*Continued*)
- elliptic, 491, 594
 - equiripple, 45, 491, 494, 594
 - FIR, 648, 670–671, 671n3, 673–674, 684
 - glass, 152, 196
 - group delay, 45, 491, 494
 - high pass, 593
 - Hilbert transform, 492–493
 - holographic, 198
 - interference, 195
 - inverse, 675–676
 - IR-blocking, 763
 - LC, 495, 592–593
 - linear phase, 467, 491, 495, 594
 - low pass, 489, 494, 673
 - matched, 495–497
 - neutral density, 183, 197, 347, 747
 - optical, 58, 137, 196, 235, 374
 - optimal, 49, 453, 674
 - order, 491–492, 671
 - phase, 467, 492, 495, 594
 - power supply, 515
 - pulse, 641
 - realizability, 595
 - Sallen-Key, 752–753
 - selectivity, 485
 - settling time, 545–546
 - shape factor, 488, 762
 - spatial, 18, 282, 345, 420
 - transmittance, 43, 767
 - tuning, 641
 - Wiener, 676
- Finesse, 37
- Fingerprints, 441, 443–444, 636
- First-in, first-out (FIFO), 606
- Fizeau wedge, 152, 427
- Flashlamp, 67–68
- Flatband, 467, 601, 706
- Flip-flop, resynchronizing, 559
- Fluctuations, 370–372
- Flying-spot system, 159, 349
- Focal distance, 167, 169
- Focus
- depth of, 13, 23, 319, 324, 402
 - sagittal, 239, 764
 - tangential and sagittal, 764
- Focusing, 174, 428–430
- Foucault knife-edge test, 175
- Four-port device, 37, 291
- Fourier imaging, 23–24, 318
- Fourier optics, 9, 235, 325, 416
- Fourier series, 21n1, 651–652, 656–657
- Fourier space, 11, 23, 450–451
- Fourier transform infrared (FTIR) spectrometer, 368
- Fourier transform plane, 17–18, 24–25, 36, 318, 340, 346, 351, 411
- Fourier transform, optical, 19, 426
- Fourier
- decomposition, 769
 - space, 11, 23, 450–451
- Frame transfer (FT) CCD, 121
- Fraunhofer, Joseph von, 145, 147, 325, 429
- Free spectral range (FSR), 37, 85, 195, 301
- Frequency
- cutoff (f_T), 534
 - instantaneous, 157, 453, 461, 546, 587
 - modulation, 305, 378, 462
 - plan, 486, 505, 562
 - pulling, 292
 - spatial, 11, 236, 271, 330–331, 334–336, 351, 353
- Frequency compensation, 571–573, 695–696
- lead-lag, 574–575
- Frequency conversion, 45, 483–487
- Frequency doubler, 564
- Frequency mixing, 363, 458–459
- Frequency multiplier, 463–464
- Frequency response, 665, 674–675, 691, 694
- Frequency synthesizer, 562, 589, 599, 749
- Fresnel coefficients, 5–6, 187
- Fresnel formulae, 4n1, 5, 7, 180, 183–184, 187, 208, 522n1
- Fresnel number, 14–15, 178, 324–325
- Fresnel rhomb, 164–165, 180, 209, 213
- Fresnel zones, 14–15
- Friedman, Edward, 354, 402n2, 689n1
- Fringe
- averaging, 36
 - visibility, 34, 60, 300–301, 384, 459
- Fringes
- etalon, 154–155, 266, 762, 767
 - white light, 36, 218, 221, 303–304, 426, 435
- Fringe surfing, 305–306, 375, 377
- Front end, 405, 688–737
- laser noise canceller, 724
 - low noise, 699, , 719
 - noise canceller, 723–725
 - ratiometric, 724, 726
- Full width at half maximum (FWHM), 37, 50, 56, 301, 304, 520, 562, 661, 717
- Function
- band-limited, 331, 605, 654, 657, 673
 - Chinese hat, 330–331, 334
 - with compact support, 651
 - Dirac (δ), 19
 - Green's, 320–322, 489n1, 686
 - Heaviside unit step, 19
 - instrument, 678–679
 - periodic, 651, 672
 - point-spread, 253
 - pupil, 276, 328
 - sampling (Cyrillic capital *sha*), 497
 - sinc, 240–241, 660–661, 663–664, 683
- Fused silica, 146–147, 149, 151, 156, 262

- Gain, insertion, 487–488
 Gambler's ruin, 204
 Gardner, Floyd, 588–589, 599
 Gauss, K. F., 11–12, 167
 Gaussian beam, 12–15
 Gaussian imaging, 311
 Gaussian statistics, 479–481
 Geometrical theory of diffraction, 316–317
 Ghosts, 237
 grating, 237
 Gibbs phenomenon, 660
 Glass
 annealed, 153
 BK7, 154, 183, 338
 Black, 201
 crown, 146–147, 151, 171–172, 191, 196, 447
 filter, 196–197, 447
 flint, 156, 171
 float, 160
 ground, 206–207
 opal, 205
 polarizing, 216
 properties, 279
 tempered, 152
 window, 215, 279, 764
 Glitch, DAC, 750
 Glue, 419, 441
 Gold plating, 568
 Goodman, Douglas, S., 349, 351
 Goodman, Joseph, W., 59, 227n2
 Goos-Hänchen shift, 8
 Gotcha
 insidious, 63
 nefarious, 698
 Gotchas, 540
 Gouy phase shift, 12
 Graduate student, 38, 274, 397, 467, 502, 723
 Grass, 354–359
 Grating equation, 177, 234
 Grating
 blazed, 243
 Bragg, 243–244, 289–290, 295–296
 concave, 239
 efficiency, 242
 holographic, 238–239
 order, 162, 240, 368, 670
 photochemical, 373
 plane, 233–235, 237, 242, 289–290
 reflection, 237–238
 ruled, 238
 transmission, 236, 237–238, 251
 Gray and Meyer, 509, 533
 Gray, Paul, A., 534n1
 Grease
 fingerprint, 443
 nose, 443n1
 Greedy optimization, 741–742
 Grid dip meter, 637
 GRIN rod, 176, 285
 Ground loops, 621–622
 Ground plane, 618–620
 Ground truth, 397, 407, 409
 Ground, virtual, 542, 693
 Group delay, equalizer, 491–492, 594
 Group velocity, 4, 8, 214, 456, 491, 494
 Hacks, 136–144
 Hair dryer, 762, 765, 767
 Half-wave voltage, 256
 Hammer, ball peen, 141, 437, 468n1
 Hand capacity, 754–755
 Hand, putting it in by, 310
 Hardline, 525
 Hariharan unwrapping algorithm, 686
 Heaviside Oliver, 19, 450
 Heaviside unit step function, 19
 Heinkel, Ernst, 413
 Helicity, 209, 214–215, 264–265, 518
 Helmholtz equation, 9, 11, 313
 Heresy zero bias, 101–102
 Hermitian operator, 450
 Heterodyne, 329, 362, 717
 Hilbert transform, 452–453, 492, 683
 Hindsight, 741, 752
 Histogram, 555–556
 Hole burning, 84, 84n1, 357
 Hologram, 233, 244
 Holographic optical element (HOE), 176–177, 244–245
 Holography, 329
 Homodyne detection, 139, 363, 410
 Hook, 250, 471, 614, 735
 Horowitz, Paul, 448, 468n3, 509, 542–543, 571
 Hot plugging, 414, 609–610
 Human eye, 54–55, 77
 Huyghens' integral, 9, 11, 23, 25, 28, 322, 326–327
 Idiot-resistant, 607
 Illumination
 asymmetric, 11
 critical, 351
 dark field, 350, 373
 diffuse, 205, 353
 flashlight, 350–351
 Köhler, 351
 radiance conservation, 57–58
 radiometry, 54–55
 telecentric, 328
 Illuminator, 52–90
 fiber, 70, 353
 Image, aerial, 275, 325, 329, 334
 Image frequency, 362–363, 410, 484–485, 492
 Image plane, 24, 173, 275, 328, 334, 336
 Image rotator, 275
 Image space, 166, 327, 336
 Image tube, 118–119

