

UNDERSTANDING LASERS

An Entry-Level Guide

Fourth Edition

JEFF HECHT


IEEE PRESS

WILEY

Copyright © 2019 by The Institute of Electrical and Electronics Engineers, 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 is available.

hardback: 9781119310648

Cover Design: Wiley

Cover Image: tiero/ Getty Images

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

To my many friends in the optics community, to the many people who have given graciously of their time helping me understand more about optics, people, and the world around us, and to the coming generation in hope this book helps you get started in the fascinating world of lasers and optics.

CONTENTS

<i>Preface</i>	xiii
CHAPTER 1 Introduction and Overview	1
1.1 Lasers, Optics, and Photonics	1
1.2 Understanding the Laser	3
1.3 What Is a Laser?	4
1.4 Laser Materials and Types	8
1.5 Optical Properties of Laser Light	10
1.6 How Lasers Are Used?	14
1.7 What Have We Learned?	17
CHAPTER 2 Physical Basics	21
2.1 Electromagnetic Waves and Photons	21
2.2 Quantum and Classical Physics	29
2.3 Interactions of Light and Matter	39
2.4 Basic Optics and Simple Lenses	47
2.5 Fiber Optics	51
2.6 What Have We Learned?	54
CHAPTER 3 How Lasers Work	59
3.1 Building a Laser	59
3.2 Producing a Population Inversion	60
3.3 Resonant Cavities	66
3.4 Laser Beams and Resonance	73
3.5 Wavelength Selection and Tuning	81
3.6 Laser Excitation and Efficiency	85
3.7 What Have We Learned?	89

CHAPTER 4	Laser Characteristics	95
4.1	Coherence	95
4.2	Laser Wavelengths	98
4.3	Properties of Laser Beams	103
4.4	Laser Power	108
4.5	Laser Efficiency	110
4.6	Pulse Characteristics	115
4.7	Polarization	120
4.8	What Have We Learned?	121
CHAPTER 5	Optics, Laser Accessories, and Measurements	127
5.1	Classical Optical Devices	127
5.2	Optical Materials	136
5.3	Optical Coatings and Filters	141
5.4	Beam Delivery, Direction, and Propagation	145
5.5	Mounting and Positioning Equipment	148
5.6	Nonlinear Optics	149
5.7	Beam Modulation and Output Control	156
5.8	Measurements in Optics	159
5.9	What Have We Learned?	164
CHAPTER 6	Laser Types, Features, and Enhancements	169
6.1	Perspectives on Laser Types	169
6.2	Laser Media	171
6.3	Pumping and Energy Storage	177
6.4	Laser Pulse Characteristics	182
6.5	Wavelength Conversion	195
6.6	Laser Oscillators and Optical Amplifiers	201
6.7	Wavelength Options	207
6.8	Laser-Like Light Sources	209
6.9	What Have We Learned?	211
CHAPTER 7	Gas Lasers	217
7.1	The Gas-Laser Family	217
7.2	Gas-Laser Basics	218
7.3	Helium–Neon Lasers	225
7.4	Argon- and Krypton-Ion Lasers	229
7.5	Metal-Vapor Lasers	232
7.6	Carbon Dioxide Lasers	235
7.7	Excimer Lasers	240

7.8	Nitrogen Lasers	243
7.9	Chemical Lasers	243
7.10	Other Gas Lasers	246
7.11	What Have We Learned?	247
CHAPTER 8	Solid-State Lasers	253
8.1	What Is a Solid-State Laser?	253
8.2	Solid-State Laser Materials	258
8.3	Solid-State Laser Configurations	265
8.4	Major Solid-State Laser Materials	271
8.5	Optically Pumped Semiconductor Lasers	284
8.6	Broadband and Tunable Solid-State Lasers	288
8.7	Pulsed Solid-State Lasers	294
8.8	What Have We Learned?	295
CHAPTER 9	Fiber Lasers and Amplifiers	301
9.1	What Are Fiber Lasers?	301
9.2	Optical Fiber Structures	306
9.3	Fiber Laser Design and Efficiency	310
9.4	Rare-Earth-Doped Fiber Lasers	318
9.5	Rare-Earth-Doped Fiber Amplifiers	328
9.6	Raman Fiber Lasers and Amplifiers	332
9.7	What Have We Learned?	335
CHAPTER 10	Diode and Other Semiconductor Lasers	341
10.1	Types of Semiconductor Lasers	341
10.2	Development of Diode Lasers	342
10.3	Semiconductor Basics	344
10.4	Comparing LED and Diode-Laser Emission	353
10.5	Confining Light and Current	359
10.6	Edge-Emitting Diode Lasers	370
10.7	Surface-Emitting Diode Lasers	375
10.8	Optical Properties of Diode Lasers	379
10.9	Diode-Laser Materials and Wavelengths	381
10.10	Quantum Cascade Lasers and Related Types	390
10.11	What Have We Learned?	393
CHAPTER 11	Other Lasers and Laser-Like Sources	399
11.1	Tunable Dye Lasers	399
11.2	Optical Parametric Sources	404

11.3	Supercontinuum Sources	408
11.4	Frequency Combs	408
11.5	Extreme Ultraviolet Sources	410
11.6	Free-Electron Lasers	416
11.7	What Have We Learned?	420
CHAPTER 12	Low-Power Laser Applications	425
12.1	Advantages of Laser Light	426
12.2	Reading with Lasers	433
12.3	Optical Disks and Data Storage	437
12.4	Laser Printing	440
12.5	Lasers in Fiber-Optic Communications	442
12.6	Laser Measurement	447
12.7	Laser Light Shows, Pointers, and Projection Displays	453
12.8	Low-Power Defense Applications	456
12.9	Sensing and Spectroscopy	459
12.10	Holography	464
12.11	Other Low-Power Applications	468
12.12	What Have We Learned?	469
CHAPTER 13	High-Power Laser Applications	475
13.1	High- Versus Low-Power Laser Applications	475
13.2	Attractions of High-Power Lasers	476
13.3	Important Considerations and Trends	477
13.4	Materials Working	481
13.5	Additive Manufacturing and Three-Dimensional Printing	489
13.6	Semiconductor Electronics Fabrication	491
13.7	Laser Medical Treatment	493
13.8	Photochemistry and Isotope Separation	501
13.9	Laser-Driven Nuclear Fusion	503
13.10	High-Energy Laser Weapons	505
13.11	What Have We Learned?	510
CHAPTER 14	Lasers in Research	515
14.1	Lasers Open New Opportunities	515
14.2	Laser Spectroscopy	516
14.3	Manipulating Tiny Objects	521
14.4	Atom Lasers and Bose–Einstein Condensates	522

14.5	Detection of Gravitational Waves	524
14.6	Laser Guide Stars for Astronomy	525
14.7	Slow Light	526
14.8	Nanoscale Lasers	527
14.9	Strange Lasers	529
14.10	Extreme Power Ultrashort Pulse Lasers	530
14.11	X-Ray Free-Electron Lasers	535
14.12	Other Emerging Research	536
14.13	What Have We Learned?	538
Answers to Quiz Questions		543
Appendix A: Laser Safety		547
Appendix B: Handy Numbers and Formulas		553
Appendix C: Resources and Suggested Readings		557
Glossary		561
Index		575

PREFACE

“For Credible Lasers, See Inside.”

THE LASER IS LESS THAN three years younger than the space age. Just days after the Soviet Union launched Sputnik I on October 4, 1957, Charles Townes and Gordon Gould had two crucial discussions at Columbia University about the idea that would become the laser. As the United States and Soviets launched the space race, Townes and Gould went their separate ways and started their own race to make the laser. On May 16, 1960, Theodore Maiman crossed the laser finish line, demonstrating the world’s first laser at Hughes Research Laboratories in California.

Bright, coherent, and tightly focused, laser beams were a new kind of light that excited the imagination. Science fiction writers turned their fictional ray guns into lasers with a stroke of the pen. Science writers inhaled deeply of the technological optimism of the early 1960s and wrote breathless predictions about the future of “the incredible laser.” An article in the November 11, 1962, issue of the Sunday newspaper supplement *This Week* revealed U.S. Army schemes for equipping soldiers with a “death-ray gun ... small enough to be carried or worn as a side-arm.” It quoted Air Force Chief of Staff Curtis E. LeMay predicting that ground-based lasers could zap incoming missiles at the speed of light.

The reality was something else. A bemused Arthur Schawlow, who had worked with Townes on the laser, posted a copy of “The Incredible Laser” on his door at Stanford University, along with a note that read, “For credible lasers, see inside.” Irnee D’Haenens, who had helped Maiman make the first laser, called the laser “a

solution looking for a problem,” a joke that summed up the real situation and became a catchphrase for the young laser industry. The infant laser had tremendous potential, but it had to grow up first.

D’Haenens’s joke lasted many years. So did the popular misconception that lasers were science-fictional weapons. If you told your neighbors you worked with lasers in the 1970s, they inevitably thought you were building death rays. That began to change as supermarkets installed laser scanners to automate checkout in the early 1980s. Then lasers began playing music on compact disks. Laser printers, laser pointers, CD-ROMs, and DVD players followed. Laser surgery became common, particularly to treat eye disease. Surveyors, farmers, and construction workers used lasers to draw straight lines in their work. Lasers marked serial numbers on products, drilled holes in baby-bottle nipples, and performed a thousand obscure tasks in the industry. Lasers transmitted billions of bits per second through optical fibers, becoming the backbone of the global telecommunications network and the internet.

The incredible laser has become credible, a global business with annual sales in the billions of dollars. Lasers have spread throughout science, medicine, and industry. Laser-generated digital signals are the heavy traffic on the fiber-optic backbone of the global information network. Lasers are essential components in home electronics, buried inside today’s CD, DVD, and Blu-Ray players. Laser pointers are so cheap that they are cat toys. It is a rare household that does not own at least one laser, though most are hidden inside electronics or other things. Yet, lasers have not become merely routine; they still play vital roles in Nobel-grade scientific research.

This book will tell you about these real-world lasers. To borrow Schawlow’s line, “For credible lasers, see inside.” It will tell you how lasers work, what they do, and how they are used. It is arranged somewhat like a textbook, but you can read it on your own to learn about the field. Each chapter starts by stating what it will cover, ends by reviewing key points, and is followed by a short multiple-choice quiz.

We start with a broad overview of lasers. Chapter 2 reviews key concepts of physics and optics that are essential to understand lasers. You should review this even if you have a background in physics, especially to check basic optical concepts and terms. Chapters 3 and 4 describe what makes a laser work and how lasers operate. Chapter 5 describes the optical accessories used with

lasers. Try to master each of these chapters before going on to the next.

Chapters 6 to 11 describe various types of lasers. Chapter 6 gives an overview of laser types and configurations and explains such critical concepts as the difference between laser oscillation and amplification, the importance of laser gain, and tunable lasers. Chapter 7 describes the workings of gas lasers and important types such as the helium-neon and carbon dioxide lasers. Chapter 8 covers solid-state lasers, from tiny green laser pointers to giant laboratory systems. Chapter 9 covers fiber lasers, the fastest-growing solid-state laser, now widely used in the industry because of its power and high efficiency, along with fiber amplifiers used in telecommunications. Chapter 10 covers the hot area of semiconductor diode lasers, ranging from tiny chips to powerful pumps for other lasers. Chapter 11 describes other types of lasers, including tunable dye lasers, extreme ultraviolet sources, and free-electron lasers.

The final three chapters cover laser applications, divided into three groups. Chapter 12 describes low-power applications, including communications, measurement, and optical data storage. Chapter 13 covers high-power applications, including surgery, industrial materials processing, and laser weapons. Chapter 14 focuses on research and emerging developments in areas including spectroscopy, slow light, laser cooling, and extremely precise measurements. The appendices, glossary, and index are included to help make this book a useful reference.

To keep this book to a reasonable length, we concentrate on lasers and their workings. We cover optics and laser applications only in brief, but after reading this book, you may want to study them in more detail.

I met my first laser in college and have been writing about laser technology since 1974. I have found it fascinating, and I hope you will, too.

JEFF HECHT

Auburndale, Massachusetts

INTRODUCTION AND OVERVIEW

ABOUT THIS CHAPTER

This chapter will introduce you to lasers. It will give you a basic idea of their use, their operation, and their important properties. This basic understanding will serve as a foundation for the more detailed descriptions of lasers and their operation in later chapters. After a brief introduction to lasers, this chapter will introduce important laser properties and applications.

1.1 LASERS, OPTICS, AND PHOTONICS

To understand lasers, you should first understand where lasers fit into the broader science and technology of light. That field was long called *optics*, but now part of it is sometimes called *photonics*. The differences in the meanings of the two words reflect how the field has changed since the mid-20th century, and understanding those differences will help you understand both lasers and the larger world of light, optics and photonics.

Optics dates back to the origin of lenses in ancient times. It is the science of telescopes, spectacles, microscopes, binoculars, and other optical instruments that manipulate light using lenses, mirrors, prisms, and other transparent and reflective objects. Isaac

Newton famously described the fundamentals of optics in his 1704 book *Opticks*. He thought light was made of tiny particles, but a century later an experiment by Thomas Young indicated light was made of waves, and opinion shifted for a while.

In the late 19th century, physicists discovered that light was a type of electromagnetic radiation, along with radio, infrared, ultraviolet, X-rays, and gamma rays. They differ in the lengths of the waves and in how fast they oscillate. The wavelength and frequency depend on each other because electromagnetic waves always travel at the speed of light. In the early 20th century, Albert Einstein showed that electromagnetic radiation could behave both as particles—called *photons*—and as waves, depending on how you looked at them. The only fundamental difference among electromagnetic waves was their wavelength, which could also be measured as frequency or (photon) energy.

The science and technology of light have also grown increasingly connected with electronics in the past century. Electronic devices can measure light by converting it into electronic signals and measuring them. Television cameras and displays include both optics and electronics. The first electronic circuits used vacuum tubes, but semiconductor devices began replacing tubes in the mid-20th century. That brought a new generation of *electro-optic* devices, including semiconductor electronics that emitted and detected light, converting signals and energy back and forth between photons and electrons.

In the late 20th century, the word *photonics* was coined to describe devices that manipulate photons, like electronics manipulate electrons. The use of the new term became controversial because many people who worked in optics in the field saw it as an attempt to “rebrand” their profession. Photonics has come to refer to things that manipulate light when it acts more like a particle (a photon) than a wave. By that definition, a laser or a sensor that converts light (a series of photons) into an electronic signal is considered photonics, but a lens that refracts and focuses light waves is considered optics. However, that definition remains somewhat hazy. Today, both terms are used, but at this writing, Google tells us that optics remains far ahead, indexed on 622 million web pages, compared to a mere 17.6 million for photonics.

Whatever you want to call the field, you should learn the physical basics of light, optics and photonics, to understand how lasers work. Chapter 2 will go into more detail.

1.2 UNDERSTANDING THE LASER

The laser was born in 1960, long before the word “photonics” came into use. Lasers retain a youthful image, thanks largely to continuing advances in the technology. They vary widely. Some lasers are tremendously sophisticated and incredibly precise scientific instruments costing tens or hundreds of thousands of dollars. Others are tiny semiconductor chips hidden inside optical disk players or pen-shaped red pointers used as cat toys. The world’s biggest laser, the National Ignition Facility at the Lawrence Livermore National Laboratory, cost over a billion dollars and fills an entire building. The tiny lasers inside CD or DVD players are the size of grains of sand and cost pennies apiece. Red laser pointers sell for only a few dollars and are often given away.

We now take many laser applications for granted. For decades, laser scanners at store checkouts have read bar codes printed on packages to tally prices and manage their inventory. Laser pulses carried through optical fibers are the backbone of the global telecommunications network. Builders use laser beams to make sure walls and ceilings are flat and smooth. Offices use laser printers to produce documents. Medical and scientific instruments use lasers to make precise measurements. Lasers cut sheets of metals, plastics, and other materials to desired shapes, so some parts of your car are likely made with a laser. Chapters 12–14 describe many more examples.

Laser light has special properties that make it useful in many ways. You can think of a laser as a very well-behaved light bulb, emitting a narrow beam of a single color rather than spreading white light all around a room. You would not use a laser to illuminate a room, but you can use a tightly focused single-color laser beam to make precise measurements, to transport information around the world at the speed of light, or to cut sheets of metal. Lasers have become tools in industry, medicine, engineering, and science, as well as components in optical systems.

Lasers come in many forms. The most common lasers are tiny semiconductor chips that look like tiny pieces of metallic confetti; untold millions of them are hidden inside electronic devices, measuring devices, and communication systems. Others are glassy or crystalline solids in the form of rods, slabs, or fibers. Some are tubes filled with gases that emit laser light. Some emit light so feeble that the eye can barely detect it; others are blindingly bright; and many

emit infrared or ultraviolet light outside the human visible spectrum. Some perform delicate surgery; others weld sheets of metal. Lasers are used by construction workers installing ceilings and by scientists detecting gravitational waves.

What makes them all lasers is that they generate light in the same way, by a process called “light amplification by the stimulated emission of radiation” that gave us the word “LASER.” We will start by explaining what makes laser light differ from that from the sun, light bulbs, flames, and other light sources.

1.3 WHAT IS A LASER?

Each part of the phrase “light amplification by the stimulated emission of radiation” has a special meaning, so we will look at it piece by piece, starting from the final word.