- Imaging
 - detectors, 116, 128
 - paraxial, 25, 309
 - perfect, 309, 336
- Impedance
 - and admittance, 512–513
 - car battery, 621
 - complex, 513, 696–697, 719
 - ground loop, 621–622
 - input, 438, 529, 536, 543, 748–749
 - load, 577, 716
 - matched, 471, 561
 - modulating, 82
 - negative, 469
 - output, 529, 536, 542, 577, 716
 - passive, 523
 - photodiode, 703
 - source, 536, 620–621, 717, 733
 - transformer, 520
 - of translinear circuit, 538
 - wave, 3
- Impossibility, 681
- Impulse response, 489, 675
- In situ coherent lidar (ISICL), 47, 47n1
- Index oil, 142, 156, 288, 442–443
- India ink, 200
- Indium tin oxide (ITO), 100, 185, 256
- Inductance, stray, 518, 615
- Inductor, ferrite bead, 624
- Infinity correction, 173–174
- Insertion loss, 640
- Inspiration/perspiration ratio (IPR), 398, 407, 410
- Instability, 38, 266, 595, 767
- Instrument function, 678–679
- Integrated circuit, 509, 534n1, 539
- Integrating sphere, 205–206
 - multipass, 206
- Integration, 743
- Intercept point, 463–465, 567–568
- Interference, AC, 368, 383
- Interface specification, 743
- Interference term, 34, 61–62, 281
- Interferometer, 36–38
 - AC, 34–35, 45
 - alignment, 36–37, 47
 - compensation plate, 36
 - corner cube, 36–37
 - each mode, 38
 - Fabry-Perot, 37
 - fiber, 8, 38
 - fringe counting, 34
 - laser, 36, 38
 - lossless, 8, 28, 32
 - Mach-Zehnder, 36–37
 - Michelson, 36–37
 - multimode, 8
 - operating point, 38–44
 - phase shifting, 23
 - scanning, 31, 47, 49
 - speckle, 49
 - white light, 36
 - with Wollastons, 32
 - Young's slits, 31
- Interline transfer (ILT) CCD, 121, 123, 126, 177
- Intermodulation, 464–465
 - two-tone test, 465, 556
- International Telecommunications Union (ITU), grid, 77
- Invariance, translational, 186, 267
- Invariant, 318
- Iris diaphragm, 85, 411, 418, 426, 439
- Isolation, 543
- Isolator, 291–292
- Isoplanatic, 319, 325
- Jackson, John David, 320
- Jacobian, 311, 314, 323, 677–678
- Jacquinet advantage in FTIR, 368
- Jensen Transformer, 734
- Johnson noise limit, 97–98, 105, 139, 231, 629, 720
- Jones matrices, 227
- Jones vector, 226
- Ju-jitsu, 384, 597
- K-space, 11–13, 209, 243, 262, 320, 322–323, 450, 687
- Kerr cell, 254–257
- Khuri-Yakub, B. T., 333
- Kirchhoff approximation, 322
- Kirchhoff integral, 321, 323
- Kneitel, Tom, 749
- Lagrange invariant, 318–319
- Lambertian source, 30
- Lambertian surface, 31
- Lamp
 - arc, 66
 - discharge, 67
 - flash, 67
 - tungsten, 58–59
- Landau, Lev, D., 56
- Laser
 - alexandrite, 73
 - diode, *See* Diode laser
 - gas, 72–73
 - helium-cadmium (HeCd), 72
 - helium-neon, 72
 - Nd:YAG, 73, 88
 - pulsed, 68, 75
 - pumping, 71–72
 - Q*-switched, 561
 - single-frequency, 72–74
 - solid state, 73–75
 - threshold, 65, 71–73, 82
 - Ti:sapphire, 73
- Laser Doppler velocimetry (LDV), 373

- Laser mode, 85–86, 88, 739
- Laser noise, 83–88
- Laser noise canceller, 379–380
 - log ratio output, 379, 725–726
- Laser resonator, 38, 71, 179, 305
- Law
 - Murphy's, 610, 696
 - Snell's, 4n1, 5–6, 213, 245, 311, 314, 316, 338
- Law, Vernon, 612
- Laziness, 25
- Least squares, 385, 676
- Lehrer, Tom, 688, 688n1
- Lens
 - achromatic, 171
 - aspheric, 170–171
 - camera, 172–173
 - cylinder, 171
 - diffractive, 176–177
 - fast, 169
 - Fourier transform, 151
 - Fresnel, 177
 - graded-index (GRIN), 176
 - microscope, 173
 - plastic, 148–149
 - relay, 176
 - sag, 167–168
 - thick, 167–169
 - thin, 166–167
- Lens bending, 169–170
- Lens design, 166, 337, 345
- Lensmaker's equation, 166–169, 172, 443
- LF356, 571, 693–695, 698
- Lidar, 47, 134, 248, 730
- Lifshitz, E. M., 56, 211n1
- Light
 - background, 134, 250, 360–361
 - black, 380, 441
 - broadband, 229, 235, 306
 - infrared, 54
 - invisible, 174, 765–766
 - parallel, 173, 344, 411
 - polarized, 208–210
 - stray, 199, 237, 759
 - thermal, 57, 133, 208–209
 - ultraviolet, 53–54
 - unpolarized, 208–209
 - visible, 52–53
- Light bucket, 171, 177, 765
- Light-emitting diodes, 63–65
- Light trap, 140, 201
- Limit cycle, 671, 671n4, 680
- Line broadening, 357
 - Homogeneous, 357
 - Inhomogeneous, 357
- Linear applications book, 560
- Linear combination, 141, 227n1, 378–379, 663, 721–723
- Linear system, 25n1, 679, 765
- Linear Technology Corp, 560, 710, 744
- Linearizing, 474, 538, 601–604
- Liquid crystal, 257–258
- Littrow incidence, 238
- Local oscillator (LO), 34, 139, 459, 483
 - in coherent detection, 34–36
- Lock-in amplifier, 93, 377, 502, 548
- Loewen, Erwin, 768
- Loop
 - AGC, 494
 - integrating, 575
 - phase-locked, 501, 587–590
- Lore, 2, 113, 509n1, 560, 645
- Loss, insertion, 45, 487, 562, 640
- Loss tangent, 9
- Low light level CCD, 124
- M-squared, beam propagation factor, 424
- Macromodel, 597, 710, 757
- Magneto-optic effect, 215
- Magnification, anamorphic, 235, 237, 242, 311
- Manufacturing, 340, 392
- Marechal formula, 343
- Market creation, 395
- Matching, laser noise cancellation and, 134
- Matching network, 520, 718–719
- Material
 - birefringent, 217, 219
 - uniaxial, 211, 225
- Matrix
 - ABCD, 25–29
 - coherency, 227–228
 - commutative, 20, 660
 - Jones, 227–228, 232, 265
 - pseudoinverse, 679
 - rotation, 210, 224
 - skew-Hermitian, 269
- Matrix operator, 27–28, 227, 678
- Maxim, 92, 544, 549, 741
- Maximum entropy method, 668, 670, 680
- Maxwell's equations, 3
- McClellan, J. H., 674
- Mean field theory, 3
- Measurement
 - AC, 35, 45, 48, 294, 382–383
 - autocorrelation, 20, 479, 682–683
 - baseband, 362
 - bright field, 35, 87, 359, 371, 413
 - capacitance, 99
 - chopped, 364
 - chopping, 59, 383
 - constrained, 432
 - crunchy, 396
 - dark field, 35, 359
 - DC, 35, 45, 266, 477, 731
 - differential, 80, 378
 - dim field, 361
 - diode laser feedback, 402