Radiation means *electromagnetic radiation*, a massless form of energy that travels at the speed of light. It comes in various forms, including visible light, infrared, ultraviolet, radio waves, microwaves, and X-rays. Light and other forms of electromagnetic radiation behave like both waves and particles (called *photons*). You will learn more details in Chapter 2.

Stimulated emission tells us that lasers produce light in a special way. The sun, flames, and light bulbs all emit light *spontaneously*, on their own, in order to release extra internal energy. Lasers contain atoms or molecules that release their extra energy when other light *stimulates* them. You will learn more about that process, called stimulated emission, in Chapter 3.

Amplification means increasing the amount of light. In stimulated emission, an input light wave stimulates an atom or molecule to release its energy as a second wave, which is perfectly matched to the input wave. The stimulated wave, in turn, can stimulate other atoms or molecules to emit duplicate waves, amplifying the light signal more. It may be easier to think of stimulated emission as one light photon tickling or stimulating an atom or molecule so it releases an identical photon, which in turn can stimulate the emission of another identical photon, producing a cascade of photons that amplifies the light.

Light describes the type of electromagnetic radiation produced. In practice, that means not just light visible to the human eye, but also adjacent parts of the electromagnetic spectrum that our eyes

cannot see because it is either longer in wavelength (infrared) or shorter in wavelength (ultraviolet.)

It took decades to put the pieces together. Albert Einstein suggested the possibility of stimulated emission in a paper published in 1917. Stimulated emission was first observed in the 1920s, but physicists long expected it to be much weaker than spontaneous emission. The first hints that stimulated emission could be stronger came in radio-frequency experiments shortly after World War II. In 1951, Charles H. Townes, then at Columbia University, thought of a way to stimulate the emission of microwaves. His idea was to direct ammonia molecules carrying extra energy into a cavity that would reflect the microwaves back and forth through the gas. He called his device a *maser*, an acronym for “microwave amplification by the stimulated emission of radiation.”

It took until 1954 for Townes and his graduate student James Gordon to make the maser work. Some ammonia molecules spontaneously emitted microwaves at a frequency of 24 gigahertz, and that spontaneous emission could stimulate other excited ammonia molecules to emit at the same frequency, building up a signal that oscillated on its own. Alternatively, an external 24-GHz signal could stimulate emission at that frequency from ammonia molecules, to amplify the signal.

In principle, the maser process could be extended to other types of electromagnetic waves, and in 1957, Townes started looking into prospects for an optical version of the maser. Early in his research, he talked with Gordon Gould, a Columbia graduate student who was using light to energize material he was studying for his doctoral research, a then-new idea called *optical pumping*. Townes soon enlisted the help of his brother-in-law, Arthur Schawlow, to work on the optical maser project. Gould, who dreamed of becoming an inventor, quietly tackled the same idea. They essentially solved the same physics problem independently, by placing mirrors at each end of a cylinder so the laser light could oscillate between them. Gould set out to patent his ideas; Townes and Schawlow published their proposal in a scientific journal, *Physical Review Letters*. Their work launched a race to build a laser, which I chronicled in *Beam: The Race to Make the Laser* (Oxford University Press, 2005).

Townes shared in the 1964 Nobel Prize in physics for his pioneering work on “the maser/laser principle.” After a long series of legal battles, Gould earned tens of millions of dollars from his patent claims. However, a third physicist won the race to make the



Figure 1-1. Theodore Maiman and Irnee J. D’Haenens with a replica of the world’s first laser, which they made at Hughes Research Laboratories in 1960. (Reprinted from Hughes Research Laboratories, courtesy of AIP Neils Bohr Library.)

first laser. On May 16, 1960, Theodore Maiman used a photographic flashlamp to excite a fingertip-sized crystal of synthetic ruby to emit pulses of red light from the world’s first laser at Hughes Research Laboratories in Malibu, California. Figure 1-1 shows Maiman and his assistant Irnee D’Haenens holding a replica of his elegant little device.

The ruby laser illustrates how a laser works. Energy from an external source, the lamp, was absorbed by chromium atoms in the ruby cylinder. A few chromium atoms spontaneously emitted photons of red light, which traveled through the ruby. Silver film coated on the ends of the cylinder reflected the red photons back into the ruby, where they stimulated other excited chromium atoms to emit identical photons in the same direction, amplifying the

light. Those photons bounced back and forth between the end mirrors, oscillating (as explained in Box 1.1) within the cavity formed by the two mirrors, with some light emerging through a hole in one coating to form the laser beam. The light was all at the same wavelength, 694 nanometers ($1 \text{ nm} = 10^{-9} \text{ meter}$) at the red end of the visible spectrum. It was also coherent, with all the waves aligned with each other and marching along in step like soldiers on parade.

Maiman's laser emitted a pulse of laser light every time the flashlamp fired, and pulsed operation proved attractive for some uses. Other lasers generated a continuous beam, which was attractive for other purposes. New laser materials followed, including crystals and glasses containing various light-emitting elements, tubes filled with mixtures of light-emitting gases, and tiny chips of semiconductor compounds such as gallium arsenide or gallium nitride, which today are the world's most common lasers.

BOX 1.1 LASER OSCILLATION

Stimulated emission amplifies light in a laser, but the laser itself is called an oscillator because it generates a beam on its own rather than amplifying light from an outside source. So you may wonder why the word “laser” comes from “light *amplification* by the stimulated emission of radiation”? There's an interesting bit of history behind that.

Charles Townes created the word “maser” as an acronym for microwave amplification by the stimulated emission of radiation. When he began thinking of a version of the maser that used light, he called it an optical maser. When Gordon Gould sat down to tackle the same problem, he wrote “laser” at the top of his notes, coining the acronym for light amplification by the stimulated emission of radiation. As the competition between Townes and Gould became intense, each side pushed its own term.

Arthur Schawlow was a jovial soul, and at one conference pointed out that because the laser was actually an oscillator, it should be described as “light oscillation by the stimulated emission of radiation,” making the laser a “loser.” Everybody laughed, but the word laser proved to be a winner.

Lasers operate at wavelengths from the far infrared all the way to X-rays. They can generate modest powers far below one watt, steady powers of thousands of watts, or concentrate light into pulses lasting less than a trillionth of a second. Chapters 2 through 5 will describe the basics of laser physics in more detail. Chapters 6 through 11 will describe various types of lasers, and Chapters 12 through 14 will explain important and important uses of lasers. But first let's take a quick look at various types of lasers and their properties.

1.4 LASER MATERIALS AND TYPES

Laser performance depends strongly on the materials from which they are made. Maiman won the laser race because he knew the optical properties of ruby and designed his laser to take advantage of them.

The ruby laser worked because Maiman used a flash lamp to produce a bright pulse of visible light that excited chromium atoms in the ruby rod to a higher energy level. The chromium atoms remained in that high energy level until they released their energy as red light and dropped to a lower energy level. Some of those photons then stimulated emission from other chromium atom, which also emitted red light, producing a cascade of red light that became a laser pulse. Figure 1-2 shows the basic idea.

Ruby is an example of a *solid-state laser*, in which light-emitting atoms are distributed in a transparent solid. In ruby, the transparent material is sapphire (aluminum oxide or Al_2O_3) and the light-emitting atoms are chromium. Such transparent solids do not conduct electric current, so the light-emitting atoms must be excited by light from an external source, such as a flash lamp or another laser, a process called *optical pumping*. Typically small quantities of light-emitting elements such as neodymium, erbium, and ytterbium are added to transparent crystals, glasses, and ceramics, which are shaped into rods, thin disks, slabs, or optical fibers for use in lasers.

Some solid-state lasers are excited with flash lamps or with bright lamps that emit continuously. Others are excited by light from other lasers, usually semiconductor diodes. Chapter 8 describes solid-state lasers in more detail. Chapter 9 describes fiber




Figure 1-2. A single spontaneously emitted photon triggers stimulated emission from excited atoms, building up a cascade of stimulated emission. In ruby, the excited atoms are chromium.

lasers, a type of solid-state laser distinct and important enough to deserve their own chapter.

A second broad class of lasers are *gas lasers*, covered in Chapter 7, in which a light-emitting gas or vapor is confined inside a hollow tube with mirrors on the ends. Passing an electric discharge through the gas excites the atoms to states in which they can generate stimulated emission. Important examples are the helium–neon, rare-gas-halide and carbon dioxide lasers, described in more detail in Chapter 7. Gas lasers have been replaced by solid-state lasers for many applications but remain in use for others.

A third broad class are *semiconductor lasers*, which, in the laser world, are considered distinct from solid-state lasers. Most semiconductor lasers are called *diode lasers* or *laser diodes* because they have two electrical terminals, and current flows in only one direction between the terminals to generate stimulated emission inside the semiconductor, as you will learn in Chapter 10. Semiconductor lasers are versatile devices that can play many roles in laser technology. Some are tiny, cheap, and low-power devices used in CD, DVD, and Blu-Ray players, and laser pointers. Others are larger devices that emit hundreds of watts and can convert more than half of the electrical energy passing through them into light, for use in pumping solid-state lasers or in some industrial applications.

Chapter 11 describes a few lasers that have found important applications but do not fall into the broad categories of gas, solid-state, fiber, or semiconductor lasers. It also covers a number of light sources that generate laser-like light but do not exactly fall under the definition of lasers. Many are nonlinear devices that shift laser light to other wavelengths in various ways.

1.5 OPTICAL PROPERTIES OF LASER LIGHT

The practical importance of lasers comes from the unusual properties of light in a laser beam. These properties are crucial for applications of lasers ranging from cutting sheets of plastic or metal to making extremely precise and sensitive measurements in scientific research. The most important of these optical properties are:

- Wavelength(s)
- Beam power and energy
- Variation of beam power with time (e.g., pulse duration)
- Beam divergence and size
- Coherence
- Efficiency

1.5.1 Wavelength(s)

Most lasers are called *monochromatic*, meaning single-colored, but that single wavelength generally can be adjusted a little or a lot, depending on the light-emitting material in the laser and on the optics used in the laser. The laser material determines the range of possible wavelengths; the optics select which of those the laser emits. The details can become complicated and are covered in Chapters 3 and 4.

Lasers typically operate in the ultraviolet, visible, and infrared parts of the spectrum. Table 1-1 lists some important lasers emitting in that range, their primary wavelengths, and the chapters that describe them. In addition, it is possible to generate additional wavelengths from these lasers, some of which can be quite important, such as the 532-nanometer green line produced by generating the second harmonic of the 1064-nm line of neodymium, as described in Section 5.6. Thousands of other laser lines have been demonstrated in the laboratory, but most are not used regularly.

Table 1-1. Some important lasers and their wavelengths

Laser name and type	Wavelengths	Chapter
Argon fluoride (ArF) gas, excimer	193 nm	7
Krypton fluoride (KrF) gas, excimer	248 nm	7
Organic dye, liquid	320–1000 nm	11
Nitride diode (InGaN), semiconductor	375–525 nm	10
Argon ion, gas	488, 514.5 nm	7
Helium–neon, gas	632.8 nm	7
InGaAlP, semiconductor	635–660 nm	10
GaAsP, semiconductor	670 nm	10
Titanium–sapphire, solid-state	700–1000 nm	8
GaAs/GaAlAs, semiconductor	780–905 nm	10
InGaAs, semiconductor	915–980 nm	10
Ytterbium, fiber, solid-state	1030–1080 nm	9
Neodymium, solid-state	1060, 1064 nm	8
InGaAsP, semiconductor	1150–1650 nm	10
Erbium, fiber, solid-state	1530–1600 nm	8, 9
Quantum cascade, semiconductor	4–12 μm	10
Carbon dioxide (CO ₂)	9–11 μm	7

1.5.2 Beam Power, Energy, and Intensity

Power is a critical quantity for laser beams, and it can be measured in three different ways that give distinctly different information.

Power measures the rate of energy delivery by a laser beam. It is important to remember that power is the amount of energy delivered per unit time. It is defined by the formula:

$$\text{Power} = \frac{\Delta \text{ energy}}{\Delta \text{ time}} \quad (1-1)$$

One *watt* of power equals one *joule* (of energy) per second. Strictly speaking, power measures how fast energy is being delivered at any given instant, so it varies with time for pulsed lasers, but is nominally constant for continuous lasers. If a laser emits a series of pulses, it can also be measured by its *average power*, the sum of the pulse energies divided by the time covered. The powers of continuous laser beams range from less than a milliwatt (0.001 watt) to over a hundred kilowatts (100,000 W), and the average powers of repetitively pulsed lasers are similar. *Peak power* measures the maximum rate of power delivery during a laser pulse and can reach much higher levels.

Energy in joules measures the total amount of energy delivered during an interval. Typically, it measures the energy delivered by a single laser pulse. The shorter the time the laser takes to deliver a given energy, the higher the peak power.

Intensity measures the power deposited per unit area. The smaller the laser spot, the higher the intensity and the more it affects what it illuminates. Think of how bright sunlight may warm a piece of paper, but focusing sunlight through a magnifier can heat the paper so it burns in a small spot.

All of these quantities are important and will be discussed in more detail later.

1.5.3 Laser Variations in Time

Some lasers can emit continuous beams, but others are limited to emitting pulses because of their internal physics. Continuous lasers can be turned on and off by modulating their output in some way. The details differ among laser types, and we will explain them when we cover the individual laser types. Inherently pulsed lasers typically fire a series of pulses at regular intervals, but some fire only a single laser shot at a time.

The length of laser pulses can vary widely, ranging from milliseconds (10^{-3} second) to femtoseconds (10^{-15} second). The pulse timing and spacing may depend on the physics of the laser, but these can also be controlled by the operator. One approach is to modulate the input power so the laser switches on and off, as when you turn a laser pointer on and off. Another way to control output is by using optical accessories described in Chapter 5. Modulating laser output can be very important and will be explained later in this book.

1.5.4 Beam Appearance, Divergence, and Size

You cannot see a laser beam in the air unless something reflects the light toward you, such as smoke or fog in the air or the beam hitting a wall or your hand. When the beam emerges from the laser, it has a diameter that depends on the size of the output optics. For small lasers such as red laser pointers, this typically is a millimeter or two and looks as thin as a string or a pencil line.

Although a laser beam looks straight to your eyes, it actually spreads at a very small angle, called the *divergence*, which is shown in Figure 1-3. The divergence depends both on the type of the laser




Figure 1-3. Calculating the size of a laser spot from the beam divergence.

and on the external optics. Most semiconductor lasers have beams that spread widely, but laser pointers typically have optics added to focus the beam into a narrow pencil-like beam.

Typically, laser beam divergence is measured in thousandths of a radian, a milliradian, a unit equal to 0.057 degree. As long as the beam divergence is small, you can estimate the radius of a laser spot at a distance D from the laser by multiplying the distance by the divergence in radians. Thus, a 2-milliradian beam spreads to a 0.2 meter spot at a distance of 100 meters. This high directionality of the laser beam is important for many applications.

1.5.5 Coherence

Stimulated emission makes laser light coherent because output photons have the same wavelength and phase as the input photons that stimulate emission. This makes laser light *coherent*, with the peaks and valleys of the waves marching in phase like soldiers on parade. Figure 1-4 compares coherent and incoherent light. The peaks and valleys of coherent light waves (top of Figure 1-4) are all the same length and have their peaks and valleys aligned. The peaks and valleys of incoherent light waves (bottom of Figure 1-4) do not line up. Laser light gets its coherence from stimulated emission. The sun, light bulbs, flames, and most other sources generate spontaneous emission, and their output is incoherent.

The degree of coherence depends on the range of wavelengths emitted, which differs among lasers. A laser that emits only a single wavelength, called *monochromatic*, generally is more coherent than a laser emitting a broader range of colors. Monochromatic light need not be coherent, but light that is not monochromatic cannot stay coherent over a long distance. Lasers are the only light sources that can readily generate light that is coherent over relatively long distances of centimeters and up.




Figure 1-4. Coherent (A) and incoherent (B) light.

1.5.6 Energy Conversion Efficiency

Lasers convert other forms of energy into laser light. This conversion efficiency can be very important for some laser applications, and many advances of recent years come from improvements in efficiency. In some early gas lasers, as little as 0.001% of the electrical power that went into the laser emerged as light in the output beam. What made the laser light valuable was that its beam was tightly focused, coherent, and monochromatic.

Many modern lasers convert 10% to 70% of input energy into laser light, and that is vital for applications that require large amounts of laser power, such as cutting and welding metals, or exciting other lasers. Semiconductor diode lasers can convert as much as 70% of the electrical energy passing through them into light. Solid-state fiber lasers can convert over 70% of the light energy powering them into a high-quality laser beam, making them particularly well suited for industrial machining applications.

1.6 HOW LASERS ARE USED?

Scientists and engineers began playing with lasers almost as soon as they could lay their hands on them. They fired laser pulses at just about everything that could not run away. They shot so many holes

in razor blades that for a while laser power was informally measured in “gillettes.” Yet few practical applications emerged quickly, and for a while the laser seemed to be, as Irnee D’Haenens told Ted Maiman soon after they made the first one, “a solution looking for a problem.”

We are long past that stage. Lasers have become standard tools in industry and research. They align construction equipment, transmit voice and data around the globe, and perform exquisitely sensitive measurements that have earned a fair number of Nobel Prizes. Table 1-2 lists some laser applications, and Chapters 12–14 cover laser applications in more detail.

Lasers are used in diverse ways. The final three chapters divide laser applications into three broad categories.

Chapter 12 covers low-power applications. One broad family of such applications uses lasers as sources of highly controlled light for transmitting and processing information, such as reading or writing data or transmitting signals. Laser light transmitted through hair-thin optical fibers is the backbone of the global telecommunications network; it carries phone calls from the cell tower nearest you to anywhere around the world. Lasers inside optical disk players read music from CDs and videos on DVDs and Blu-Ray disks. The coherence of lasers makes it possible to create and display three-dimensional holographic images.

Another broad category of low-power applications is measurement. Laser beams can draw straight lines in space to help construction workers align walls or pipes. Precision techniques use the coherence of lasers to measure distances to within a fraction of the wavelength of light. Laser radars can create three-dimensional profiles of our environment, digitizing dinosaur fossils for paleontologists and helping to steer self-driving cars away from potential roadside hazards.

Chapter 13 covers high-power applications, in which a laser beam delivers energy that alters the material it hits. Lasers deliver small bursts of energy to mark painted metal surfaces; the laser vaporizes the paint, exposing the shiny metal. More powerful lasers can drill holes through materials ranging from baby-bottle nipples to sheets of titanium. The laser beam does not bend soft materials like latex nipples and does not grow dull like a drill bit cutting into a hard material.

Laser surgery works in the same way. Pulses from an ultraviolet laser can vaporize tissue from the lens of the eye, precisely removing just the right amount to correct vision defects. By selecting

Table 1-2. A sampling of laser applications

Information handling
Fiber-optic communications
Laser printers for computer output
Playing DVD or Blu-Ray video
Playing CD audio
Reading and writing computer data on CDs and DVDs
Reading printed bar codes for store checkout and inventory control
Measurement and inspection
Exciting fluorescence from various materials
Illuminating cells for biomedical measurements
Measuring concentrations of chemicals or pollutants
Measuring small distances very precisely
Measuring the range to distant objects
Measuring velocity
Projecting straight lines for construction alignment and irrigation
Studies of atomic and molecular physics
Helping guide self-driving cars
Medicine and dentistry
Bleaching of port-wine stain birthmarks and certain tattoos
Clearing vision complications after cataract surgery
Dentistry
Refractive surgery to correct vision
Reattaching detached retinas
Shattering of stones in the kidney and pancreas
Treatment of diabetic retinopathy to forestall blindness
Surgery on tissue rich in blood vessels
Materials working
Cutting, drilling, and welding plastics, metals, and other materials
Cutting titanium sheets
Non-contact machining
Drilling materials from diamonds to baby-bottle nipples
Engraving wood
Heat-treating surfaces
Marking identification codes
Semiconductor photolithography
Three-dimensional printing or additive manufacturing
Military
Range-finding to targets
Simulating effects of nuclear weapons
Target designation for bombs and missiles
War games and battle simulation
Antisatellite weapons
Anti-sensor and antipersonnel weapons
Defense against rockets, artillery, mortars, drones, and small boats

Table 1-2. (Continued)

Other applications
Basic research
Controlling chemical reactions
Theater displays
Holography
Laser light shows
Laser pointers
Laser paint removal from aircraft

the right laser wavelength, surgeons can bleach dark birthmarks or tattoos.

The ultimate in high-power lasers are high-energy laser weapons. You can think of them as performing materials working on unfriendly objects. A laser weapon might blind the sensor that guides a missile, causing it to go astray. Or a higher-energy laser weapon can heat explosives in a rocket to their detonation temperature, soften high-pressure fuel tanks so they fail, or ignite gasoline fumes to catch an engine on fire. Lasers can also detonate unexploded shells left on a battlefield, or defend ships against attacks by drones or small boats.

Chapter 14 covers laser applications in scientific research. Laser techniques can slow atoms to a virtual crawl and probe their energy states with exquisite precision. Laser beams can manipulate tiny objects, from bacteria to single atoms. These laser applications have led to many Nobel Prizes.

1.7 WHAT HAVE WE LEARNED?

- Optics is the science of light. Photonics is another term for optics, usually covering devices that work on photons rather than on light waves.
- LASER is an acronym for “light amplification by the stimulated emission of radiation.”
- Stimulated emission of light by excited atoms generates laser radiation.
- Most lasers are tiny semiconductor chips.
- Lasers have become so commonplace in many places.
- Charles Townes conceived of the amplification of stimulated emission for microwaves.

- Theodore Maiman demonstrated the first laser using a ruby rod pumped by a photographic flashlamp.
- Successful operation of a laser requires both an optical resonator and a suitable gain medium to amplify light.
- The three main classes of lasers are gas, semiconductor, and solid-state lasers.
- Fiber lasers are a type of solid-state laser in which the laser is an optical fiber.
- Solid-state is not equivalent to semiconductor in the laser world.
- Lasers can emit a very narrow range of wavelengths.
- Laser light is concentrated in a beam, which is generally tightly focused.
- Laser light is coherent.
- Low-power laser applications include measurement and information processing.
- High-power laser applications modify materials for tasks including surgery, machining, and weapons.
- Lasers can make precision measurements for scientific research.

WHAT'S NEXT?

The first step in understanding lasers is to learn the basic principles of physics and optics that are involved in laser operation. Chapter 2 introduces the essential physical concepts. Some of this material may be familiar if you have been exposed to physics, but you should review it because later chapters assume that you understand it.

QUIZ FOR CHAPTER 1

1. The word laser originated as
 - a. A military code word for a top-secret project
 - b. A trade name
 - c. An acronym for Light Amplification by the Stimulated Emission of Radiation
 - d. The German word for light emitter
2. Which statement is not true
 - a. Light sometimes acts as a wave and sometimes acts as a photon
 - b. Only laser light acts as photons
 - c. Light is a form of electromagnetic radiation

- d. Visible light differs from ultraviolet in wavelength
 - e. Electromagnetic radiation includes radio waves
3. Most lasers today are
 - a. Semiconductor devices used in electronic equipment
 - b. High-power weapons used to deter drone attacks
 - c. Gas-filled tubes emitting red light
 - d. Ruby rods powered by flash lamps
 - e. Ruby rods powered by LEDs
 4. Laser light is generated by
 - a. Spontaneous emission
 - b. Gravity
 - c. Stimulated emission
 - d. Microwaves
 - e. Mirrors
 5. What emits light in a ruby laser?
 - a. Aluminum atoms
 - b. Sapphire atoms
 - c. Oxygen atoms
 - d. Chromium atoms
 - e. Mirrors on the ends of the rod
 6. Why are most semiconductor lasers sometimes often called diode lasers?
 - a. Because the first diode lasers had to be installed in vacuum tubes so the semiconductor would not evaporate.
 - b. Because they conduct light between two terminals to generate stimulated emission.
 - c. Because it is powered by light from an external light-emitting diode.
 - d. Because it is an acronym for “damn idiotic optical device exploded,” which is what happened to the first one.
 7. You have a laser that emits one watt of light and is 1% efficient. How much of the input power ends up as heat rather than light?
 - a. One watt
 - b. Three watts
 - c. 10 watts
 - d. 30 watts
 - e. 99 watts
 8. You have a laser that emits one watt of light and is 25% efficient. How much of the input power ends up as heat rather than light?
 - a. One watt
 - b. Three watts

- c. 10 watts
 - d. 30 watts
 - e. 99 watts
9. Stimulated emission generates light waves that are in phase with each other. This makes them
- a. A beam
 - b. Coherent
 - c. Pulsed
 - d. Span a range of wavelengths
10. How many lasers do you own? There is no single “right” answer, but it can be fun to take a mental inventory. Do not forget that some devices may contain multiple lasers, such as a Blu-Ray player that can also play DVDs and CDs.

CHAPTER 2

PHYSICAL BASICS

ABOUT THIS CHAPTER

Lasers evolved from concepts of modern physics that emerged early in the 20th century. To understand lasers, you need to understand basic concepts including light, atomic energy levels, quantum mechanics, and optics. This chapter starts with the nature of light, then moves on to how light is generated, the interactions of light and matter, and some fundamentals of optics and fiber optics, to give you the background you need to understand lasers themselves.

2.1 ELECTROMAGNETIC WAVES AND PHOTONS

Early physicists debated long and loud over the nature of light. In the 17th century, Christian Huygens believed light was made up of waves, but Isaac Newton held that light was made up of tiny particles. Newton's theory was thought to be right for a century, until Thomas Young showed the interference of light, which could only be explained by waves.

Today, we know that both theories are partially right. Much of the time, light behaves like a wave. Light is called an *electromagnetic* wave because it consists of electric and magnetic fields perpendicular to each other, as shown in Figure 2-1. Because the electric and magnetic fields oscillate perpendicular to the direction in which the waves travel, they are called *transverse waves*.

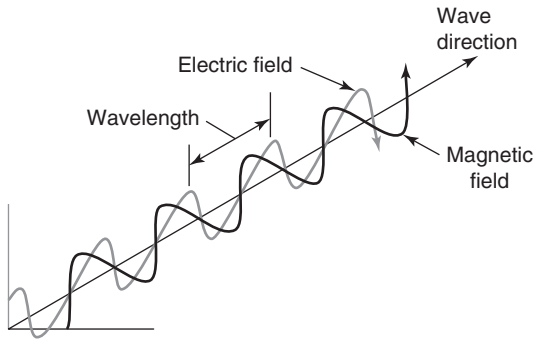


Figure 2-1. Structure of an electromagnetic wave.

At other times, light behaves like massless particles called *photons*. A photon is a quantum of electromagnetic energy, not the hard particle that Newton envisioned. The quantum view that light comes in discrete chunks rather than a continuous wave is also critical to understanding light. We will move back and forth between the wave and particle views of light, but you should remember that they are just different ways to look at the same thing. We also include some simple formulas, explaining them in words to help you understand what they mean. You should remember these ideas if you have had a physics course.

2.1.1 Waves and Photons

We describe electromagnetic waves by one of two related values. The *wavelength* is the distance between successive peaks, shown in Figure 2-1, and denoted by the Greek letter lambda (λ). The *frequency* is the number of wave peaks passing a point in a second, denoted by the Greek letter nu (ν). Frequency is measured in oscillations (from peak to valley and back) per second, a unit called hertz, after Heinrich Hertz, the 19th-century discoverer of radio waves. Multiplying the wavelength by the number of waves per second gives the wave velocity, as a distance per second:

$$\text{Wavelength} \times \text{Frequency} = \text{Velocity} \quad (2-1)$$

The velocity of light and other electromagnetic waves is a universal constant, denoted by c , the speed of light in vacuum, which is close to 300,000 kilometers per second or 186,000 miles

per second. (The exact value is 2.99792458×10^8 meters/second, listed in Appendix C, but the round number usually is good enough for practical purposes.) This means the shorter the wavelength, the higher the frequency. Using the standard symbols for wavelength, frequency, and the speed of light, we can write simple formulas to convert between wavelength and frequency in a vacuum:

$$\lambda = c/\nu \quad (2-2)$$

and

$$\nu = c/\lambda \quad (2-3)$$

We can also view light as *photons* or *quanta*, massless particles that travel at the speed of light. The energy of a photon E equals the frequency of the wave multiplied by a constant h , which is called *Planck's constant* after the German physicist Max Planck, who worked out the formula. Planck's constant equals 6.63×10^{-34} joule-second, so photon energy is

$$E = h\nu = 6.63 \times 10^{-34} \nu \quad (2-4)$$

The frequency ν is measured in waves per second, so multiplying it by the number of waves (sometimes called cycles) gives energy per photon in the same joule-second units as Planck's constant.

You can use this formula and the relationship between wavelength and frequency to show that wavelength times photon energy equals Planck's constant multiplied by the speed of light c , or 1.99×10^{-25} joule-meter:

$$\lambda E = hc = 1.99 \times 10^{-25} \text{ joule-meter} \quad (2-5)$$

You may need these conversion factors from time to time when you work with light and other electromagnetic waves, so they are listed in Appendix B. For now, the most important lesson is that the nature of a light wave can be measured in three ways: photon energy, wavelength, or frequency. For example, light with a 1- μm wavelength roughly has a frequency of 3×10^{14} Hz, and photon energy of 2×10^{-19} joule.

Physicists often measure photon energy in electron volts, the energy an electron acquires by moving through a one-volt potential. One electron volt equals 1.6022×10^{-19} joule, so a photon energy of

2×10^{-19} joule equals 1.24 electron volts, a more convenient unit of measurement.

It is easy to make mistakes in converting units, so it helps to remember these simple rules of thumb:

- The higher the frequency, the shorter the wavelength.
- The higher the frequency, the larger the photon energy.
- The shorter the wavelength, the larger the photon energy.

Electromagnetic waves are often called *electromagnetic radiation* because objects emit or *radiate* them into space. If the word “radiation” sounds distressingly like something that comes from a leaky nuclear reactor, it is. However, that loses the vital difference between nuclear and electromagnetic radiation. The electromagnetic spectrum spans a broad range of photon energies, and only the most energetic photons—X-rays and gamma rays, which have the highest frequencies and shortest wavelengths—pose any radiological threat. Ultraviolet light with longer wavelengths and less energy can cause sunburn with long-term risk of skin cancer.

2.1.2 The Electromagnetic Spectrum

We usually think of the spectrum as the colors we see in a rainbow or those that appear when we pass sunlight through a prism. Colors are the way our eyes and brains sense the differences in wavelength across the visible spectrum, from the short violet waves to the long red waves. However, our eyes can see only a narrow slice of the whole *electromagnetic spectrum*. Table 2-1 lists the components of

Table 2-1. Wavelengths and frequencies of electromagnetic radiation. Boundaries are approximate.

Name	Wavelengths (m)	Frequencies (Hz)
Gamma rays	under 3×10^{-11}	over 10^{20}
X-rays	3×10^{-11} to 10^{-8}	3×10^{16} to 10^{20}
Ultraviolet light	10^{-8} to 4×10^{-7}	7.5×10^{14} to 3×10^{16}
Visible light	4×10^{-7} to 7×10^{-7}	4.2×10^{14} to 7.5×10^{14}
Infrared light	7×10^{-7} to 10^{-3}	3×10^{11} to 4.2×10^{14}
Microwaves	10^{-3} to 0.3	10^9 to 3×10^{11}
Radio waves	0.3 to 30,000	10^4 to 10^9
Low-frequency waves	over 30,000	under 10,000

Table 2-2. Prefixes used for metric units. Note that abbreviations for mega and above are capital letters.

Prefix	Abbreviation	Meaning	Value
exa	E	quintillion	10^{18}
peta	P	quadrillion	10^{15}
tera	T	trillion	10^{12}
giga	G	billion	10^9
mega	M	million	10^6
kilo	k	thousand	1,000
deci	d	tenth	0.1
centi	c	hundredth	0.01
milli	m	thousandth	0.001
micro	μ	millionth	10^{-6}
nano	n	billionth	10^{-9}
pico	p	trillionth	10^{-12}
femto	f	quadrillionth	10^{-15}
atto	a	quintillionth	10^{-18}

that spectrum, ranging from extremely low-frequency waves many miles long to gamma rays less than a trillionth of a meter long.

Table 2-1 lists all frequency and wavelength values in scientific units to simplify comparison over orders of magnitude. In practice, the values of frequency, wavelength, and other quantities such as time and power are expressed in metric units with the standard prefixes shown in Table 2-2. You probably already know some of these prefixes, but you are likely to discover others as you explore the world of lasers. Virtually everything optical is measured in metric units, and this book follows that practice. (Some older books give visible wavelengths in Ångstroms, Å, a unit equal to 10^{-10} meter or 0.1 nanometer, but the Ångstrom is not a standard metric unit.)

Parts of the electromagnetic spectrum were discovered at different times, and the boundaries between them remain poorly defined. Radio waves, X-rays, and gamma rays were discovered separately, and only later did physicists realize that all were electromagnetic waves. Even the limits of human visibility are not rigidly defined because the eye's sensitivity drops slowly at the red and violet ends of the spectrum, as shown in Figure 2-2. The human visible range usually is defined as 400 to 700 nanometers, but the human eye can faintly sense wavelengths from 380 to 780 nm.

In practice, the terms "light" and "optical" are not limited to visible wavelengths. Parts of the infrared and ultraviolet are also

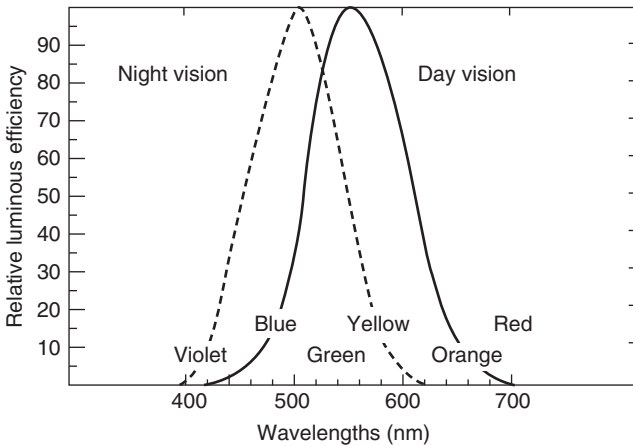


Figure 2-2. Relative sensitivity of the human eye to different wavelengths during daylight when we see colors, and at night when only one visual receptor is active.

called optical or light because those invisible wavelengths act much like visible light.