- Measurement (*Continued*)
- doubly-dark field, 373
 - dubious, 539, 563
 - extinction, 360
 - film, 404
 - fluorescence, 295, 374
 - flying capacitor, 569–570
 - FM, 477, 730
 - fringe-based, 373, 504
 - ground truth, 397, 409
 - heterodyne, 362, 382
 - intracavity, 402
 - Johnson noise limited, 97–98, 359, 720
 - laser noise canceller, 379, 724
 - linear combination, 141, 378, 721–723
 - low light, 359
 - model independent, 331
 - modulation mixing, 383, 504
 - monochromator, 373
 - multiple scale, 304, 366
 - noise, 386, 475
 - nonlinear, 374
 - nulling, 377–378
 - phase, 365–366
 - photoacoustic, 202, 364, 374, 383
 - photocurrent, 731n2
 - photon counting, 116, 362, 650
 - photon labeling, 381
 - polarization, 228
 - polarization flopping, 230, 376
 - pretty pictures, 397, 398
 - principle, 1, 46, 91, 404, 407, 415
 - pulsed, 49, 497–498
 - ratiometric, 104, 371–372
 - ring-down, 592
 - scanned, 503
 - scanning, 368–369, 398
 - shot noise limited, 363–364
 - SNR of, 61
 - squishy, 397
 - SSB, 487
 - stable, 134, 359
 - subranging, 366
 - thin film, 404
 - time-gated, 503
 - time-resolved, 408
 - two-point correlation, 570, 685
 - ultrasensitive, 730
 - undersampled, 686
 - zero background, 258, 373–376
- Measurement physics, 372, 383, 387, 396
- Medium, isotropic, 3, 9, 15, 211n1, 213, 314
- Mendacity, odor of, 354n1
- Meyer, Robert, G., 509, 533, 534n1
- Microbench, 202, 413, 417–418
- Microbending, 294, 766
- Microchannel plate (MCP), 119–120
- Microscope
- cobbling, 22, 220, 632
 - Nomarski, 364
 - scanning, 173–174
- Miller effect, 701, 712
- Miller, John Lester, 354, 402n2, 689n1
- Milligan's law, 407
- Milne, Alan Alexander, 415
- Milne, Christopher Robin, 415
- Mini-Circuits Labs, 734
- Mirage, 15, 316
- Mirror
- Al, 183
 - dielectric, 184–186
 - enhanced Al, 191
 - focusing, 174
 - metal, 182–184
 - off-axis, 418
 - protected Al, 189
- Mismatch, 32, 523–524
- Missing code, 551, 554, 557
- Mixer, 483, 485, 492, 541
- diode, 532, 544, 559, 566, 568, 585
 - single sideband (SSB), 363, 410, 460, 492–493
- Mode
- beat length, 270
 - coupled, 243, 271, 278, 288
 - degenerate, 300
 - dispersion, 272n1
 - fiber, 8, 282, 285
 - guided, 277, 291
 - high-angle, 271
 - laser, 85–86, 88, 739
 - leaky, 277, 766
 - PM fiber, 273–274
 - stripping, 277
 - transverse electric (TE), 6n1
 - transverse electromagnetic (TEM), 267
 - transverse magnetic (TM), 6n1
- Mode beats, 86–87
- Mode hopping, 84–85, 759
- Model
- Analytical, 9, 409
 - numerical, 9
 - SPICE, 710
- Modularity, 742
- Modulation
- amplitude, 459
 - double sideband (DSB), 460
 - frequency (FM), 462
 - I/Q, 457
 - mixing, 458–459
 - non-quasistatic, 461–462
 - phase, 460–462
 - single-sideband (SSB), 460
 - vector, 492
- Modulation frequency, 456

- Modulation-generated carrier, 301, 305–306, 501, 630
- Modulation index, 239, 305, 384, 458–460, 462, 464, 499
- Modulation mixing, 383
- Modulator
 - acousto-optic, 234n1, 717
 - liquid crystal, 257
 - optical, 254
 - photoelastic, 260–261
- Monochromator, 235–237, 240, 242, 275, 346, 373
 - double, 237
- Moore, E. H., 679
- Morgan le Fay, 316
- Motorboating, 599
- Mounts, 417
- Multiplexer (mux), 757
- Multiplexing, 307–308
 - frequency, 368
- Multiplier
 - frequency, 463–464
 - Gilbert cell, 486, 498, 500, 538, 585, 752
- Murphy's Law, 610, 696
- Name
 - most jaw-cracking, 118n1
- Narcissus, 140
- National Semiconductor, 540, 540n1, 560, 596, 603, 708
- Negative electron affinity (NEA), 110, 374
 - in PMTs, 110
- Negative impedance converter, 469
- Network
 - discrete time, 651
 - LC, 715
 - lead-lag, 462, 574, 588, 708
 - linear phase, 491
 - linear time-invariant, 450–451, 462, , 504, 513
 - matching, 520, 718–719, 733, 735, 758
 - stub, 528
 - T-coil, 719–720
 - peaking, 715–717
 - pulse forming, 256
 - stable, 453, 573
 - time invariant, 450–451, 462, 504, 513
 - two-port, 449, 522
- Network analyzer, 637, 641, 749
- Nitrogen, clean, 444
- Noise
 - $1/f$, 356, 358, 381, 467–468
 - 1-Hz, 474
 - ADC, 506, 557
 - additive, 84, 379, 495, 723
 - amplifier, 97, 363n1, 597, 710, 731
 - background, 137, 724
 - beat, 290
 - BJT, 703
 - in cascade, 536
 - coherence fluctuation, 266, 305, 380
 - combining, 473–475
 - correlations, 475
 - current, 456, 474, 685, 699–700
 - deceptive, 375
 - diode laser, 400
 - double Rayleigh, 281, 380
 - electronic, 136, 358
 - ergodic, 478
 - excess, 113–114, 293, 356, 723
 - fixed pattern (FPN), 124, 126, 128, 356, 646
 - flatband, 601
 - flicker, 467–468, 468n1
 - Gaussian, 40–41, 49, 475–476, 480
 - image frequency, 362, 485, 492
 - intensity, 83–84, 724n1
 - intermodulation, 84, 601, 729
 - Johnson, 689, 692, 708n1
 - lamp, 357
 - laser, 39, 83–88, 356
 - measuring, 628
 - mode partition, 85, 759
 - multiplicative, 84, 477–478, 726
 - of source, 136, 293, 308
 - op amp, 542
 - phase, 464, 477, 487, 558–559
 - PMT, 135
 - polarization, 274
 - popcorn, 468, 470, 515, 550, 762
 - power spectral density (PSD), 45n1, 135
 - in power spectrum estimation, 666–670
 - power supply, 710
 - prerecorded, 747
 - quantization, 45, 358, 553n2, 557
 - reducing, 485
 - shot, 35, 47, 132–134, 497–498
 - amplified, 95
 - full, 133, 470, 538, 703–704
 - statistics, 481–482
 - source, 131–136, 689
 - speckle, 49
 - spectral density, 45n1, 63, 132, 455–456
 - statistically stationary, 478
 - statistics, 136, 356, 478
 - supply, 558, 710
 - technical, 478, 621–625
 - temperature fluctuations, 135
 - thermal (Johnson), 468–470
 - threshold crossings, 49, 480
 - TIA, 708, 711–712
 - timing jitter, 498, 558
 - voltage, 468–470, 472–473
 - white, 467–468
 - wiggle, 88, 759
- Noise budget, 563, 746–747
- Noise current, 472–474, 692, 702, 704, 723
- Noise Eater, 371
- Noise-equivalent photon flux (NE Δ J), 43

- Noise-equivalent power (NEP), 95–97
- Noise-equivalent temperature difference (NETD or NE Δ T), 38
- Noise figure, 470–472, 642, 737
- Noise floor, 557, 561–564, 727, 735
- Noise gain, 457, 644–645
- Noise intermodulation, 84, 601, 724, 726, 729
- Noise measurement, 39, 386, 471–472, 475
- Noise model, 469, 472–473, 697, 702–703
- Noise sidebands, 132, 355, 357, 477–478, 597
- Noise statistics, 136, 356, 478–479
- Noise temperature, 470–472
- Noise voltage, 468–470, 472–473, 716
- Nomarski, 221, 258, 364, 739–740
- Nomenclature, 7, 337
 - confusing, 334
- Nonlinearity
 - ADC and DAC, 554
 - differential (DNL), 554
 - integral (INL), 554
 - intermodulation, 541–542
 - in mixers, 541
 - phase, 466, 641
- Nonreciprocal device, 215, 230
- Nonstationary process, 45, 351, 356, 467, 478
- Norton circuit model, 468–469, 472–473
- Nose grease, 443n1
- Numerical aperture (NA), 10–13, 15
- Numerical Recipes, 645, 658, 668, 677n1, 679
- Nyquist, 336, 504, 506, 655
- Nyquist frequency, 45, 658, 671
- Nyquist limit, 45, 123