The electromagnetic spectrum is a continuum spanning many orders of magnitude in frequency and wavelength. The divisions that we make depend largely on how waves in different parts of the spectrum interact with matter. For example, wavelengths shorter than about 200 nm are called “vacuum ultraviolet” because air absorbs them so strongly that they must be observed in a vacuum. Section 2.3 tells more about how light interacts with matter.

Lasers normally operate in the infrared, visible, and ultraviolet bands. Masers, the microwave counterparts of lasers, operate at microwave frequencies, which are lower than those in the infrared. The few research lasers that operate in the X-ray band are described in Sections 11.6 and 14.11.

2.1.3 Interference and Waves

As you learned in Section 2.1.1, light has a dual personality, sometimes behaving as waves and sometimes as particles. The best illustration of the wave nature of light is the interference of waves seen when light from one source passes through two parallel slits. As shown in Figure 2-3, waves emerging from the two slits combine to form a pattern of alternating light and dark bands. That pattern




Figure 2-3. Bright light illuminating two slits causes interference.

comes from adding the amplitudes of the waves emerging from the two slits when they reach the screen at the back. The light is bright where two peaks or two valleys overlap and add together, a process called *constructive interference*. But the light vanishes where a peak of one wave meets a valley of the other, so they cancel each other to leave no light, a process called *destructive interference*. Figure 2-4 shows how the two waves add up in both cases.

Interference is a general effect and occurs whenever waves combine. If the two waves are precisely in phase, they add together without loss. If two equal-amplitude waves are exactly one-half wavelength (180 degrees) out of phase, they cancel each other out at that point. In either case, the waves keep on going unless something blocks them, so they can interfere with other waves at other points in space.

No light energy is really lost in destructive interference. If nothing blocks the waves, they just keep on going as if they had never encountered each other, like ripples spreading across a pond. If something does block the light so you see a pattern of light and dark zones, the total amount of energy across the pattern remains the same, but interference changes its distribution.

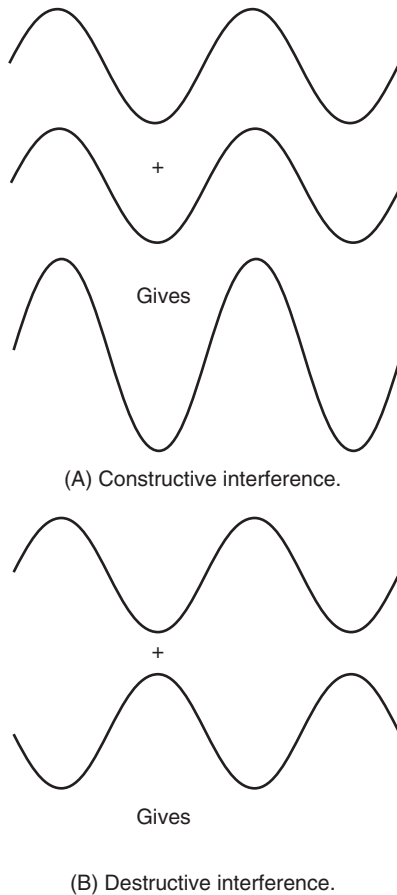


Figure 2-4. Addition and subtraction of light-wave amplitudes causes interference. (A) Constructive interference. (B) Destructive interference.

Interference is a very important property of light, and it is particularly important for lasers because their coherent light remains in phase over relatively long distances.

2.1.4 Light as Photons

The photon side of light's dual personality shows most clearly when light interacts with matter.

One of the most famous examples is called the *photoelectric effect*, which causes certain metals to emit electrons if light strikes them in a vacuum. Initial experiments were puzzling because they

showed a peculiar dependence on the wavelength. The metal did not emit electrons when illuminated by long wavelengths, even if the light was very bright. But at wavelengths shorter than a threshold level that differed among metals, the metal emitted a number of electrons that depended on the light intensity. That is, doubling the illumination doubled the number of electrons released.

This did not make sense if light was purely a wave; adding more energy at a long wavelength eventually should build up enough energy to free an electron. But in 1905, Albert Einstein explained that the threshold occurred because light energy was bundled as photons, and the photons had to have enough energy to free the electrons. Once you had a wavelength short enough to get a photon with enough energy to free the electron, adding more light could free more electrons. Einstein's explanation of the photoelectric effect helped lay the groundwork for quantum mechanics, and it eventually earned him the Nobel Prize in Physics.

2.2 QUANTUM AND CLASSICAL PHYSICS

So far, we have mostly considered light, but laser physics is not just about light; it is about how we use matter to produce a beam of coherent light. That means we need to look at quantum mechanics and the atom.

The operation of a laser depends on the quantum mechanical properties of matter. The classical physics described by Isaac Newton assumed that energy could vary continuously, like an absolutely smooth liquid. In contrast, quantum mechanics tells us that energy comes in discrete chunks called *quanta*, so everything in the universe can have only discrete amounts of energy. In classical physics, energy can change continuously; in quantum mechanics, it changes in steps. The wave picture of light is classical; the photon picture is quantum mechanical.

The central concern for the laser is the quantization of energy levels within atoms. We will start with atomic energy levels, then look at how that leads to the atomic physics behind the laser.

2.2.1 Energy Levels

The simplest atom is hydrogen, in which one electron circles a nucleus that contains one proton. The hydrogen atom looks like




Figure 2-5. Electron orbits and the corresponding energy levels of the hydrogen atom.

a very simple solar system, with a single planet (the electron) orbiting a star (the proton). The force holding the atom together is not gravity but the electrical attraction between the positive charge of the proton and the negative charge of the electron.

In a real planetary system, the planet could orbit at any distance from the star. However, quantum mechanics allows the electron in the hydrogen atom to occupy only certain orbits, shown in Figure 2-5. We show the orbits as circles for simplicity, but we cannot see the exact shape of the orbits. Their nominal sizes depend on a “wavelength” assigned to the electron because matter, too, sometimes can act like a wave. The circumference of the innermost orbit equals one wavelength, the next orbit equals two wavelengths, and so on.

If we added energy to a planet, it would speed up and move further from its star. The electron in a hydrogen atom can also move to more distant orbits when it absorbs extra energy. However, unlike a planet, the electron can occupy only certain orbits, usually called energy levels, which are plotted at the side of Figure 2-5, with labels indicating the corresponding orbits. The atom is in its lowest possible energy level—the ground state—when the electron is in the innermost orbit, closest to the proton. (The electron cannot fall onto the proton.)

The spacing between energy levels in the hydrogen atom decreases as the energy levels become higher above the ground state and, eventually, the energy difference becomes vanishingly small. If the electron gets enough energy, it escapes from the atom altogether, a process called *ionization*.

Physicists often define the energy of an ionized hydrogen atom as zero so they can write the energy of the un-ionized hydrogen atom E as a negative number using the simple formula

$$E = -R/n^2 \quad (2-6)$$

where R is a constant (2.179×10^{-18} joule) and n is the quantum number of the orbit (counting outward, with one being the innermost level).

Hydrogen is the simplest atom in the universe, so physicists use it as a model to explain quantum mechanics and atomic energy levels. Naturally, things get considerably more complex in atoms with more electrons. Each electron occupies an energy state that is specified by multiple quantum numbers, which physicists interpret as identifying quantities including the shell, the subshell, position in a shell, and spin. You do not need to know the details to understand the basics of laser physics, so we will skip them.

One thing that is important is a rule called the Pauli exclusion principle, which says that each electron in an atom must occupy a unique quantum state, with a unique set of quantum numbers. Shells and subshells usually are identified by a number and letter, such as $1s$ or $2p$. The number is the primary quantum number. The letter identifies subshells formed under complex quantum mechanical rules. Each shell can contain two or more electrons, but each electron must have a unique set of quantum numbers. In the simplest case of two-electron shells, the two electrons spin in opposite directions.


If atoms are in their ground state, electrons fill in energy levels from the lowest-energy or innermost shell until the number of electrons equals the number of protons in the nucleus. The primary quantum number of the last ground-state electron corresponds to the row the element occupies on the periodic table. Figure 2-6 shows how electrons fill up a series of energy levels for neodymium, which has an atomic number 60 and is widely used in lasers.

These energy levels are the same ones that play a crucial role in chemistry, determining the chemical behavior of the elements and

60

Neodymium

Nd



Atomic mass: 144.24

Electron configuration: 2, 8, 18, 22, 8, 2

Figure 2-6. Electron shells of neodymium, which has 60 electrons. (iStock)

the structure of the periodic table. Electrons generally fill up the shells from the inside out, although it gets more complex for rare-earth elements like neodymium. In lasers, as in chemistry, the outermost electrons are the easiest ones to excite, and the extra energy takes them to higher energy levels not shown in Figure 2-6.

2.2.2 Transitions and Spectral Lines

The transitions that atoms and molecules make between energy levels are crucial to laser physics. To understand their importance, let us look again at the neodymium atom.

The electron needs an extra increment of energy to move from the ground state to a higher energy level. Conversely, the electron must release energy to drop from a higher level to a lower one. Electrons can do this by absorbing or emitting photons with energy equal to difference between the two states involved in the transition. (They can also transfer energy in other ways.)

To move one step up the energy-level ladder, an electron must absorb a photon with energy equal to the difference between the two levels, called the transition energy. To drop back down to the lower state, the electron must emit a photon with the same energy. These transition energies differ between pairs of levels.

For the simple case of hydrogen, this means that the atom's single electron can absorb or emit light only at wavelengths corresponding to transitions between the simple energy levels shown in Figure 2-5. Neodymium is considerably more complex. Most of the 60 electrons are close enough to the nucleus that they have little effect on the outermost electrons, and the laws of quantum mechanics let electrons make transitions only between certain energy levels, but neodymium still has far too many energy levels and transitions to show in full. Figure 2-7 simplifies the energy levels in neodymium in one crystalline host to show what happens in a laser. Light from an external source excites the neodymium atom to a higher level, shown at left. Neodymium has a number of higher-level states, so light at 808 nm can raise the electron to one excited level, and light at 750 nm can excite the electron to a higher level. (Light at shorter wavelengths can excite the electron to higher levels not shown.) Electrons excited to those two higher levels spontaneously release some energy and drop to the upper laser state, where they accumulate until they are stimulated to




Figure 2-7. Laser action in neodymium. Light from an external source excites an electron in the outer shell of neodymium to a higher energy level, where it stays briefly before dropping to the upper laser level, which is long-lived, so many neodymium atoms wind up with electrons in that state. Stimulated emission then produces a cascade of light and the neodymium atoms drop down first to a lower laser level, then to the ground.

release their energy on the 1064-nm laser line, dropping to a lower energy level that is still not all the way down to the ground state.

The energy levels that electrons occupy in atoms are simple to describe, but there are also many more types of energy levels. Molecules, like atoms, have electronic energy levels, which increase in complexity with the number of electrons and atoms in the molecule. Molecules also have energy levels that depend on vibrations of atoms within them and on the rotation of the entire molecule. All these energy levels are quantized.

2.2.3 Types of Transitions

Atoms and molecules have many different types of energy-level transitions. Table 2-3 lists some important examples and their wavelength ranges.

Transitions like those of the neodymium atom are called *electronic transitions* because they involve electrons moving between two electronic energy levels. Specifically, they are transitions within the outer valence shell of electrons that are involved in chemical reactions. Hydrogen and helium only have one populated electron shell, so all their electrons are valence electrons, but heavier elements like neodymium also have electrons in inner shells, which do not take place in chemical reactions. Outer-shell transitions span a range of wavelengths from about 100 nm in

Table 2-3. Representative energy-level transitions and their wavelengths

Transition type	Wavelengths	Spectral range
Nuclear	0.0005–0.1 nm	Gamma ray
Inner-shell electronic (heavy element)	0.01–100 nm	X-ray
Inner-shell electronic (Cd ⁺²⁰)	13.2 nm	Soft X-ray
Electronic (Lyman-alpha in H)	121.6 nm	Ultraviolet
Electronic (argon-ion laser)	488 nm	Visible, green
Electronic (H, levels 2–3)	656 nm	Visible, red
Electronic (neodymium laser)	1064 nm	Near infrared
Vibrational (HF laser)	2.7 μm	Infrared
Vibrational (CO ₂ laser)	10.6 μm	Infrared
Electronic Rydberg (H, levels 18–19)	0.288 mm	Far infrared
Rotational transitions	0.1–10 mm	Far infrared to microwave
Electronic Rydberg (H, levels 109–110)	6 cm	Microwave
Hyperfine transitions (Interstellar H gas)	21 cm	Microwave

the ultraviolet to a couple of micrometers in the near infrared. Electronic transitions can occur in molecules as well as atoms, although most electronic laser transitions are in atoms.

Transitions involving electron shells inside the valence shell of a heavy element have much higher energy than those of valence electrons. Inner-shell electrons are tightly bound to a highly charged nucleus, so ejecting them from the inner shell takes a lot of energy, and dropping back into the inner shell likewise releases a lot of energy. Laser action at wavelengths down to about 10 nm has been produced by blasting atoms with intense pulses of visible or infrared light that strip many electrons from the atoms, leaving a highly charged ion. Electrons that drop into inner shells emit short-wavelength photons. For example, blasting 20 electrons from cadmium atoms produces Cd^{+20} , which produces laser light at 13.2 nm when an electron falls back into the exposed inner shell.

In theory, transitions between energy levels in atomic nuclei could release gamma rays, photons with even more energy than X-rays, but no gamma-ray laser has been made.

On the other hand, very high-lying electronic energy levels (say, levels 18 and 19 of hydrogen) are spaced so closely that transitions between them involve very little energy, with wavelengths in the far infrared, microwave, or even radio-frequency range. These are called Rydberg transitions and can only occur under rare conditions.

Molecules have two other sets of quantized energy levels with lower energy than electronic transitions. Transitions involving vibrations of atoms in a molecule typically have wavelengths of a few to tens of micrometers. Those involving rotation of the entire molecule typically correspond to wavelengths of at least 100 μm . Laser action can occur in both vibrational and rotational transitions.

Transitions in two or more types of energy levels can occur at once. For example, a molecule can undergo a vibrational and a rotational transition simultaneously, with the resulting wavelength close to that of the more energetic vibrational transition. Many infrared lasers emit families of closely spaced wavelengths on such vibrational–rotational transitions. Molecules that make vibrational transitions generally make rotational transitions at the same time, which can spread the combined transition over a range of wavelengths.

Transition energies or frequencies add together in a straightforward manner:

$$E_{1+2} = E_1 + E_2 \quad (2-7)$$

where E_{1+2} is the combined transition energy, and E_1 and E_2 are the energies of the separate transition. The same rule holds if you substitute frequency ν for energy. However, you need a different rule to get the wavelengths of combined transitions:

$$\frac{1}{\lambda_{1+2}} = \frac{1}{\lambda_1} + \frac{1}{\lambda_2} \quad (2-8)$$

2.2.4 Absorption

For an atom to make a transition to a higher energy level, it must absorb a photon of the right energy to raise it to that higher energy level. You may think it unlikely that a photon of exactly the right energy would wander by just when an atom is in the right state to absorb it. In fact, it is not improbable because of other details we have not explained.

One is that atoms tend to occupy the lowest possible energy state, especially at cold temperatures. Hydrogen atoms floating in intergalactic space are likely to be cold and in the ground state, so they are prepared to absorb light. Sure enough, spectra of light from distant galaxies show dark lines where interstellar hydrogen atoms in the ground state absorbed the light.

A second is that atoms around us rarely are truly isolated. Atoms and molecules in the air are moving constantly, bumping into each other, and transferring energy. Atoms in solids and liquids are constantly in contact with their neighbors, transferring energy back and forth. An atom in a solid does not have to wait for a photon with exactly the right energy to match a transition; if the energy is close, the atom can transfer some energy to or from a neighbor to match the transition energy. As a result, absorption is not sharply defined when matter is packed together.

2.2.5 Stimulated and Spontaneous Emission

Light emission seems simpler than absorption. An atom in a high-energy state can make a transition to a lower state whenever it “wants to” simply by spontaneously emitting a photon of the right energy. Spontaneous emission, shown in Figure 2-8A, produces the




Figure 2-8. Spontaneous and stimulated emission.

light we see from the sun, stars, television sets, candles, and light bulbs. An atom or molecule in an excited state releases energy and drops to a lower energy state. On average, it releases the excess energy after a characteristic spontaneous emission lifetime (t_{sp}), which like the half-life of a radioactive isotope measures the time it takes half the excited atoms to drop to a lower state. No outside intervention is needed.

Albert Einstein recognized another possibility. A photon with the same energy as the transition could stimulate an excited atom to release its energy as a second photon with exactly the same energy. That second photon would be a wave with precisely the same wavelength as the first and would be precisely in phase with the stimulating photon, as shown in Figure 2-8B. You can visualize the process as a wave tickling the excited atom so it oscillates at the transition energy, increasing the chance it will release a second identical photon by stimulated emission.

Now we are almost at the point of having a laser. You learned earlier that “laser” was coined as an acronym for “Light Amplification by the Stimulated Emission of Radiation.” We have seen that, like a laser beam, stimulated emission is all at the same wavelength and in phase, or coherent. But we have not made a laser beam yet. It took decades for physicists to clear the crucial hurdle needed to amplify stimulated emission.