- Object space, 25, 166, 327, 336
- Objective lens
 - infinity corrected, 173–174
 - microscope, 173
 - Schwarzschild (mirror), 173, 179
- Obliquity
 - drinking straw test, 30, 323
- Obliquity factor, 323
 - calculating, 323
 - at high NA, 323
 - reciprocal of, 311, 323
- Offset
 - in transimpedance amplifier, 693–714
- Ohm's law, 509–510, 512–513, 569
- Oil
 - contamination, 445
 - index, 156, 164, 256, 442–443
 - snake, 401
- On-off keying (OOK), 503
- Onion problem, 416, 750, 755
- Op amp, 542
 - ideal, 542
- Open/shut ratio, 211, 215–216, 291
- Operating point, 220, 220n1
- Oppenheim, A. V., 645, 676

- Optic axis, 211–214, 217–219, 221–222, 224–225
- Optical activity, 4, 214
 - dispersion, 214
- Optical adhesive, 441–442
- Optical diode, 229, 231
- Optical fiber, 262–308
 - annealed, 153
 - annealing, 307
 - bend birefringence, 277–278
 - bending loss, 277
 - circular birefringence, 4, 214
 - cladding, 277
 - cleaving, 287
 - connectors, 359
 - cost, 287
 - dispersion, 271–272
 - dispersion-shifted, 272
 - étendue, 272–274, 283, 296–297
 - few-mode, 273
 - graded index, 272–273
 - gradient-index, 271
 - high birefringence (HiBi), 274, 299
 - inertness, 147, 212
 - instability, 266
 - launching, 283–285
 - loss, 278–279
 - mode structure, 264
 - modes, 266–268
 - mounting, 286
 - multimode (MMF), 272–273
 - piezo-optic, 278
 - pointing stability, 264
 - polarization instability, 266
 - polarization maintaining (PM), 273–274
 - polarizing (PZ), 306
 - single mode (SM), 273, 300
 - solarization, 279
 - splice, 281, 286–287
 - splicing, 286, 301
 - step index, 262, 266–267, 271–273
 - strain effects, 280
 - strength, 279
 - structure, 307–308
 - theory, 266–272
 - two-mode, 300
 - ultraviolet, 53–54
 - vices of, 265
 - waveguide dispersion, 272, 272n1
- Optical flat, 152, 284, 764n2
- Optical isolator, 215, 230–231
- Optical materials, 145–149
- Optical molasses, 369
- Optical multichannel analyzer (OMA), 367, 670
- Optical path difference, 340
- Optical system
 - scanning, 174
 - staring, 349
 - telecentric, 327

- poor man's, 353
- Optical time-domain reflectometry (OTDR), 295
- Optically thick, 67–68
- Opticlean polymer, 445–446
- Optics
 - coherent, 363n1
 - Fourier, 9, 25n1, 235, 322–323, 325
 - paraxial, 57, 318
 - physical, 311, 316–317, 322
 - ray, 9
 - scalar, 9
- Optimization
 - greedy, 741–742
- Oscillation, 753–755
 - differential, 754
 - fixing, 744
 - motorboating, 599
 - spurious, 739, 750, 753, 763
 - troubleshooting, 744–745
- Oscillator, 597
 - ALC, 597, 599
 - phase noise, 597
 - voltage-controlled (VCO), 537, 587
- Oscilloscope, 440, 748
- Outliers, 480, 503, 762
- Output, DC, 735, 765
- Overdriving inputs, 500, 585, 756
- Overflow, 680
- Overloading outputs, 756
- Oversampling, 555, 655
- Overshoot, 575, 580, 695, 737
- Overtone, 88, 279, 367
 - spectral line, 54–55, 68, 72, 87, 145, 357
- Paint
 - black, 199
 - white, 204, 206, 351, 420
- Paraxial optics, 57
- Paraxial wave equation, 11–12, 456
- Parks, T. W., 674
- Particle, 738
- Passband, 195, 484, 486–488, 491, 494n1, 593–594
- Peaking coil, 715
- Pease's principle, 575, 737
- Pease, Robert, A., 745n1
- Penrose Roger, 679
- Pentaprism, 162
- Perdition, ticket to, 47, 736
- Perfection, 569–571
- Periodogram, 668
- Perspiration, 398
- Perturbation theory, 268
- Perturbation, 15, 268–269, 317, 462, 476
- Phase
 - acoustic, 260
 - insertion, 487, 576
 - Lorentz scalar, 4
 - matching, 4–5
 - topological, 209–210, 265
 - unwrapping, 685–686
- Phase detector, calibration, 590–591
- Phase-frequency detector (PFD), 565, 588–590
- Phase-locked loop (PLL), 501, 587–590
 - acquisition, 590
 - book, 596–597
 - debugging, 606
 - demodulator, 599
 - design, 560–568
 - detector, 577–587
 - frequency compensation, 571–573
 - frequency synthesizer, 589
 - lock detection, 589
 - multiplier, 589, 599
 - noise, 589
 - offset, 585
 - PM demodulator, 587
 - sideband selection, 565
 - synthesizer, 589, 590–591
- Phase margin, 573–574, 588, 695, 737, 754
- Phase matching, 4–9
 - tangential, 5
 - wideband, 750
- Phase modulation, 460–462
 - calibrating, 590–591
- Phase sensitive detector, 498, 585–586
- Phase velocity, 4, 8, 212, 214
- Philips Semiconductor, 560
- Philosopher
 - incompetent, 1
- Photocathode, 110–113, 116, 118–120
- Photodetector
 - infrared, 108
- Photodiode
 - avalanche, 113–115
 - biasing, 710, 753
 - capacitance, 714
 - front end, 690–692
 - germanium, 46
 - IR, 132, 134, 690, 733
 - lateral effect, 104–105
 - PIN, 96, 517
 - position-sensing, 105
 - quadrant, 438
 - quantum efficiency, 95
 - shunt resistance, 721
 - split, 104
- Photoelastic modulator, 260–261
- PhotoFlo, 444
- Photomultiplier, 109–110
 - nonlinearity, 110
- Photomultiplier tube (PMT), 41, 109, 120n1
- Photon budget, 38–44
 - and Strehl ratio, 49
- Photon counting, 111–112, 115–116, 373–374
- Photon labeling, 381

- Photons
 - considered harmful, 35
 - labeling, 382
- Photophoresis, 83
- Photorefractive, 74, 271, 289
- Phototransistor, 103
- Physical theory of diffraction, 316n3
- Physicists, 19, 54–55, 186n1, 689, 750
- Picasso, Pablo, 560
- Pickup, 623, 750–753
- Piezoelectric actuator, 599
- Piezo-optic effect, 258, 296
- Ping pong ball, 206, 351
- Pinhole, 59–60, 292, 345–348, 385, 435, 437
- Pistol, 749
- Plane, of incidence, 5–7, 222
- Plane wave, 3, 10, 322
- Plastic, black, 200, 351
- Plasticizer residue, 445
- Play, 389–390
- Plumber, admirable, 1
- Pockels, 254–257
- Pockels cell, half-wave voltage, 256
- Pockels effect, 256
- Point spread function, 11, 20, 253, 328–329
 - amplitude, 329
- Pointing stability, 264
- Poisson statistics, 41, 482
- Poking, 755
- Polarcor, 216, 760
- Polarization
 - circular, 8, 209
 - dielectric, 118
 - diversity, 153, 306
 - eigenstate, 260, 280
 - elliptical, 8, 232, 760
 - flopping, 230, 376
 - full, 227
 - insensitivity, 232
 - Jones matrix calculus, 227–228
 - labeling, 385
 - linear, 210, 214, 225–226
 - orthogonal, 34, 210
 - partial, 227–228
 - plane of, 4, 7
 - purity, 211, 217, 219–223
 - in scalar optics, 9–11
 - selectivity, 265
 - shifts, 160, 163, 189, 210, 217, 298, 411
 - TE and TM, 61
- Polarization compensator, 229, 227–278
 - Soleil, 229
- Polarization rotator, 229
- Polarizer
 - birefringent, 217–218
 - Brewster, 216–217
 - circular, 226
 - crossed, 361, 373, 375, 762
 - cube, 217
 - double-refraction, 218–221
 - fiber, 232
 - film, 215
 - multilayer, 217
 - open/shut ratio, 211
 - Polarcor, 216
 - TIR, 221–223
 - walkoff, 213–214, 217
 - wire grid, 215–216
- Polarizing film, 229
- Poles and zeros, 451, 573–574, 592
- Poling, 118
- Polynomial, 341, 463, 592
 - Zernike, 341
- Polystyrene latex (PSL), 40, 50
- Polystyrene latex spheres, 420, 739
- Pooh, Winnie the, 415
- Porsche, Ferdinand, 413
- Portrait, 16–17
- Post-processing, 45, 400, 644–687
 - obese, 400
- Potentiometer, 511
- Power supply, quiet, 600
- Poynting vector, 3, 211, 213, 313
- Pressure-broadening, 66
- Pretty pictures, 397–398
- Principal axes, 211–212
- Principal planes, 167–168
- Prism
 - anamorphic, 171, 175, 346
 - constant deviation, 162, 165, 437
 - Dove, 161
 - Glan-Taylor, 223
 - Glan-Thompson, 223
 - glass, 160–165
 - Littrow, 161–162
 - Nomarski, 364, 413, 740
 - pathologies, 165
 - polarizing, 32, 221
 - right-angle, 160–161
 - Rochon, 219–220
 - roof, 163
 - Savart plate, 218
 - Senarmont, 219
 - TIR, 160, 180
 - unfolding, 143
 - wedge, 162–163
 - Wollaston, 218–219
- Probe, 373–374, 627, , 636, 749
- Problems, chapter, 125, 413
- Project, choosing, 389
- Propagation constant, 3, 56n2, 270, 642n1
- Propagation, calculating, 320–321, 557
- Propagator
 - Huyghens, 12, 17, 20, 322
 - Kirchhoff, 12, 322
 - Rayleigh-Sommerfeld, 322

- Prototype, 629–634
 banging on it, 419
 brassboard, 742
 breadboard, 416, 742
 not breaking, 632
 building, 629, 631, 742
 dead-bug, 630–631, 752
 electronic, 612
 flakiness, 634–635
 height adjustment, 635
 initial, 634
 layout, 623
 limits of, 634
 SMT, 629, 635
 testing, 632
 Pulling, frequency, 292
 Pupil, 318
 entrance, 24–25, 318, 325, 351
 exit, 18, 24, 32, 221, 318, 322
 function, 25, 158, 276, 324–325, 328
 moving, 253
 Pupil function, 328
 annular, 276
 Purse, Silk, 681
 Pyramids, 392

 Quadrature pair, 270, 579
 Quadrature, 457
 Quantization, 358, 504–505
 Quantum efficiency, 95
 Quenching, 89–90, 298

 Radiance
 conservation of, 57–58
 spectral, 54–55, 57
 Radiometry, 54–55
 Railed, 576, 588, 642, 745
 Raman effect, 297
 Random process, 111, 478–479, 666
 Raster, 249, 252, 649
 Raster scan, 253
 Ray
 extraordinary, 212, 214
 ordinary, 211, 316
 Ray intercept error, 321, 341
 Ray optics, 9, 15–17
 ABCD matrix, 25–29
 Rayleigh range, 13–15, 281–282, 318, 762
 Rayleigh scatter, 150, 279, 281, 355, 586
 Rayleigh scattering, 150, 227, 266, 278, 281, 290
 Rayleigh-Sommerfeld integral, 318, 321, 323–324
 Rayleigh-Sommerfeld propagator, 322
 Rayleigh's $\lambda/4$ criterion, 83n1, 343
 Rayleigh's theorem, 21, 656, 656n2
 Reactance, 512–514, 518–520, 602, 640–641, 715, 737
 Reasoning, 7, 23, 746, 748

 Receiver, superheterodyne, 139
 Reciprocal dispersive power V , 146
 Reciprocity, 282, 283
 Reflectance, 6, 37, 123, 154, 189–190, 199–200, 204–206
 Reflection
 at birefringent surface, 222
 helicity change, 209
 multiple, 199, 216, 519
 at surface of birefringent material, 222
 total internal (TIR), 7
 Reflection coefficient, 23, 180–181, 187
 Reflectivity, 7, 102, 156, 183–184, 204
 Refraction
 critical angle, 7, 164
 double, 32, 218–221
 Refractive index n , 146, 212
 temperature coefficient, 148, 280
 Reinholdtsen, Paul, 333
 Relaxation oscillation, in lasers, 72
 Remote sensing, 54, 124, 409
 Resampling, 676–678
 Resistance
 emitter, 534, 537, 608, 702, 754
 ESR, 515, 625
 extrinsic base, 472, 702
 évenin, 468–469, 472
 video, 582
 Resistor, 510–512
 metal, 510
 noisy, 510
 Ohm's law, 510, 512–513
 thermoelectric coefficient, 510
 Resolution
 diffraction-limited, 241
 model-independent, 331
 Rayleigh, 260
 Sparrow, 260
 spectral, 164, 237, 239, 242, 346
 spurious, 334
 Resolvable spots
 AO cell, 260–261, 730
 and Fresnel number, 14–15
 Resolving power, of grating, 241
 Resonance
 LC, 67, 509
 linearization, 501, 590, 602
 self, 625
 Resynchronizing, 558–559, 606
 Retardation, 212–213
 static, 256
 stress, 762
 Retarder
 Fresnel rhomb, 164–165
 full-wave, 225
 Jones matrix, 227–228
 multi-order, 225
 quarter-wave, 224

- Retarder (*Continued*)
 - variable, 229
 - wave plate, 224
 - zero-order, 225–226
- Retroreflector, specific brightness, 246
- Return loss, 525, 528, 640–641
- Richardson Grating Laboratory, 243
- Risk, technical, 2, 387–388, 400, 415, 738, 742
- Roundoff error, 650, 671, 674, 680
- Rubicon, 743
- Rule of One, 35–36
- Rule of thumb, 168–169, 694
- Runyon, Damon, 644
- Rutherford, Ernest, Lord, 467