2.2.6 Populations and Population Inversions

The problem with stimulated emission is that it does not work well when matter is in *thermodynamic equilibrium*, a state in which




Figure 2-9. Relative populations of energy levels as a function of energy above the ground state at thermal equilibrium.

atoms and molecules tend to be in the lowest possible energy level. Physicists traditionally considered that to be the normal state of matter. Not all atoms are in the ground state because they always have some thermal energy above absolute zero. But at equilibrium the number of atoms in an energy level always decreases with the energy of the state, as shown in Figure 2-9. In fact, the decline can be steep.

Why does this make stimulated emission difficult? Because the best way to stimulate emission from an excited atom is with spontaneous emission from an identical excited atom. That would seem easy if you had a lot of atoms in the excited state, but in equilibrium more atoms would be in the low-energy state than in the excited state. In that case, spontaneous-emission photons would be more likely to encounter atoms in the lower state that would absorb them than atoms in the upper state that they could stimulate to emit another photon. So physicists were able to show that stimulated emission was possible in the 1920s, but they could not stimulate enough emission to be of any practical use.

For stimulated emission to dominate, the majority of atoms must be in an excited state, so spontaneously emitted photons are more likely to stimulate emission than to be absorbed by atoms in the ground state. This is called a population inversion, because the population is inverted from the normal situation in which more atoms are in lower levels than in higher levels. Stimulated emission can continue as long as the population inversion exists, but stops once the ground-state population becomes larger and photons are more likely to be absorbed than to stimulate emission.

Nobody knew how to produce a population inversion in the 1920s and 1930s. It seemed so counterintuitive that they called such a condition a “negative temperature” because in theory the only way to produce a population inversion was to have a temperature below absolute zero. Only after World War II did physicists realize they could produce a population inversion for more than a fleeting instant. You will learn more about that in Chapter 3.

2.3 INTERACTIONS OF LIGHT AND MATTER

So far, we have described light and its interaction with individual atoms. However, what is important in most laser applications is the interaction of light with bulk materials. That is the science of *optics*, and it depends on both the wavelength of light and the properties of the materials.

2.3.1 Optics and Materials

Light interacts with individual atoms and molecules in the air, but atoms are packed much closer together in liquids and solids. Visible wavelengths of 400 to 700 nm are more than a thousand times larger than an average atom (around 0.1 nm), so they interact with large numbers of atoms spread through the bulk material. Those interactions depend on both the wavelength and the composition of the material.

We usually sort objects into three classes according to how they interact with light:

1. Transparent objects (e.g., glass) transmit light
2. Opaque objects (e.g., dirt or rocks) absorb light
3. Reflective objects (e.g., mirrors) reflect light

However, nothing is perfectly transparent, opaque, or reflective. Thin slices of objects we consider opaque actually transmit some light. If you are reading a paper copy of this book, hold its pages together, and they are opaque, but hold a single page up to a bright light and some light will pass through. That is because absorption builds up quickly. If each page transmits 10% of the light striking it, 10 pages stacked together allow just 10^{-10} or one tenth of a billionth of the original light to get through. Likewise, geologists

can see through rocks that look opaque by cutting them into slices so thin they transmit some light.

Everything reflects some light, whether we consider it transparent, opaque, or reflective. We see the light that objects reflect, and our eyes interpret the amount of reflection as dark or light. Sometimes our eyes can fool us. A full moon reflects only about 6% of the incident sunlight, but it looks bright when seen against a black night sky. White paint looks white because it reflects light uniformly across the visible spectrum. Green leaves look green because they reflect much more light in the green part of the spectrum than at other wavelengths. (Section 2.3.5 explains why mirrors reflect light differently than white paint.)

Transparent materials are important for optics, because light can pass through them. Nothing is perfectly transparent except for a perfect vacuum, which is, of course, really nothing. As Section 2.3.6 explains, light transmission depends on the wavelength. Glass is transparent through the visible spectrum, but it absorbs light in much of the infrared and ultraviolet. Ruby crystals and rose-colored glasses absorb blue, green, and yellow light but transmit red, so light emerging from them looks red to the eye. Finally, any transparent material also reflects some light at its surface; that is why you see reflections from indoors when you look out through a window at night.

Let us look a bit more at the details, starting with transparent materials.

2.3.2 Refractive Index

The most important property of a transparent material is its *refractive index*, n , which tells how fast light travels in the material relative to that in a vacuum. The velocity of light in a vacuum, denoted as c , is a universal constant, precisely defined as 299,792.458 kilometers per second, but often rounded to 300,000 km/sec or 186,000 miles/sec. Light slows down in matter, even something as tenuous as air. The refractive index measures how much longer light takes to pass through a solid by dividing the speed of light in air by the speed of light in the material:

$$n_{\text{material}} = \frac{c_{\text{vacuum}}}{c_{\text{material}}} \quad (2-9)$$

Because the speed of light is faster in a vacuum than in a material, the refractive index normally is greater than one. (Some very

Table 2-4. Refractive indexes of common materials for wavelengths near 500 nm, except where noted

Material	Index
Air (1 atmosphere)	1.000278 (usually approximated as 1.0003)
Water	1.33
Magnesium fluoride	1.39
Fused silica	1.46
Zinc crown glass	1.53
Crystal quartz	1.55
Optical glass	1.51–1.81*
Heavy flint glass	1.66
Sapphire	1.77
Zircon	2.1
Diamond	2.43
Silicon	3.49 @ 1.4 μm
Gallium arsenide	3.5 @ 1 μm

*Denotes depends on composition. Standard optical glasses have refractive indexes in this range.

interesting exceptions exist on the cutting edge of optical science, but they involve special materials with internal structures smaller than the wavelength of light, so you do not need to worry about them now.) Air is so tenuous that it has a very low refractive index of 1.000278 at atmospheric pressure, which is so small that it can be ignored for most optical systems.

It is important to note that the refractive index varies with wavelength. This must be considered in optical devices, particularly those built to transmit the whole visible spectrum, like binoculars. Table 2-4 lists the refractive indexes of some common materials near 500 nm, in the middle of the visible spectrum.

The frequency of light waves in solids is the same as in a vacuum, but the wavelength depends on the refractive index. Just as the refractive index slows the speed of light in a solid, it also decreases the wavelength, which equals the wavelength in vacuum divided by the material's refractive index n_{material} :

$$\lambda_{\text{material}} = \frac{\lambda_{\text{vacuum}}}{n_{\text{material}}} \quad (2-10)$$

2.3.3 Refraction

Light bends as it passes between two materials with different refractive index, such as from air into glass. This effect, shown in

Figure 2-11, is called *refraction* and depends on the refractive index. The closely spaced parallel lines in the figure represent the peaks of successive waves, one wavelength apart in the direction in which the light travels. In the low-index material at the top, the peaks are far apart, but the wavelength shrinks when the light slows down in the higher-index material at the bottom. The peaks stay in phase, so the change in refractive index bends the light in the direction shown. If the light wave passes into the same low-index material on the other side of high-index material, the light waves become longer, and the light bends back. You can see how if you turn the page upside down.

Refraction only occurs at the border between two transparent materials, not within the materials themselves. The amount of refraction is measured by comparing the direction of the light with the *normal angle*, which is perpendicular to the surface where refraction takes place, as shown in Figure 2-10. The angle is simple to calculate using Snell's law, which states that if light traveling in




Figure 2-10. Refraction of light waves as they pass from a low-index medium (such as air) with refractive index n_1 to higher-index medium with index n_2 .

a medium with refractive index n_1 strikes the surface of a material with index n_2 at an angle of θ_1 to the normal, the direction of light in the second material is given by θ_2 :

$$n_1 \sin \theta_1 = n_2 \sin \theta_2 \quad (2-11)$$

You can also rewrite this to directly calculate the angle of refraction θ_2 :

$$\theta_2 = \arcsin \left(\frac{n_1 \sin \theta_1}{n_2} \right) \quad (2-12)$$

where the arcsine is the inverse of the sine.

Suppose, for example, that light in water ($n_1 = 1.33$) strikes the surface of crystal quartz ($n_2 = 1.55$) at 30 degrees to the normal. The formula tells us that $\theta_2 = 25.4$ degrees, meaning that the light is bent closer to the normal. If the light were going in the other direction, it would be bent further from the normal. If the difference in refractive index were larger, as when going from air into quartz, the change in angle would be larger.

Something peculiar happens if light goes from a high-index material into a low-index material at a steep angle. Suppose, for example, that light in a quartz crystal hits the boundary with air ($n = 1.0003$) at a 50-degree angle to the normal. Plug the numbers into the equation above, and you find that θ_2 equals the arcsine (1.19). However, the sine of an angle cannot be greater than one, so your calculator will flag that as an error. If you try an experiment, you will find that the light does not escape into the air; it is all reflected back into the quartz, an effect called *total internal reflection* that makes diamonds sparkle and is the basis of light guiding in optical fibers.

2.3.4 Transparent and Translucent Materials

You may have been confused at some point by the distinction between transparent and translucent materials. Transparent materials are clear, like a window. Translucent materials transmit light, but are cloudy, like milk, wax paper, or ground glass. The difference is that transparent materials let light pass straight through, whereas translucent materials scatter the light rays, blurring them, so you cannot see clearly. This is similar to the difference

between white and mirror-like reflective surfaces; the white surfaces scatter light diffusely and the shiny surfaces are smooth and reflect light directly, so you see images on a mirror but not on a white wall.

2.3.5 Reflection

Reflection redirects light, as we see every day in mirrors. The law of reflection is simplicity itself:

$$\text{Angle of Incidence} = \text{Angle of Reflection} \quad (2-13)$$

This means that if light strikes a mirror at a 50-degree angle to the normal, the mirror reflects it at the same angle.

This law applies to mirror-like or *specular* reflection from a surface that is smooth on a scale smaller than the wavelength of light. Such smoothness is possible for polished glass and metal surfaces. A surface that may look smooth to the eye but is rough on the scale of a wavelength reflects light *diffusely*. Each point reflects light at the angle of incidence, but the surface is uneven enough that reflection scatters light in all different directions. You can compare diffuse reflection to balls bouncing off a pile of rocks and specular reflection to balls bouncing off a flat floor. The paper of this book reflects light diffusely.

Household mirrors reflect light from a metal film on their back side, so light is refracted going back and forth through the glass to the reflective surface. High-performance optical systems require *front-surface mirrors* which reflect from their front side. Front-surface mirrors require more care than household mirrors, because their reflective surfaces are directly exposed to the environment. However, they reflect more light and are much more accurate optically. Front-surface mirrors may be solid metal or glass coated with a reflective surface layer.

The surface of a transparent material also reflects some light due to an effect called Fresnel reflection. People learn to recognize this reflection as a subtle cue to avoid walking through glass doors, although a bit of dirt also helps. (Dogs, birds, and distracted people do not notice the faint reflection and may crash into glass doors.) The amount of surface reflection can be reduced by coating glass with a lower-index material.

Table 2-5. Refractive index of silicon dioxide glass as a function of wavelength

Wavelength (nm)	Index
404.656	1.46961
508.582	1.46187
589.262	1.45847
706.519	1.45515
852.111	1.45248
1013.98	1.45025
1395.06	1.44584

2.3.6 Wavelength-Dependent Effects

As mentioned earlier, optical properties of materials vary with the wavelength of light. For example, the refractive index of silicon dioxide glass decreases slightly as wavelength increases, as listed in Table 2-5. Looking at the numbers, the overall variation seems so small that it would not be noticeable. However, this small variation causes glass prisms to refract different colors at angles that differ enough to display a full spectrum of colors. Rainbows arise from refractive-index variations in tiny water droplets in the air.

Light transmission through transparent materials also depends on its wavelength. This effect is not obvious in window glass because the materials are chosen to be reasonably transparent across the visible spectrum. However, it is more obvious in colored glass or crystals, which get their color from uneven absorption. Variations in light transmission are also evident in other parts of the spectrum.

Reflection also depends on wavelength. We pick paints and inks because they reflect specific colors. These are things you have to think about when you work with lasers. When supermarket scanners were first introduced, a Boston-area dairy printed bar codes on its white milk cartons in red ink. The symbol was clearly visible to the human eye, but the scanner could not see the symbol because the red ink reflected almost as much red light as the white paper.

2.3.7 Optical Attenuation and Losses

Some optical materials are remarkably clear. Optical fibers used for telecommunications are so transparent at 1550 nm in the infrared that 10% of the light that enters them remains after passing through 50 kilometers of glass. Yet no optical material is totally transparent.

Some losses come from light absorption by atoms in the material. How much light the material absorbs depends on the wavelength of light and the material composition. The clearest optical fibers are pure silica (SiO_2); impurities such as iron or hydrogen can increase their absorption dramatically.

Other losses come from scattering by atoms in the material. Rayleigh scattering, the process that makes the sky blue, makes light waves bounce off objects much smaller than the waves. The shorter the wavelength, the bigger the atoms look to the light, and stronger the scattering—which is why you see red light as the sun sinks down the horizon, while the blue light is scattered to light the sky.

Absorption and scattering add together to cause loss or attenuation, the reduction in light power per unit distance as it travels through a material. If the material is uniform, the same fraction of the input light is lost each unit distance. Loss is measured by the attenuation per unit distance (α), and total loss is calculated with an exponential equation. For example, if 10% of the light (0.1) is lost in 1 cm, the amount passing through 10 cm is the amount passing through one cm (0.9) raised to the tenth power, $(0.9)^{10}$ or 35%. The general formula is an exponential equation:

$$I = I_0 \exp^{-\alpha D} \quad (2-14)$$

where I is the intensity of light that has passed through a distance D of the material (90% in this case), I_0 is the initial intensity, and α is a measure of the fraction transmitted through a unit distance.

2.3.8 Diffraction

Diffraction is an effect that arises from the wave nature of light. Look back at how light passes through the slits in Figure 2-3, and you see how light waves radiate from each slit. The light does not travel in a straight line; it spreads out in waves from the edge. You can also see diffraction if one edge of the slit is removed, leaving the other edge to cast a shadow; the diffracted waves spread toward the shaded area. You can see a similar effect on the surface of water if a barrier blocks waves.

As in the two-slit experiment, diffracted waves interfere with each other as they overlap. The most obvious diffraction effect comes from white light diffracted from an array of many parallel lines. The angles at which the waves diffract depend on wavelength

and spacing of the lines, and interference among the diffracted waves concentrates different wavelengths at different angles, spreading out a rainbow of light. The easiest way to see the effect is by reflecting light from a CD, DVD, or Blu-Ray disk; the spiral patterns of data recorded on the disks act like a diffraction grating.

Diffraction effects are important in lasers because they put a lower limit on the spreading of a laser beam with distance, an effect called beam divergence, described in Section 4.3.1.

2.4 BASIC OPTICS AND SIMPLE LENSES

The next step is to move from basics such as refraction and reflection to simple optics such as lenses, which focus light and are important in laser applications. We cannot go through them all, but we can give some simple examples.

A good starting point is the operation of a lens, which can be visualized by following the paths of a number of imaginary rays that initially are parallel. You can think of them as straight lines representing the path of light waves. Refraction bends the rays at the surface of the lens, according to Snell's law (Equation 2-12). Rounding the surface of the lens can focus the light to a point or an image, collimate it into a beam, or spread it out.

Figure 2-11 shows some simple lenses and how they affect parallel light rays entering them from the left. Lenses that are thicker at the center than at the sides are called *positive lenses* and bend the light rays so they come together at a *focal point*. If the lens is thicker at the edges—a *negative lens*—it spreads out or diverges the light rays as shown. Properties of the lenses depend on their surface curvature, refractive index, and size.

Curved mirrors work similarly to focus or disperse light.

2.4.1 Positive Lenses

The most important parameter of any lens is its *focal length*. For a positive lens, this is the distance from the center of lens to the point at which it focuses parallel light rays to a point, as shown in Figure 2-12. Most simple lenses have spherical surfaces, with their shape measured by a "*radius of curvature*," which is the distance from the lens surface to the point that would be the center of a glass sphere with the same shape. One or both surfaces are curved.

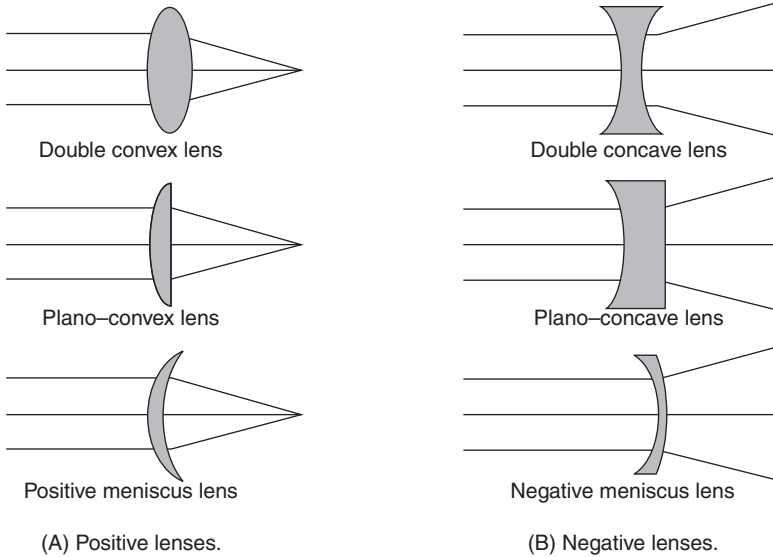


Figure 2-11. Six basic types of simple lenses.