- Sag, 167–168
- Sagitta, 168, 239, 764
- Sagnac interferometer, 265, 301–303
- Sagnac, 301
- Sales pitch, 392, 394, 397, 415
- Sampled data, 496, 651, 653, 656, 670, 750
- Sampled data system, 336, 462, 652–654
- Sampling theorem, 506, 605, 654–655
- Sampling, track/hold, 550, 671
- Sapphire, *c*-axis normal, 153, 212, 740
- Savart plate, 218, 413, 739–740
- Scalar optics, 9–11
- Scallop loss, 579, 664, 666
- Scan, constant speed, 248, 250
- Scanner
 - acousto-optic, 258–259
 - constant speed, 250–251
 - galvo, 247
 - hologon, 251–253
 - jitter, 247, 250, 252–253, 260
 - linear, 250–251
 - mechanical, 253
 - polygon, 248–249
 - wobble, 247
- Scanning
 - pre-and post-objective, 250
 - telecentric, 259n1
- Scatter/Scattering
 - dipole, 40
 - forward, 360, 504
 - Rayleigh, 150, 227, 266, 279, 290, 586
 - total integrated (TIS), 151, 360
- Schlieren effect, 316, 412, 765
- Scratch/dig, 150
- Self-resonance, 625
- Sensors, distributed, 263–264, 307
- Serendipitous information, 426
- Servo, constant false alarm rate (CFAR), 50
- Settling time, 545–546
- Sha (sampling) function, 497, 654
- Shack-Hartmann sensor, 177–178
- Shafer, R. W., 645, 676
- Shape factor, 488, 762
 - of a filter, 488
- Shear plate, 427, 764
- Shield, 525–526, 529–530, 621–623, 626
- Shielding, magnetic, 626
- Shims, 416, 418
 - metal, 416
 - sticky note, 418
- Shoe horn, 749
- Shot noise, 35, 132–134, 470, 481
- Shot noise limit, 47, 363, 730n1
- Sideband
 - angle modulation, 458, 462
 - double (DSB), 460
 - lower (LSB), 460
 - single (SSB), 460
 - upper (USB), 456, 460
- Sideband folding, 457
- Sidelobe, 194–195, 241, 663–665
- Siegmán, A. E., 29n1, 71n1
- Sign convention
 - Fourier transform, 449
- Signal
 - analytic, 451–453
 - baseband, 455–456, 462, 498, 759, 760n1
 - CW, 453
 - modulation, 460, 462, 478
 - spurious, 86n1, 355–356
- Signal averaging, 496, 503–504, 683–685
 - analogue, 44–45
 - design, 44–45
 - digital, 4n2
 - front end, 44–46
 - messy, 211n2, 309
 - prototype, 750
 - RF, 463, 543–544, 560, 640
 - statistical, 467, 747
 - strategy, 44–51
- Signal to noise ratio (SNR), 92–94
 - optical vs electrical, 94
- Signal-to-noise ratio (SNR), 92–94
 - budget, 97, 143
 - calculating, 131
 - capacitance and, 99
 - deconvolution, 675–676
 - degradation, 141
 - digital, 143
 - feedback and, 101, 103
 - front end, 97
 - high, 102
 - improving, 113, 143
 - limitation, 97, 100
 - low, 127, 476, 499, 681–682
 - minimum, 97
 - plotting, 698
 - pulses and, 102–103
 - Strehl ratio and, 343
 - tradeoff, 121, 124–125
 - whitening, 198, 462

- Signature analysis, 747
- Simplicity, 15, 176, 291, 308, 340, 560
- Simulation, SPICE, 624, 710
- SINAD, 556–557
- Single sideband (SSB), 363, 410, 455, 460, 492–493
- Singular value decomposition (SVD), 385, 679
- Singularity, 313, 451
- Sirius cybernetics corporation, 738
- Skin depth, 488n1, 525, 619–620
- Slit, 235, 241, 368
- Slowly-varying envelope equation, 11, 327
- Snake oil, 401
- Snell's law, 4n1, 5–6, 10n1, 213, 245, 311
- Software radio, 558
- Soldering, 629
- Solid angle, 30–31
- Source
 - continuum, 55–59
 - diffuse, 273
 - isotropic, 30
 - Lambertian, 30
 - laser, 434, 436, 501
 - LED, 130, 507
 - superluminescent, 65
 - thermal, 59–60, 66, 69, 435
- Spark gap, 256, 607
- Spark plug, 747
- Spatial filter, 17–18, 345
- Spatial frequency, 11, 24, 327–328, 330–331, 334–336
- Spatial light modulator, 258
- Specification, operating, 38–44
- Speckle, 61, 207, 429
- Speckle interferometry, 355, 404, 430
- Specsmanship, 402n1, 540
- Spectral density, 56, 133, 455–456
- Spectral encoding, 296–297
- Spectral leakage, 659
- Spectral radiance, 54–55, 57
- Spectrometer
 - crude, 678
 - Czerny-Turner, 373
 - diode laser, 84, 379–380, 503, 730
 - FFT, 551, 664, 750
 - fiber, 669–670
 - Fourier transform infrared (FTIR), 368
 - grating, 201n2, 237, 297, 484n2
 - OMA, 367, 670, 677
 - resolution, 39, 83, 275
 - tunable laser, 646
- Spectroscopy
 - Doppler free, 2, 413, 430, 504
 - fiber plate, 297
 - FM, 82, 383–384
 - Fourier transform infrared (FTIR), 368
 - photochemical grating, 373
 - Raman, 197, 237, 369, 373–374
 - slitless, 368
 - sub-Doppler, 369
 - thermal lens, 316
 - tunable laser, 646
- Spectrum, 52–54
 - angular, 23–24, 238, 240–241, 346, 350–351
 - baseband, 72, 86
 - black body, 56–57
 - continuum, 55–59
 - emission, 65–66
 - from grating, 72–73, 77–78, 89
 - line, 67–69
 - modulation, 457
 - noise, 99, 132, 455–456, 599
 - normalization, 663
 - plane wave, 9, 12, 61, 322–323
 - power, 20, 368, 479–480, 666–670
 - solar, 145
 - visible, 177, 678
- Spectrum analyzer, 748–749
- Speed of sound, in AO cells, 260, 730
- Splatter, 463, 756
- Spur, 86, 475
- Spurious-free dynamic range, 556
- Spurious signal, 86n1, 88, 753
- Stability, measurement, 688
- Stabilization
 - intensity, 371, 405
 - pointing, 264, 370–371, 405
- Stagger tuning, 192, 195
- Startup, 600–609, 756
- Startup circuit, 756
- Statistics, gaussian, 479–481
- Steep, 169
- Stefan-Boltzmann formula, 57
- Stegun, Irene, 331
- Stellar magnitude, 42–43, 201n2
- Stick-slip, energy storage, 286, 417, 434, 759
- Sticker shock, 307
- Stokes, 259, 297–298
- Stokes parameters, 227, 227n1
- Stopband, 195, 198, 487–488, 491, 594, 639
- Stops, 18–19, 169, 405
- Stray, 292, 639
- Stray capacitance, 518, 616
- Stray inductance, 518, 615
- Stray light, 49, 198–199, 201, 444, 759, 762, 766
- Strehl ratio, 343
 - and SNR, 49
- Student
 - graduate, 38, 274, 397, 467, 502
 - optics, 322
 - physics, 769
- Subsystem, 560–611, 743
- Successive approximation, 500, 552, 595–596
- Summing junction, 473, 542, 704, 754–755
- Sunlight, 32, 419, 430
- Superheterodyne receiver, 139
- Support, 395, 651

- Surface, Lambertian, 31
- Surface normal, 5, 210, 218, 251
- Surface Plasmon, 291
- Susceptance, 448, 513
- Suspiciousness, 65, 424n1, 629n1, 687, 710, 745
- Symmetry, 20, 264, 339
- System integration, 388, 738–739, 743
- System specification, 742, 762n1
- Systems engineers, 561–564, 741–742

- T-coil, 719–720
 - constant-resistance, 720
- T-shirts, 3
- Tape
 - double sticky, 416, 441
 - Magic Invisible, 65, 205
- Technical taste, 398–402
- Telecentric, 259n1, 327–328
- Telescope, 43
- Temperature coefficient, 280–281
 - of expansion, 146
 - of index, 146
 - of optical path length, 147
- Temperature compensation, 306–307
- Temperature control, 115, 156, 571, 629, 756
- Temperature, controller, 571
- Temperature gradient, 316, 510, 629
- Termination, 472, 527–529
- Test
 - bar pattern, 425, 765
 - beam quality, 764
 - component, 763–764
 - duck, 134
 - Foucault knife-edge, 175
 - imaging, 764
 - knife-edge, 175, 764
 - misleading, 763
 - MTF, 335
 - optical stability, 381
 - optical, 422
 - signature, 409, 747
 - subsystem, 743
 - two-tone, 465, 556
- Test equipment, 93, 612, 753, 758
- Test point, 39, 627, 636, 749
- Test points, 627
- Theorem
 - autocorrelation, 334, 479n1
 - Bode's, 718–719
 - central-limit, 479
 - convolution, 20, 479, 676
 - derivative, 656
 - equipartition, 468n2, 470
 - fluctuation-dissipation, 468, 471, 569, 733
 - Green's, 321
 - Hermitian matrix, 269n1
 - memorizing, 19
 - Nyquist's, 504
 - optical, 2, 360–361, 416
 - Parseval's, 21n1, 664
 - power, 21, 330, 656
 - Poynting's, 212n1
 - Rayleigh's, 21, 656, 656n2
 - sampling, 505n1, 506, 605, 654–655
 - Shannon's, 681–682
 - shift, 656, 666
 - Taylor's, 382, 654
 - Widrow's, 505n1
 - Wiener-Khinchin, 479, 479n1
- Theory
 - diffraction, 11, 309, 312, 320
 - linear systems, 25n1, 679
- Theory wins, 747
- Thermal equilibrium, 56, 209
- Thermal lensing, 74, 201, 316, 759
- Thermodynamic, 56–57, 66–67, 561–564
 - second law, 56–58
- Thermoelectric coefficient, 510, 758
- Thermometer, 266, 294, 297–298, 307
- Thin lens approximation, 166
- Thin object, 23–24, 151, 324–325, 330
- Thresholding, 50, 111, 116, 480, 482
- Time-bandwidth product, 260, 681, 689n2
- Time delay integration (TDI), 124, 367
- Time-invariant, 9, 449–451, 455, 513, 651
- Timing, resynchronizing, 558–559, 606
- Time to market, 741, 743–744
- Tire iron, 437
- Titchmarsh, E. C., 490, 490n1
- Tone burst, 47–50, 480, 561–562, 591, 683
- Total internal reflection (TIR), 7, 160
 - frustrated, 7
- Track/hold amplifier, 550–551
- Tradeoff, lifetime vs brightness, 58
- Transconductance, 533, 538–539, 702, 704, 708
- Transducer, 261, 364, 622, 757
- Transfer function
 - coherent (CTF), 328–33
 - modulation (MTF), 335
 - optical (OTF), 333–334
 - stable, 490
- Transformer
 - series connection, 721
 - transmission line, 529–530
- Transform, Hilbert, 452–453, 489, 492, 683
- Transimpedance amplifier, 693–714
 - cascade, 701–702
 - noise in, 702
- Transistor
 - bipolar, 103, 533
 - BFG25A/X, 578, 704, 706, 709, 712–714
 - cutoff frequency f_T , 534
 - inverted, 535
 - model, 701
 - noise model, 702
- Transition band, 487, 491, 496, 674, 762

- Translinear circuit, 531, 538
- Transmission coefficient, 6, 17, 24, 187, 323, 346
- Transmission line
 - electrical length, 523, 525, 614
 - microstrip, 527, 613–614
- Transmittance
 - external, 197
 - internal, 197
- Transverse electric (TE) mode, 6n1
- Transverse electromagnetic (TEM) mode, 267
- Transverse magnetic (TM) mode, 6n1
- Trend removal, 656–657
- Trick, sleazy, 501, 579
- Triton surfactant, 444–445
- Troubleshooting
 - digital, 755
 - oscillations, 749, 753–755
- Truncation error, 652n1, 677
- Tube length, 173
- Tuning, laser, 77, 85, 561

- Ultraviolet, 53
 - and fiber, 64
- Uniaxial material, 211, 225

- Variational principle, 315–316
- Vector, Poynting, 3, 211, 213
- Velocity
 - group, 4, 8, 214, 271, 456, 491, 494
 - phase, 4, 212, 214, 271
- Verdet constant, 215, 298
- Vernier, 434–435
- Vertical illuminator, 158
- Vibration isolation, 408
- Video cameras, 121, 126–127
- Vignetting, 19
 - and fringes, 766
 - and noise, 730
 - nonuniform, 24
 - and translational invariance, 186, 267
- VSWR, 524–525, 643

- Walkoff, 213–214
- Walkoff plate, in CCD camera, 214
- War story, 144n1, 738
- Washington, George, 333
- Wasteland, vast, 126
- Water, ultrapure, 372, 444

- Wave
 - forward, 186, 523, 530
 - plane, 3, 320
 - Plasmon, 291
 - propagating, 7–8, 234, 526
 - reverse, 186, 523–524, 617
 - vector, 3, 49, 57, 234
- Wave equation
 - Helmholtz, 320
 - paraxial, 11–12, 456
 - scalar, 313, 320
 - vector, 3, 9
- Wave impedance, 3
- Wave number, 54, 79, 159, 298, 379
- Wave plate, 224–226
 - full, 225
 - multi-order, 225
 - quarter, 224
 - zero-order, 225–226
- Wave propagation, calculating, 9–32
- Wave vector, 3, 49, 57, 234
- Wavefront, 178, 260311–312
- Waveguide, 285
- Wavelength, 127, 170
- Weak distortion limit, 463
- Web site, 613
- Wedge angle, 29, 151–152, 222
- Whitening, 198, 462
- Williams, Jim, 403n1, 618n1, 630n2, 720n1, 744, 744n1
- Window, 151–152
 - Brewster, 73, 231, 418, 437, 444
 - function, 663–664, 666, 673, 684
 - rectangular, 683
 - von Hann, 667
- Windup, 575–576
- Wolf, Emil, 7n1, 15n1, 31, 60, 211n2, 227, 313
- Wood's Horn, 201–202
- Working distance, 46, 137, 167, 173, 250, 285

- Yariv, A., 266n1, 281n1
- Young's slits, 59

- Zero padding, in DFT, 664–66
- Zero path difference, 218, 364, 384, 435
- Zero-order hold, 550n1, 671–672
- Zone plate, 176, 244, 312–313

Notation

Electric field	E
Magnetic induction	B
Wavelength	λ
Frequency (in hertz)	f electrical ν optical
Radian frequency	$\omega = 2\pi f$
Wave vector	k
Cartesian unit vectors	$\hat{x}, \hat{y}, \hat{z}$
Poynting vector	S
Eikonal	S
Logarithms	$\ln x$ base e , $\log x$ base 10
Forward Fourier transform (electronics)	$H(f) = \int_{-\infty}^{\infty} h(t)e^{-i2\pi ft} dt$
Forward Fourier transform (optics)	$E(\mathbf{k}, \nu) = \iiint_{-\infty}^{\infty} E(\mathbf{x}, t)e^{i(-\mathbf{k}\cdot\mathbf{x}+2\pi\nu t)} d^3x dt$
Heaviside unit step function $\mathbf{U}(x)$	$= 0 \quad x < 0$ $= 0.5 \quad x = 0$ $= 1.0 \quad x > 0$
Rectangle function $\text{rect}(x)$	$= 1 \quad -0.5 < x < 0.5$ $= 0 \quad \text{otherwise}$
Sinc function	$\text{sinc } x = \sin(\pi x)/(\pi x)$
Sinc Fourier transforms	$\text{rect}(x) \supset \text{sinc}(f)$ $\text{sinc}(x) \supset \text{rect}(f)$
Circ function $\text{circ}(u, v)$	$= 1 \quad u^2 + v^2 < 1,$ $= 0 \quad \text{otherwise}$
Jinc function (Airy pattern)	$\text{jinc } x = J_1(2\pi x)/(\pi x)$
Jinc Fourier transforms	$\text{circ}[(u^2 + v^2)/(\text{NA})^2] \supset \text{jinc}(r\text{NA}/\lambda) \quad (2D)$
Sha (sampling) function	$\text{III}(x) = \sum_{n=-\infty}^{\infty} \delta(x - n)$
Sha Fourier transform	$\text{III}(x) \supset \text{III}(f)$
1 megaton	$10^{15} \text{ cal} = 4.184 \times 10^{15} \text{ J}$
1 parsec	$3.086 \times 10^{16} \text{ m}$ or 3.26 lightyears
1 cm	1.1126 pF
500 nm	$600 \text{ THz} = 20,000 \text{ cm}^{-1}$
1 cm^{-1}	30 GHz
$1 \text{ BTU}\cdot\text{in.}/\text{ft}^2/\text{h}^\circ\text{F}$	0.144 W/m/K
$1 \text{ BTU}/\text{h}/\text{ft}^\circ\text{F}$	1.73 W/m/K
Speed of light c	299792458 m/s (exactly)
Planck's constant h	$6.6262 \times 10^{-34} \text{ Js}$
Boltzmann's constant k	$1.3806 \times 10^{-23} \text{ J/K}$
Electron charge e	$1.6022 \times 10^{-19} \text{ coulomb}$
1 eV photon hc/e	$1.240 \mu\text{m}, 242 \text{ THz}$
#photons/joule = λ/hc	$5.03 \times 10^{15} \lambda \text{ (nm)}$
Thermal voltage	$kT/e = 25.7 \text{ mV}$ at 300 K