The focusing power depends on the shape of the lens. You can see this from the general formula for the focal length f of a lens with refractive index n and two curved surface, with radii of curvature of R_1 and R_2 :

$$f = \frac{1}{(n - 1) \left[\frac{1}{R_1} + \frac{1}{R_2} \right]} \quad (2-15)$$




Figure 2-12. Important parameters of a positive lens.




Figure 2-13. A positive lens forms a real image of an object.

If one surface of the lens is flat, its radius of curvature is infinite, so the $1/R$ term equals zero. If a surface is curved inward (concave), it has a negative radius of curvature.

These formulas are simplified to consider only light rays that are parallel, as if they came from the sun or some other very distant object. Light rays diverge outward from nearby objects, so if the lens is placed at a certain point as shown in Figure 2-13, it focuses light rays from the object to a point beyond the focal point where the light forms an image of the object. This image is called a *real image* because it can be focused onto a screen at a distance D_1 that depends on the focal length f of the lens and the distance D_o of the object from the lens.

$$\frac{1}{D_1} = \frac{1}{f} - \frac{1}{D_o} \quad (2-16)$$

The image size depends on the distance. The ratio of image height h_i to object height h_o is called the magnification ratio m . It also depends on the distance between the lens and the object:

$$m = h_i/h_o = D_i/D_o \quad (2-17)$$

As the formula indicates, the magnification is greater than one if the object is near the lens and the image is far away, meaning that the image is larger than the object. If the image is nearer the lens than the object, the magnification is smaller than one, and the image is smaller than the object. In Figure 2-13, the object and the image are at equal distances from the lens, and the magnification is one.




Figure 2-14. Important parameters of a negative lens.

2.4.2 Negative Lenses

The same optical laws apply for negative lenses as for positive lenses, but they yield different results. Concave surfaces are assigned negative radii of curvature, which can lead to negative focal lengths, which means that negative lenses do not focus parallel light rays so they come together. Pass parallel light rays through a negative lens, and they seem to spread from a point behind the lens, as shown in Figure 2-14. The distance from the midplane of the lens to this point is the focal length; it has a negative value because it is on the opposite side of the lens from the focal point of a positive lens. (Optical designers use sign conventions to keep these things straight.)

2.4.3 Mirrors

Curved mirrors, like lenses, can focus parallel light rays or cause them to spread apart. Mirror surfaces, like lens surfaces, can be either concave or convex. A *concave mirror* seems hollow, like the inside of a bowl. Like a positive lens, a concave mirror with a spherical surface can form a real image or focus light rays to a point. A *convex mirror* has a rounded surface, like the outside of a ball; reflection spreads out light rays like a negative lens, and it cannot form a real image.

Like lenses, spherical mirrors have focal lengths, which indicate the distance from the mirror at which they focus parallel rays. The focal length f is half the mirror's radius of curvature R :

$$f = R/2 \quad (2-18)$$

Following the same sign convention as for lenses, we assign a positive radius of curvature to a concave surface and a negative radius to a convex surface. This gives a concave mirror a positive focal length and a convex mirror a negative focal length.

2.4.4 Optical Complexities

Optics is a complex field that goes far beyond lasers and the simple concepts you have learned so far, but this book is about lasers, so it only covers what you need to know to get started with lasers. You will find more laser-related optics in Chapters 5 and 6. To learn more about optics in general, check the Further Readings in Appendix C.

2.5 FIBER OPTICS

Fiber optics are long thin transparent fibers of glass or plastic that guide light along their length. Individual fibers are flexible, so they can transmit light to and from otherwise inaccessible locations, including inside the human body for medical diagnosis and treatment. Fiber optics are made for a number of purposes, including for use as lasers (see Chapter 9) and for laser beam delivery.

This section will first describe the mechanisms behind light guiding, then describe the types of optical fibers that are most important for use in and with lasers. Fiber-optic technology continues to grow rapidly, and this book cannot cover the wide variety of optical devices based on fiber optics that are used with lasers, such as special fibers designed for sensing.

2.5.1 Light Guiding in Optical Fibers

Light guiding in optical fibers is based on total internal reflection, a concept introduced in Section 2.3.3. The basic idea is that light in one material that hits the boundary with another material with lower refractive index at a glancing angle will be totally reflected back into the higher-index material.

Figure 2-15 shows how this works inside an optical fiber. Light enters a high-index material that forms the central core of the fiber, then hits the boundary with the lower-index cladding layer at a glancing angle, where it is reflected entirely back into the core. The




Figure 2-15. Internal structure of an optical fiber, showing how total internal reflection guides light through the central light-guiding core.

light ray can bounce off the cladding a number of times when the fiber is bent.

You may wonder why bother cladding the fiber? Would not it just be simpler to use a bare glass fiber in air, which has a refractive index of 1.0003 compared to 1.5 for the glass core? It might seem simpler, but it does not work well. If the total internal reflection surface was in air, light could leak out wherever that surface touched something, including fingerprint oils. For that reason, most optical fibers have at least two layers, an inner light-guiding core with a high refractive index and an outer cladding with a lower refractive index, as shown in Figure 2-15. Fibers typically also have an outer plastic coating to protect the glass from damage during handling.

Fiber designs differ considerably, depending on the intended use. Fibers used in long-distance telecommunications typically have a 10- μm glass core surrounded by a glass cladding with refractive index only about 1% lower than that of the core. Beam-delivery fibers used in industrial machining typically have a much larger core to handle the high power required, with a larger refractive-index difference. The fibers used to make fiber lasers and fiber amplifiers have cores doped with a light-emitting element such as erbium or ytterbium, which are described in Chapter 9.

The design of the light-guiding core is critical for fiber applications. Technically, the core of the fiber functions as a *waveguide* that

confines light and directs it along the length of the fiber. Important types include:

- *Single-mode fibers*, with a 10- μm core that transmits only one mode, used for fiber-optic communications, and a 125- μm cladding covered by a plastic coating.
- *Multimode fibers*, with glass cores 50 μm or larger surrounded by a cladding layer, usually glass. One application is short-distance data links. Another is beam delivery for industrial lasers, including fiber lasers. Typical beam-delivery fibers have cores 50, 100, or 200 μm in diameter. Special types have been developed for very high-power lasers.
- *Dual-core fibers* with an active light-emitting core used for fiber lasers or amplifiers as described in Chapter 9.
- *Flexible fiber-optic bundles* are loose fibers bundled together, generally for imaging or illumination applications. A bundle with fibers arranged in matching configurations on the two ends can transmit images and is called coherent; such bundles are used for medical endoscopies and colonoscopies. Fiber bundles used to illuminate objects typically are arranged randomly.
- *Fused fiber-optic bundles* are fibers arranged in bundles which are heated to the melting point and drawn to fuse them together and shrink the bundle diameter. These fiber bundles are rigid and may be used for imaging applications.

2.5.2 Passive Fibers and Beam Delivery

Most fibers described in this section are considered passive because they transmit light without amplifying or otherwise modifying it. Important applications include:

- *High-power laser beam delivery* directs high-power laser beams to objects being machined with lasers. One end of a large-core fiber collects light from a laser; the other end delivers beam to the object being machined.
- *Laser beam delivery for minimally invasive surgery* involves threading optical fibers through small holes and into the body to avoid large incisions that take a long time to heal. One example is laser lithotripsy, in which an optical fiber is threaded through the urethra, bladder, and ureter into the kidney, where it delivers laser pulses to break up kidney stones into pieces small enough to pass naturally through the urethra and ureter.

- *Optical communications* through optical fibers, which carry voice, video, and data around the globe.

2.5.3 Active Fibers, Lasers, and Amplifiers

Active fibers contain materials that amplify, generate, or otherwise modify laser light. The primary uses of active fibers are in the fiber lasers and amplifiers described in Chapter 9.

2.6 WHAT HAVE WE LEARNED?

- Light sometimes acts like waves and other times acts like massless particles called photons.
- Light is electromagnetic radiation that can be classified by its photon energy, wavelength, or frequency.
- Wavelength is the distance between wave peaks; frequency is the number of wave peaks passing a point per second.
- Frequency equals the speed of light divided by wavelength.
- The speed of light in vacuum is roughly 300,000 kilometers per second.
- Photon energy equals Planck's constant (h) times frequency.
- The electromagnetic spectrum ranges from radio waves to gamma rays.
- Visible wavelengths are 400 to 700 nm.
- Waves interfere constructively when they are in phase and their amplitudes. Waves interfere destructively when they are out of phase by 180 degrees, canceling each other.
- Quantum physics recognizes that energy comes in discrete chunks or quanta.
- Quantum numbers specify electronic energy levels.
- An electron loses energy if it drops to a lower energy level. It must absorb energy to make a transition to a higher level.
- Electronic transitions occur in atoms and molecules.
- Vibrational and rotational transitions occur in molecules and involve less energy than electronic transitions.
- Excited states can emit energy spontaneously or when stimulated by a photon with the same energy.
- A population inversion occurs on a transition when more atoms or molecules are in the upper state than in the lower state. It is needed for a laser to operate.

- Laser light is stimulated emission.
- Stimulated emission amplifies light.
- A laser oscillates when it amplifies light that bounces back and forth within a resonant cavity.
- Objects may transmit, absorb, or reflect light.
- The refractive index of a material equals the speed of light in vacuum divided by the speed of light in the material.
- Refraction bends light rays as they pass between transparent media. Snell's law shows how this depends on the refractive index and the angle of incidence.
- The angle of incidence equals the angle of reflection.
- The refractive index and light absorption of a material vary with the wavelength of light.
- Light transmission depends on transparency and thickness of a material; the intensity drops exponentially with distance traveled through the material.
- Diffraction is the spreading of waves from an edge.
- A positive lens focuses parallel light rays to a spot. A negative lens causes parallel rays to diverge.
- The focal length of a positive lens is the distance from the lens to a point at which it focuses parallel light rays.
- A positive lens can project a real image.
- A virtual image cannot be projected but can be seen by the eye.
- Fiber optics guide light through a central core by total internal reflection.

WHAT'S NEXT?

In Chapter 3, we will use what we have learned about optics and physics to describe how lasers work in general. Later, we will learn about specific types of lasers, which differ greatly in detail.

QUIZ FOR CHAPTER 2

1. A carbon-dioxide laser has a nominal wavelength of $10.6 \mu\text{m}$. What is its frequency?
 - a. 300,000 hertz
 - b. 2.8×10^{13} hertz
 - c. 1.06 gigahertz

- d. 2.8×10^{10} hertz
 - e. None of the above
2. What is the photon energy for an infrared wave with frequency of 10^{14} hertz?
 - a. $10.6 \mu\text{m}$
 - b. 6.63×10^{-34} joule
 - c. 6.63×10^{-20} joule
 - d. 10.6×10^{20} joules
 - e. About one joule
 3. What is the wavelength of the infrared wave in Problem 2?
 - a. 300 nm
 - b. $3 \mu\text{m}$
 - c. $10.6 \mu\text{m}$
 - d. $30 \mu\text{m}$
 - e. 10.6 nm
 4. How do two light waves of the same wavelength and amplitude interfere if they are 180 degrees out of phase?
 - a. Destructively
 - b. Constructively
 - c. Partially, producing a wave with amplitude $\sqrt{2}$ times the two input waves
 - d. Not at all
 5. Calculate the wavelength of the transition in the hydrogen atom from the $n = 2$ energy level (the second orbit) to the $n = 3$ level. This is the first line in the Balmer series of spectral lines.
 - a. 121.6 nm
 - b. 91.2 nm
 - c. 656 nm
 - d. 632.8 nm
 - e. 900 nm
 6. An electron in the ground state first absorbs one photon at 500 nm, then a second at 1000 nm, to reach an excited state. What is the wavelength of the photon it would emit if it dropped all the way back to the ground state?
 - a. 200 nm
 - b. 225 nm
 - c. 300 nm
 - d. 333 nm
 - e. 1500 nm

7. Light in a medium with refractive index of 1.2 strikes a medium with refractive index 2.0 at an angle of 30 degrees to the normal. What is the angle of refraction (measured from the normal)?
 - a. 42 degrees
 - b. 20.1 degrees
 - c. 18.0 degrees
 - d. 17.5 degrees
 - e. 15.6 degrees
8. A 0.5-cm thickness of a material absorbs half the light that enters it and has negligible scattering. What fraction of incident light can pass through a 2-centimeter thickness?
 - a. 0.5
 - b. 0.25
 - c. 0.018
 - d. 0.01
 - e. None; the material is opaque
9. A positive lens with focal length of 10 centimeters forms a real image of an object 20 centimeters away from the lens. How far is the real image from the lens?
 - a. 5 cm
 - b. 10 cm
 - c. 15 cm
 - d. 20 cm
 - e. 25 cm
10. What is the ratio of image size to object size for the case in Problem 9?
 - a. 2
 - b. 1.5
 - c. 1
 - d. 0.667
 - e. 0.5

CHAPTER 3

HOW LASERS WORK

ABOUT THIS CHAPTER

Chapter 2 explained the basic principles of physics and optics that are needed to understand how lasers work. This chapter describes how lasers work. Lasers transform energy rather than producing it. Most lasers transform electric current into a beam of light in one or more steps. In the simplest case, a laser uses the input energy to excite atoms or molecules, extracts that energy as a cascade of photons, and transforms that light into a well-controlled beam of light. The rest of this chapter will teach you the general principles of laser operation. Chapter 4 will explain important laser characteristics.

3.1 BUILDING A LASER

The best way to learn how a laser works is to build one. We cannot actually build a laser in this book, but we can follow the path that early laser developers took to make lasers work. We have a big advantage in knowing where we are going. Hindsight is always 20–20, so we do not have to pursue all the dead ends that wasted the original developers' time and energy, but describing the key steps will help you understand how the pieces go together. Each step is a separate section of this chapter.

The first step toward our conceptual laser is figuring how to produce the population inversion needed for stimulated emission, covered in Section 3.2.

Once we have the population inversion needed for stimulated emission to dominate, Section 3.3 tells how a resonant cavity helps amplify light to a point where the laser oscillates.

Then Section 3.4 explains how oscillation creates resonances and generates a laser beam.

Section 3.5 tells how the resonator and laser medium combine to select a laser wavelength.

Finally, Section 3.6 goes back and looks more closely at exciting the laser medium to produce a population inversion.

To keep the discussion simple, we will talk about atomic states rather than molecular states, but remember that the same principles apply to molecules.

3.2 PRODUCING A POPULATION INVERSION

The starting point for making a laser is producing the population inversion needed for stimulated emission. To make it simple, consider a population inversion between two energy levels. The *upper laser level* is the high-energy state from which an atom drops when it emits stimulated emission. The *lower laser level* is the lower-energy state into which the atom drops after emitting light. Real atoms have many more energy levels.

Under normal conditions, most atoms tend to be in their lowest possible energy level, with only a few in higher energy states, as you saw earlier in Figure 2-9. This condition is called *thermodynamic equilibrium*, and you can think of it as what happens after things you stirred up settle down to their lowest possible energy levels. At equilibrium, more atoms always are in the lower level of our two-level system than in the upper level. In that case, a photon spontaneously emitted by an atom in the upper state is more likely to encounter an atom in the ground state that will absorb it than it is to encounter an atom in the higher level that it can stimulate to emit another photon. So no laser action will occur.

For laser action to occur, more atoms must be in the upper laser level than in the lower one, so the first spontaneously emitted photon is likely to encounter an upper-state atom that it can stimulate to emit an identical photon. As long as more atoms are in the upper state than in the lower one, this can produce a cascade of stimulated emission that becomes a laser beam. An excess of atoms in the higher energy level is called an *inverted population* because it

is the inverse of the equilibrium condition, with more atoms in the lower level.

Real atoms are far more complex than our simple two-level system, so they have many more energy levels and possible transitions. It is possible for population inversions to occur on two or more different transitions simultaneously, but generally only one can oscillate at a time.

3.2.1 Excitation Mechanisms

It takes energy to produce a population inversion, but heat alone will not work. Heating increases the average energy of atoms, and thus their temperature. But heating is not selective, so it always leaves more atoms in lower energy levels and cannot produce a population inversion unless something else is done, such as flowing hot gas at high speed through nozzles into a low-pressure chamber.

Producing a population inversion requires selectively exciting atoms from lower energy levels into higher energy levels. To visualize what is needed, think of energy levels as narrow flat steps on a staircase and the atoms as little marbles. Put a marble on a step, and soon it will roll to the edge and drop to a lower step, just as atoms drop to lower energy states by spontaneous emission. The only way to keep a number of marbles on the upper steps is to keep putting them there. Producing a population inversion likewise requires exciting atoms to the upper laser level faster than they can drop back down to a lower state.

Selective excitation is tricky, because an atom must have a certain energy to move to the right energy level. That energy could come from light of the right wavelength, or from electrons in an electric current passing through a gas or semiconductor. We will examine these ideas in Section 3.6, but for now let us consider excitation in general.

3.2.2 Masers and Two-Level Systems

The first major step on the road to the laser was the microwave-emitting maser. Charles Townes produced a population inversion by passing ammonia molecules through a magnetic field to separate molecules in a high-energy level from those in a lower energy state. He then directed the excited molecules into a chamber where spontaneous emission from some excited ammonia molecules

could stimulate other excited molecules to emit microwaves on a transition at 24 gigahertz. The walls of the chamber reflected the microwaves back and forth within the cavity, and some microwaves escaped from the cavity.

Physically separating excited molecules showed that a population inversion could generate stimulated emission and oscillation in a suitable chamber. But it was a cumbersome, brute-force approach. More practical masers, and the later development of the laser, required ways to selectively excite atoms to particular states and keep them there. To see how that works, we need to look at the lifetimes of excited states.

3.2.3 Metastable States and Lifetimes

Normally, atoms in excited states release their extra energy by spontaneous emission within nanoseconds (billionths of a second) after reaching the excited state. That is so short a time that most excited atoms do not remain in the high-energy state long enough for anything to stimulate emission. What is needed is a longer-lived excited state.

Such states are called *metastable* because they are unusually stable on an atomic timescale, although they may only last for a millisecond (a thousandth of a second) or a microsecond (a millionth of a second). They exist because the quirks of quantum mechanics affect how fast atoms can drop to a lower energy level. Atoms must change their quantum states to shift between energy levels, and some changes are less probable than others. Fast transitions are ones that are easy for atoms; metastable transitions are less likely.

To go back to our mechanical analogy, imagine that short-lived states are steps with surfaces that curve down slightly, so the marbles quickly roll off them. Metastable steps with surfaces that curve up a bit, let the marbles stay in place longer.

A metastable state makes a good upper laser level because atoms stay in that state long enough for a large population to accumulate in that state. Spontaneous emission always lets some atoms drop to the lower level, but the longer the spontaneous-emission lifetime, the slower the leakage, and the easier it is to maintain the population inversion needed for stimulated emission.

3.2.4 Three- and Four-Level Lasers

It is hard to excite atoms directly into a metastable state, but population inversions can be produced by exciting atoms indirectly

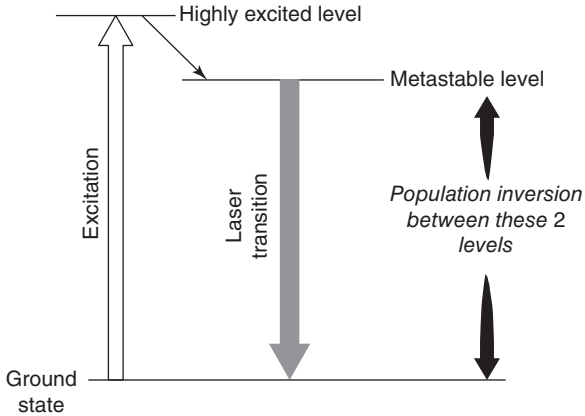


Figure 3-1. Energy levels in a three-level laser.

to metastable states. A simple approach involving three distinct energy levels is shown in Figure 3-1. Atoms start in the ground state, which is also the lower laser level. Energy from an external source excites ground-state atoms to a short-lived energy level slightly above the metastable level. Then, the atoms quickly drop down to the metastable upper laser level, where they typically remain a thousand times or longer than in the higher level. The atoms collect in the metastable level, creating a population inversion between it and the lower laser level. A few photons spontaneously emitted by atoms in the metastable state then stimulate a cascade of laser emission from atoms remaining in the metastable state.

Maiman's ruby laser is an important example of this three-level scheme. However, it is not ideal because the ground state is also the lower laser level, so most of the atoms must be excited to the upper laser level to produce a population inversion. Maiman did that successfully using a flashlamp, but ruby and other three-level lasers are generally limited to pulsed operation and cannot emit steady beams.

The four-level laser shown in Figure 3-2 allows continuous operation and has other advantages. As in the three-level laser, the atom is excited from the ground state to a short-lived, highly excited level, then drops quickly to a metastable upper laser level. Stimulated emission on the laser transition drops atoms to the lower laser level, which in a four-level laser is above the ground level, as shown in Figure 3-2. The atoms then drop from the lower laser




Figure 3-2. Energy levels in a four-level laser.

level to the ground state, releasing energy by spontaneous emission or other processes.

The key difference is that the lower level of a four-level laser is not the ground state. This is important because it means the population of the upper laser level only has to exceed that of the lower laser level, not the ground state. That avoids the need to excite most atoms out of the ground state, opening the door to continuous output.

Four levels in a laser are not enough to ensure continuous laser emission because atoms can accumulate in the lower laser level, ultimately stopping laser emission when the population of the lower laser level exceeds that in the upper laser level. To prevent such a bottleneck, atoms should drop out of the lower laser level faster than from the upper laser level. Rapid depopulation of the lower laser level allows the laser to emit continuously (called *continuous-wave* emission). Accumulation of too many atoms in the lower laser level limits the laser to pulsed operation, so the lower level has time to empty before the next pulse.

If you look carefully at the real lasers described in later chapters, you will see that their actual energy-level structures can be quite complex. Excitation is not always to a single high level; it may be to a group of levels, which decay to the same upper laser level.

That improves efficiency by exciting atoms over a wider range of energies than would be possible if they could be raised to only a single state. The depopulation of the lower laser level can also be complex and may involve a series of steps. The upper and lower laser levels can be much farther above the ground state than indicated in Figure 3-1 or 3-2, which makes it easier to control their populations, but limits overall operating efficiency.

Laser power-conversion efficiency can also be improved by adding other types of atoms to the laser mixture. For example, the helium in a helium–neon laser captures energy from electrons passing through the gas, and those excited helium atoms transfer their energy to neon atoms, capturing more energy from the electric discharge and generating more laser light than possible with pure neon. That excitation process produces a population inversion in the neon gas and emission on transitions of neon. In other gas lasers, a gas may be added to depopulate lower laser levels.

3.2.5 Natural Masers and Lasers

A major reason that decades passed between Einstein's prediction of stimulated emission and the first working laser was that physicists thought population inversions would be very hard to produce. They had no idea that stimulated emission might be visible in nature, but in the 1960s radio astronomers discovered that stimulation emission occurred naturally in outer space.

Natural masers occur in gas clouds near hot stars. Starlight excites simple molecules such as water to high-energy levels, and the molecules drop down the energy-level ladder to a metastable state. In a dense gas, excited molecules would quickly lose their extra energy by collisions, but the gas in space is so tenuous that collisions are unlikely, and the molecules stay in the upper laser level long enough to produce stimulated emission at microwave and infrared frequencies.

Unlike their man-made counterparts, cosmic masers do not generate beams. Instead, they radiate stimulated emission in all directions, like other hot clouds of interstellar gas radiate light. In fact, without special instruments they look just like other gas clouds. Astronomers did not realize the light came from stimulated emission until they analyzed the spectra of the gas and noticed some emission lines were unusually bright.

3.3 RESONANT CAVITIES

Cosmic masers show that population inversions are possible in nature, and that stimulated emission can produce light on narrow lines. However, it takes more than a population inversion to make a laser. It also takes a resonant cavity to produce oscillation.

The resonant cavity for a microwave maser is a metal box that resonates at the microwave frequency. That is, the cavity echoes, like sound waves bouncing off the walls, floor and ceiling of a shower or bathroom if you hit the right note. Such resonances happen when an integral number of wavelengths equals the distance the wave has to travel to make a round trip in the box. Microwave cavities fit only a small number of waves because microwaves are measured in centimeters and the cavities typically are no more than a few times wider than the wavelength, as shown in Figure 3-3 for cavities holding one, two, and four waves.

Laser cavities differ from maser cavities in important ways. Laser wavelengths are much shorter, 0.4 to 0.7 micrometer for visible light, so thousands to more than a million waves can fit in a reasonable length cavity. Placing mirrors on the two opposite ends of a long thin laser cavity creates a single resonance between those two mirrors, not the multiple resonances between the six walls of a boxy microwave maser cavity.

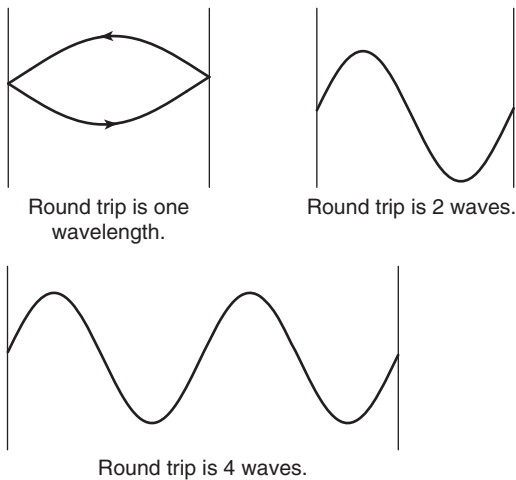


Figure 3-3. Waves resonate in a laser cavity when a round trip is 1, 2, or 4 waves.

The rest of this section explains how the resonance between the two mirrors produces a laser beam. This explanation is general, and as you will learn in later chapters, lasers come in many types, and not all of them fit this model exactly. However, this simple model will give you a basis to understand all types of lasers.

3.3.1 Amplification and Gain

Stimulated emission amplifies light in a laser by generating photons that are exact copies of input photons. (Remember that “laser” comes from light amplification by the stimulated emission of radiation.) Laser oscillation begins with spontaneous emission by the laser medium on the same transmission that is to be amplified.

A population inversion provides the optical spark that triggers the cascade of stimulated emission that generates a laser beam. A spontaneously emitted photon stimulates an atom in the upper laser level to emit another photon on the same transition with the same wavelength and transition as the first. Both the original spontaneously emitted photon and the stimulated-emission photon can then stimulate other atoms to emit additional photons. As long as the population inversion lasts, each emitted photon is more likely to stimulate emission than to be absorbed by an atom in the lower laser level. One photon stimulates the emission of a second, those two stimulate the emission of two more, those four stimulate the emission of four more, and so on in a rapidly growing cascade.

Laser physicists measure the amplification as *gain*, the amount of stimulated emission a photon can generate as it travels a unit distance. For example, a gain of 2 per centimeter means that a photon generates an average of 2 stimulated-emission photons each centimeter it travels. The amount of gain depends on the atom, the transition, and the environment inside the laser.

Power increases exponentially with gain because the photons produced by stimulated emission in turn stimulate emission further along their path. So the power increases, rather like an investment increases with compounded interest, but the power compounds continually rather than incrementing at certain intervals.

The *amplification factor* A , which measures the increase in power through a length of laser medium L with gain G per unit length, is an exponential:

$$A = \exp(GL) = e^{GL} \quad (3-1)$$

The exponential means that a gain of 1 per centimeter produces an amplification factor of 2.72 after 1 cm and 148 after 5 cm. Many lasers have lower gain, but their amplification factor grows with distance. For a gain of 0.05/cm, the calculated amplification factor is 1.0513 after 1 cm, 1.284 after 5 cm, and 1.6487 after 10 cm. That exponential increase cannot be sustained indefinitely because the gain term depends on conditions within the laser medium, including changes in the populations of the upper and lower laser levels and the flux of photons as the laser power level grows.

The normal value given for gain is the *small-signal gain*, the gain when the light being amplified is still weak and many atoms remain in the upper laser level. As the amplified power increases, the gain declines and the power eventually *saturates* because the excited atoms are producing stimulated emission just as fast as they can. In other words, the input light is stimulating emission from all available atoms in the upper laser level, so the power cannot be increased further.

The actual gain dynamics can be complex. When a laser switches on, the initial cascade of stimulated emission grows at the rate of small-signal gain. What happens next depends on the geometry of the laser medium and the design of the laser cavity, described in Sections 3.3.2 and 3.3.3.

3.3.2 Geometry of the Laser Medium

The geometry of a laser depends both on the physics of stimulated emission and on the optical gain of the laser material. Stimulated emission continues in the same direction as the photon that stimulated the emission. Thus, the whole cascade of emission stimulated by one photon goes in the same direction. The further the light goes through the laser medium, the higher the stimulated-emission power (neglecting saturation effects).

Gas and most solid-state (non-semiconductor) lasers have relatively low gain, so those laser media usually are shaped as long thin cylinders to take advantage of how total gain increases with distance. Spontaneous emission is by its nature random, so it goes in all directions, but most inevitably leaks out the sides of the tube, as shown in Figure 3-4. One mirror, called the output mirror, transmits a fraction of the laser light to produce the beam at right. Other light reflected back and forth between the resonator mirrors stays within the laser cavity to be amplified strongly by stimulated emission; light going toward the sides of the tube is lost.




Figure 3-4. Spontaneous emission starts in all directions, but only that directed along the length of the laser cavity is reflected back and forth and amplified to become the laser beam. Other spontaneous emission leaks out of the laser rod or tube.

Stimulated emission builds up along the length of a long thin rod, so it tends to concentrate light in that direction. If a laser tube or rod has no mirrors, as in Figure 3-5A, this length effect concentrates stimulated emission into an angle θ , defined by an inverse sine function:

$$\theta = \arcsin\left(\frac{Dn}{2L}\right) \quad (3-2)$$

where D is the rod diameter, L is length, and n is the refractive index.




Figure 3-5. Stimulated emission builds up the most along the longest dimension of the laser, which concentrates light along the length of a rod or tube. Top drawing shows output from a long rod with no mirror at the back; bottom drawing shows how putting a mirror on the rear end concentrates emission somewhat more.

The effective length of the laser rod can be doubled by adding a mirror on one end, so stimulated emission emitted toward the mirror is reflected back in the other direction and out of the other end. The further the light travels through the cylinder, the more the amplification. As shown in Figure 3-5B, this reflection also concentrates the emitted light into a narrower angle θ' :

$$\theta' = \arcsin\left(\frac{Dn}{4L}\right) \quad (3-3)$$

A few lasers with very high gain can operate in this way with only a single round trip through a laser cavity, including the excimer lasers in Section 7.7 and the edge-emitting semiconductor lasers in Section 10.6. However, most laser materials do not have high enough gain, so light must make multiple passes through the gain material to build up useful powers. This is done by placing mirrors at each end of the laser rod or tube, with one mirror transmitting a small fraction of the light, which becomes the laser beam. This approach not only generates higher power; it also gives much better control over the emitted light and gives lasers many of their special features such as coherence and single-wavelength output.

3.3.3 Mirrors, Resonant Cavities, and Oscillators

Putting a pair of parallel mirrors on opposite ends of a laser medium creates a *laser resonator*, *laser cavity*, or *resonant cavity*. Light waves bounce back and forth between the mirrors like sound waves in an echo chamber, but unlike the walls of an echo chamber or a maser cavity, the mirrors in a laser are separated by many optical wavelengths. This is the simplest type of resonator; we will describe important variations in Section 3.4.3.

The two mirrors in a laser cavity usually are not identical. One, called the *output mirror*, transmits part of the light resonating in the laser cavity, which becomes the laser beam, and reflects the rest back into the laser material. The fraction of power transmitted depends on the type of laser, its internal gain, and the desired output. The rear cavity mirror normally reflects all light back into the laser material, although in some cases it may transmit a small fraction of the light to a power monitor or measurement system.

The laser itself is called an *oscillator*, because like an electronic oscillator it generates an electromagnetic wave at optical frequencies, which are hundreds of thousands of times higher

than microwave frequencies. The idea that an oscillator generates a signal on its own is important. An amplifier normally boosts the power of a signal from an external source. An oscillator generates its own signal, determined by internal resonances. Optical amplifiers exist; like lasers, they generate power by the stimulated emission of radiation but, unlike lasers, they do not have a resonant cavity and cannot generate light on their own. You will learn more about them and how they are used in Section 6.4.

As you learned in Section 3.3.1, the seed that starts laser oscillation is the spontaneous emission of a photon along the length of the laser cavity, which stimulates emission that bounces back and forth between the mirrors. Figure 3-6 shows how one photon




Figure 3-6. Growth of stimulated emission in a resonant laser cavity.

starts a cascade of stimulated emission that bounces between the mirrors, with some photons escaping through the output mirror to form the beam. The amount of stimulated emission grows on each pass through the laser medium until it reaches an equilibrium level. It takes a lot of amplifying to multiply an initial few photons into a laser beam; every second, a 1-milliwatt red laser emits 3.2×10^{15} photons.

3.3.4 Balancing Gain and Loss in a Laser

From the time a continuous-output laser turns on until it reaches a stable output power, the power bouncing back and forth between the mirrors grows a bit on each pass through the laser cavity. Then it reaches a point where the gain is balanced by losses so the output remains at a steady level, which is often called *continuous wave*. This means that the power increase during a round trip of the laser cavity must equal the sum of the power emerging in the beam and power lost inside the laser:

$$\text{Power Increase} = \text{Lost Power} + \text{Output Power} \quad (3-4)$$

Losses can be low, but cannot be reduced to zero. No mirror can reflect every photon that reaches its surface. Material in the laser cavity absorbs some light. These losses can be important if the gain is low, as in many continuous-wave gas lasers, and they must be considered in selecting transmission of the output mirror.

For example, suppose that light is amplified 5% during a round trip of a continuous-wave laser cavity. One percent of the light is absorbed in the rear cavity mirror, and another 1% is absorbed by the laser medium. The output mirror transmits 2% of the light (the laser beam), absorbs 1%, and reflects the remaining 97% back into the laser cavity.

Output mirror transmission determines the relative power of laser light inside the laser cavity and in the external beam. In our example, the output mirror transmits 2% of the light and beam power is 1 mW, and the power inside the laser cavity is about 48.5 times higher, 48.5 mW. If the output mirror transmitted only 1% of the light, the power inside the laser cavity would be about 98 times higher than the beam power.

Lasers with higher gain inside the cavity require output mirrors with higher transmission. For example, if the internal gain within

the cavity was 25% and internal losses were 5%, the output mirror would need to transmit 20% of the light, and the internal power would be about five times the beam power.