WILEY SERIES IN PURE AND APPLIED OPTICS

Founded by Stanley S. Ballard, University of Florida

EDITOR: Bahaa E. A. Saleh, Boston University

BARRETT AND MYERS • *Foundations of Image Science*
BEISER • *Holographic Scanning*
BERGER-SCHUNN • *Practical Color Measurement*
BOYD • *Radiometry and The Detection of Optical Radiation*
BUCK • *Fundamentals of Optical Fibers*, Second Edition
CATHEY • *Optical Information Processing and Holography*
CHUANG • *Physics of Optoelectronic Devices*
DELONE AND KRAINOV • *Fundamentals of Nonlinear Optics of Atomic Gases*
DERENIAK AND BOREMAN • *Infrared Detectors and Systems*
DERENIAK AND CROWE • *Optical Radiation Detectors*
DE VANY • *Master Optical Techniques*
ERSOY • *Diffraction, Fourier Optics and Imaging*
GASKILL • *Linear Systems, Fourier Transform, and Optics*
GOODMAN • *Statistical Optics*
HOBBS • *Building Electro-Optical Systems: Making It All Work*, Second Edition
HUDSON • *Infrared System Engineering*
IIZUKA • *Elements of Photonics, Volume I: In Free Space and Special Media*
IIZUKA • *Elements of Photonics, Volume II: For Fiber and Integrated Optics*
JUDD AND WYSZECKI • *Color in Business, Science, and Industry*, Third Edition
KAFRI AND GLATT • *The Physics of Moire Metrology*
KAROW • *Fabrication Methods for Precision Optics*
KLEIN AND FURTAK • *Optics*, Second Edition
MALACARA • *Optical Shop Testing*, Third Edition
MILONNI AND EBERLY • *Lasers*
NASSAU • *The Physics and Chemistry of Color: The Fifteen Causes of Color*, Second Edition
NIETO-VESPERINAS • *Scattering and Diffraction in Physical Optics*
OSCHE • *Optical Detection Theory for Laser Applications*
O'SHEA • *Elements of Modern Optical Design*
OZAKTAS • *The Fractional Fourier Transform*
SALEH AND TEICH • *Fundamentals of Photonics*, Second Edition
SCHUBERT AND WILHELMI • *Nonlinear Optics and Quantum Electronics*
SHEN • *The Principles of Nonlinear Optics*
UDD • *Fiber Optic Sensors: An Introduction for Engineers and Scientists*
UDD • *Fiber Optic Smart Structures*
VANDERLUGT • *Optical Signal Processing*
VEST • *Holographic Interferometry*
VINCENT • *Fundamentals of Infrared Detector Operation and Testing*
WILLIAMS AND BECKLUND • *Introduction to the Optical Transfer Function*
WYSZECKI AND STILES • *Color Science: Concepts and Methods, Quantitative Data and Formulae*,
Second Edition
XU AND STROUD • *Acousto-Optic Devices*
YAMAMOTO • *Coherence, Amplification, and Quantum Effects in Semiconductor Lasers*
YARIV AND YEH • *Optical Waves in Crystals*
YEH • *Optical Waves in Layered Media*
YEH • *Introduction to Photorefractive Nonlinear Optics*
YEH AND GU • *Optics of Liquid Crystal Displays*

Physical Constants and Rules of Thumb

Solar constant (at zenith, above atmosphere)	1.36 kW/m ² , 20 MW/m ² /sr, 136 klx
Transmittance of clear atmosphere	0.8
Minimum luminance for easy reading	5 lx
Bright room lights	400 lx
Bright desk lamp for close work	7000 lx
Peak luminous efficiency of light-adapted eye	683 lm/W @552 nm, -3 dB @510&610
Brightest stars ($m_v = 0$)	2.0 μ lx at ground
Faintest naked-eye star ($m_v = 6$)	$8 \cdot 10^{-9}$ lx at ground
Black body radiation	56.7 kW/m ² at 1000 K
Earth's magnetic field B	0.3 to 0.6 gauss
Circular cone of half-angle θ	$NA = n \sin \theta$, $\Omega' = \pi(NA)^2$
Airy disc radius of circular aperture of radius a	$0.61\lambda/a$ radians ($a \gg \lambda$)
Étendue of Gaussian beam	$(\pi\lambda/4)^2$
Waist radius of Gaussian beam	$w = \lambda/(\pi NA)$
Airy disc diameter in the visible	$a \approx f\# = 0.5/NA$ (a in microns)
Defocus Tolerance	$ \Delta Z < 0.5\lambda/(NA)^2$
Peak efficiency of an optical system + photodiode	0.4 to 0.8 depending on coatings
Image flux density at $f/8$ (0.063 NA), distant object	1% of object flux density (Lambertian)
Strehl ratio with rms wavefront error E waves	$S \approx \exp(-E^2/2)$
Diffraction limit	$\lambda/4$ rms wavefront error \rightarrow 0.8 Strehl
Hyperfocal distance of lens of diameter D	D^2/λ
Aberration scaling with NA and field angle	Spherical $\propto (NA)^3$, Coma $\propto (NA)^2\theta$, Astigmatism & Field Curvature $\propto (NA)\theta^2$, Distortion (barrel or pincushion) $\propto \theta^3$
Things invariant under magnification:	radiance, $n^2 A \Omega'$, # resolvable spots, phase shift, total power
Fibre Étendue: $n^2 A \Omega'$ (cm ² ·sr)	$\sim 3 \cdot 10^{-6}$ (SM), $3 \times 10^5 - 10^{-3}$ (step MM)
Responsivity of photodiode ($\eta = 1$)	$\mathcal{R}_{\max} = \lambda/1.240 \mu\text{m A/W}$
1 dB increase	26% power, 12% voltage
Additional noise producing 1 dB SNR reduction	5.87 dB below noise floor
Shot Noise limit	$i_{\text{photo}} R_L > 2kT/e$ (50 mV at 300 K)
Shot Noise Rule of One	1σ AC shift with 1 photon/s in 1 Hz coherently added
Shot Noise of 1 mA	17.90 pA/ $\sqrt{\text{Hz}}$
Resistor with Johnson noise of 1 nV/ $\sqrt{\text{Hz}}$ (300 K)	60.4 Ω
Resistor with Johnson noise of 1 pA/ $\sqrt{\text{Hz}}$ (300 K)	16.56 k Ω
Noise power with matched source (NF = 0 dB for pure Johnson noise)	$P_J(\text{dBm}) = -173.8 + \text{NF}$ $+ 10 \log \left(\frac{T}{300\text{K}} \right) + 10 \log \text{BW}$
Quantization noise	$1/\sqrt{12}$ ADU
Sine wave power (50 Ω)	$P(\text{dBm}) = 4 + 20 \log_{10} V_{\text{pp}}$
Transconductance of bipolar transistor	$g_m = e i_C / (kT) = i_C / 25.7 \text{ mV}$ (300 K)
Sheet resistance of 0.5-oz copper (0.017 mm)	1.0 m Ω/\square at 25 $^\circ$ C
Inductance of 1 inch component lead	$L \approx 20 \text{ nH}$
Capacitance of 1mm ² pad on 4-layer card	0.08 - 0.2 pF (1 & 3 layers from ground)
Low frequency capacitance of RG-58 cable	100 pF/m