3.4 LASER BEAMS AND RESONANCE

Now that we have “built” our laser on paper, let us go back and look more closely at some important parts of laser operation that were glossed over. We will cover some of these in more detail later, but we should explain a few things before going further.

3.4.1 Beam Characteristics

You may think of a laser beam simply as a collection of parallel light rays that form a bright spot on a screen. However, if you drew a line across the laser spot and measured the light intensity along that line, you would find that intensity varies, peaking in the middle and dropping off at the sides. Because the power drops off gradually, it can be hard to define the edge of the beam, so, normally, the *beam diameter* is defined as the point at which power drops to a certain fraction of the central intensity (usually $1/e^2$, where e is the root of natural logarithms).

A laser beam spreads with distance once it leaves the laser. This is not obvious over the scale of a room for most lasers, but it is important over longer distances. If you know the spreading angle, called *beam divergence*, you can calculate the size of a laser spot at a certain distance from the laser. Divergence, like beam diameter, is measured to points at which the beam intensity has dropped to a certain level. You will learn more about these laser characteristics in Chapter 4.

3.4.2 Cavity Resonances and Modes

Resonances arise from the wave nature of light. As explained in Section 3.3, a round trip of the cavity has to equal an integral number of wavelengths or, equivalently, the cavity length must equal an integral number of half-waves. Figure 3-3 showed how a few waves fit into a microwave maser cavity, but with few exceptions, laser cavities are hundreds to hundreds of thousands of wavelengths long.




Figure 3-7. Light waves are resonant if twice the length the laser cavity equals an integral number of wavelengths.

Figure 3-7 shows how light waves fit into a laser cavity, although the wavelengths are shown much larger than life so you can see them clearly. Note how neatly the light waves fit into the cavity. This coherence arises from a resonance produced by the nature of stimulated emission, feedback from the laser cavity, and the interference of light waves. Stimulated emission amplifies light across a range of wavelengths (called the *gain bandwidth*), and the cavity mirrors reflect all of them back and forth. However, as the light waves bounce back and forth, they undergo interference. Waves that fit exactly into the cavity interfere constructively, with their amplitudes adding together. Waves that do not fit exactly are reflected out of phase, so they interfere destructively, reducing their power. This resonance selects light with wavelengths λ that fit into the laser cavity of length L so that N waves fit into the laser cavity:

$$N\lambda = 2L \quad (3-5)$$

Destructive interference does not cancel other wavelengths completely, but it does reduce their gain. As Equation 3-1 shows, amplification depends exponentially on gain, so the wavelengths with higher gain soon overwhelm those with low gain. For example, suppose gain on resonance is 0.05/cm, and gain off resonance only half that value, 0.025/cm. The amplification factor after just one meter will be 148 on resonance but only 12 off resonance. The longer light oscillates back and forth in the laser, the more the off-resonance wavelength fades away.

The need for precise resonance might seem to be very restrictive, but that restriction is offset by the fact that light wavelengths are much smaller than most laser cavities. For example, 30 centimeters, a typical round-trip distance in a small helium–neon laser, equals about 474,000 waves at the laser’s 632.8-nanometer red line. Resonance is possible not just at 474,000 wavelengths, but also at 474,001, 474,002, and so on. Each of these resonances is called a

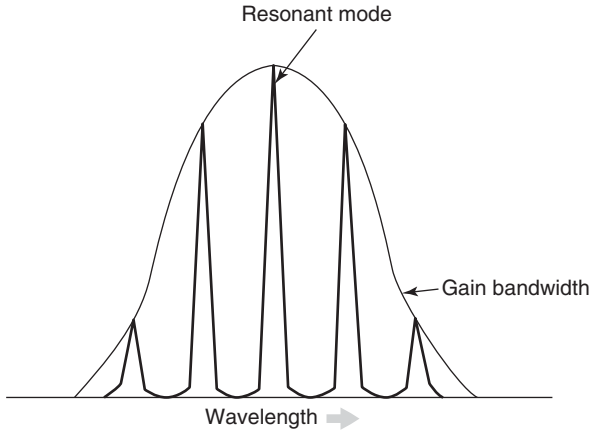


Figure 3-8. Several cavity resonances can fit within the gain bandwidth of a typical gas laser.

longitudinal mode of the laser, and each has a slightly different wavelength, with an integral number of waves fitting in the cavity, as you saw in Figure 3-7.

In addition, several effects combine to spread laser gain across a range of wavelengths, called the *gain bandwidth*, which can span several resonant modes, as shown in Figure 3-8 for a helium–neon laser. Many other lasers show the same effect. Cavity optics can be designed to limit laser emission to one longitudinal mode, but operation on multiple modes or a broad band is fine for many purposes.

3.4.3 Types of Laser Resonators

So far, we have said nothing about the shape of the mirrors on the ends of a laser cavity. You might expect both mirrors to be flat and parallel to each other, a simple arrangement called a Fabry–Perot resonator that is shown in Figure 3-9A. However, several other resonator designs shown in Figure 3-9 work better.

Although the Fabry–Perot resonator is simple, it is hard to use in long, low-gain laser cavities because a slight misalignment of one mirror can cause photons reflected from it to miss the other mirror. If two flat mirrors are only half a degree out of parallel, a light ray directed along a 15-cm laser cavity would be reflected to a point 1.3 millimeters from the center of the other mirror and would miss the edge of a cavity mirror 2 mm in diameter. Even




Figure 3-9. Laser resonators can use several different mirror combinations.

with smaller misalignment, light bouncing back and forth between the two mirrors would soon be lost.

Using one or two curved mirrors in the resonator can avoid such light losses. Following the light rays in Figure 3-9B–E shows how the focusing power of the curved mirrors directs light back into the cavity even if the light rays are not aligned precisely with the resonator axis. Curved mirror combinations that reflect light back toward the other mirror are called *stable resonators* because light rays will keep bouncing between the mirrors indefinitely if there are no losses. Stable resonators are most attractive for low-gain lasers.

Other resonator designs work better in high-gain lasers. Look carefully at Figure 3-9, and you can see that light rays reflected between cavity mirrors in some stable resonators do not pass through all of the excited laser medium, so they cannot collect the energy stored there. The resonator design shown in Figure 3-9F, called an *unstable resonator*, can collect light from those regions because it lets light rays drift off to the sides, making the cavity unstable. Eventually, the light rays drift off to the sides and leak out of the cavity, but those losses are acceptable in high-gain lasers because they are offset by the collection of light from more of the laser medium.

The reason for this difference is that light does not need to make as many round-trips between cavity mirrors in high-gain lasers as does the light in low-gain lasers. Low-gain lasers require highly reflective mirrors and low cavity loss to sustain oscillation, so the light may be reflected 20 times before it exits the laser cavity in the beam. In contrast, high-gain lasers have less-reflective mirrors, so the average photon will be reflected at most a few times before being emitted.

A true plane-parallel Fabry–Perot mirror cavity is neither stable nor unstable, because it neither focuses nor diverges light rays. Although they do not work well for large lasers, their simplicity is very attractive for semiconductor lasers, which are very small and have high gain.

Some laser resonators differ from the simple linear configuration shown in Figures 3-6 and 3-9. Some cavities include one or more folding mirrors to keep the light following a zigzag path between cavity mirrors. Three, four, or even more mirrors may form a ring cavity, with one mirror transmitting some light and serving as the output mirror, as shown in Figure 3-10. In this design,

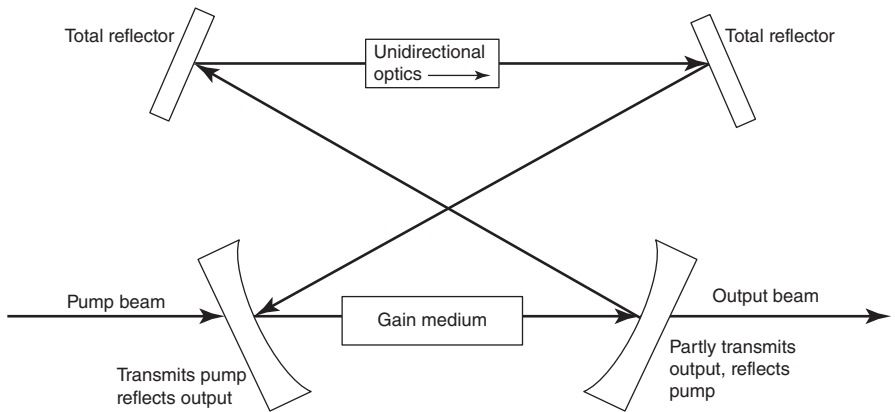


Figure 3-10. A series of mirrors may reflect light around a closed cavity, called a “ring,” although many types have other shapes, such as Figure 3-8. Pump light enters from left and is reflected around the cavity to stimulate emission in the gain medium in the bottom arm. Special optical components in the top arm limit laser oscillation to one direction, and the beam emerges through the output mirror at left.

pump light enters through one mirror, but pump light could also be focused onto the gain medium in another direction.

We will not go into detail on the theory of laser resonators. It bogs down in the sort of mathematical complexity you are reading this book to avoid, and many of the results are at best arcane. However, laser resonators do (quite literally) shape both the intensity distribution in the output beam and the rate at which the beam diverges.

3.4.4 Intensity Distribution and Transverse Modes

The intensity distribution across a laser beam is crucial for many applications. That distribution depends on a set of *transverse modes* that exist across the width of the laser beam, separate from the longitudinal modes described in Section 3.4.2. Those transverse modes arise from the *boundary conditions* for the electric and magnetic fields of light waves in the laser cavity. Transverse modes may have one, two, or more intensity peaks in the central part of the beam, as you can see in Figure 3-11.

The simplest or “lowest order” transverse mode is a smooth beam profile with a peak in the middle, as shown at the top of Figure 3-11. The profile at left shows the intensity of a cross-section

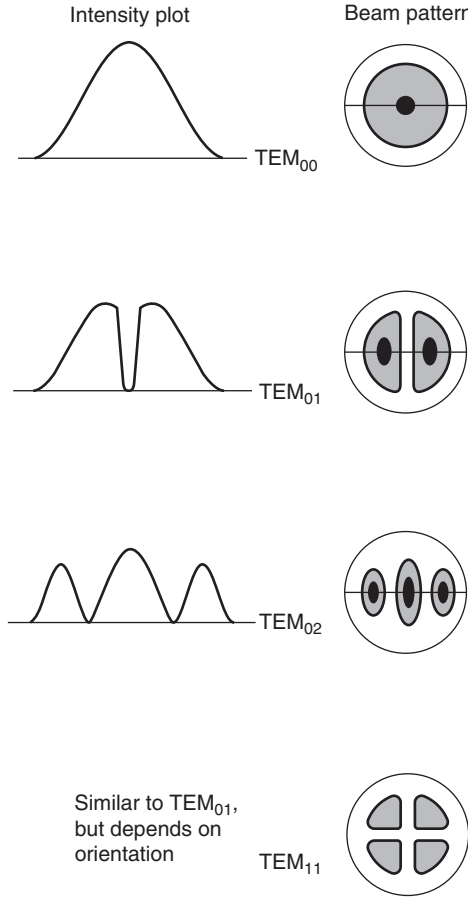


Figure 3-11. TEM₀₀ mode beam and a sampling of other transverse beam modes.

through the middle of the beam, which is called a Gaussian curve, after the famed mathematician Carl Friedrich Gauss.

The math describing the intensity across the beam may look more complex than the beam, but is worth showing because Gaussian beams are important. The power at a point $I(r)$ in the beam is given by

$$I(r) = \left(\frac{2P}{\pi d^2} \right) \exp(-2r^2/d^2) \tag{3-6}$$

where P is beam power, r is the distance from the beam center, and d is the size of the laser spot measured to a point where the intensity is $1/e^2$ intensity of that in the middle. This first-order

mode is called the TEM_{00} mode, where the T stands for transverse and the E and M stand for electric and magnetic modes, respectively. The numbers that follow are the number of zero-intensity points inside the pattern, with the first number giving the number of null points in the electric field and the second giving the number of nulls in the magnetic field. Recall that the electric and magnetic fields are perpendicular to each other.

As you might expect, there are a large family of TEM_{mn} modes, where m and n are integers denoting the number of nulls in electric and magnetic fields, respectively. Thus, a TEM_{01} beam has a single minimum dividing the beam into two bright spots. A TEM_{11} beam has two perpendicular minima (one in each direction) dividing the beam into four quadrants, as shown in the bottom part of Figure 3-11, and so forth. The TEM_{00} mode is desirable because it suffers less spreading than higher-order modes. A low-gain laser medium with a stable resonator can readily produce this lowest-order mode; proper design adds losses to suppress the oscillation of higher-order modes. However, some stable-resonator lasers operate in one or more higher-order modes, especially when they are designed to maximize output power.

You may see similar modes in other situations, where electromagnetic waves are guided through long pipes or structures called *waveguides*. An optical fiber is an optical waveguide; metallic waveguides are sometimes used for microwaves.

Unstable resonators have fundamentally different mode patterns. You can see why in Figure 3-12, which examines the unstable-resonator example from Figure 3-8F in more detail. In this




Figure 3-12. Near-field emission from an unstable resonator has a central minimum.

simple unstable resonator, the output mirror is a solid metal mirror that reflects some light back into the laser medium while the beam escapes around it. Thus, near the laser the beam profile has a “doughnut” cross-section—a bright ring surrounding a dark circle where light was blocked by the mirror. Far from the laser, this intensity distribution averages out to a more uniform pattern.

The details of unstable-resonator modes are even more complex than those of stable-resonator modes, but we can safely skip their many variations because they generally have little practical impact. In Chapter 4, we will learn how transverse modes affect the properties of laser beams.

3.5 WAVELENGTH SELECTION AND TUNING

Section 1.5.1 stated that a laser’s wavelength depends mostly on the laser material and secondarily on the optics. That is a good starting point, but it is time to dig deeper into the details, including laser transitions, subtle differences among laser materials, the role of laser cavities, tuning laser emission across a wider range of wavelengths, and converting laser output to other wavelengths. You do not need to know all the details, but you should understand what wavelengths are available and what influences them.

3.5.1 Picking a Transition

Atoms and molecules have many energy levels, with transitions possible between some pairs of levels, but impossible between others. Many thousands of laser transitions have been demonstrated in the laboratory, but only a few of them have become viable as commercial laser products.

Efficiency is crucial for commercial lasers, and Figure 3-13 shows four key factors affecting the efficiency of a laser transition:

- An efficient way to excite the atom to a high-energy level
- Efficient transfer of those excited atoms to a metastable upper laser level where they accumulate to produce a population inversion
- The availability of a lower laser level, which can easily be depopulated
- No other states that strongly absorb stimulated emission on the laser transition




Figure 3-13. Requirements for efficient four-level laser action.

Other factors are also important, particularly the availability of suitable optics, beam quality, power, and the availability of competing lasers. For example, an inexpensive red diode laser works perfectly well as a pointer, but if you need red light that is highly coherent, you should buy a red helium–neon gas laser. Other applications are less sensitive to wavelength, such as industrial laser machining, so a higher power is usually preferred even if the wavelength is not ideal.

Some lasers can emit light on several transitions at the same time, including the helium–neon laser. Typically, higher-gain transitions come to dominate after the laser is switched on and the light has gone a long way in the cavity, as shown in Table 3-1. Although the difference between gain of 0.04 and 0.05 per centimeter is small at 5 centimeters, after a meter it reaches a factor of 2.7, and by 10 meters it exceeds 22,000.

3.5.2 Laser Linewidth

Nominally, all laser transitions are at a sharply defined wavelengths, but in practice the transitions span a range of wavelengths, as you saw in Figure 3-8. In that case, for a helium–neon gas laser,

Table 3-1. Effect of gain on relative power of light traveling various distances

Distance	Relative power		Ratio
	Weak line	Strong line	
0.0 cm	1.00	1.00	1.00
1.0 cm	1.04	1.05	1.01
5.0 cm	1.22	1.28	1.05
10 cm	1.49	1.65	1.11
100 cm	54.60	148.00	2.71
1000 cm	2.35×10^{17}	5.18×10^{21}	22,043

the random motion of gas atoms spread the wavelengths emitted by neon atoms, but resonances in the laser cavity selected limited the amplified wavelengths to a few narrow ranges, producing the narrow peaks.

Solid-state lasers also often have wide linewidth, but for different reasons. The light-emitting atoms interact with the atomic lattice that holds them in a solid or liquid host material in different ways. For example, neodymium emits at 1054 nm in phosphate glass, 1062 nm in silicate glass, 1064 nm in a widely used crystal called yttrium–aluminum garnet (YAG), and 1080 nm in pure silicon dioxide glass. Some solid-state lasers have transitions that are inherently broad because of their interactions with the host material. You will learn more about those broadband solid-state lasers in Section 8.6.

Lasers with very high gain may also emit a wide range of wavelengths. That happens when the output mirror transmits a large fraction of the light oscillating in the laser cavity to form the beam. That keeps the light from being amplified long enough for the wavelengths with lower gain to fade away.

Lasers with broad linewidth offer two different opportunities. One is to tune the laser across that wavelength range, allowing a single laser to change its output color to meet varying needs. The other is to emit light across a broad range, which can allow the generation of extremely fast pulses of light, as you will learn in Section 4.6.

3.5.3 Cavity Optics and Laser Wavelength Selection

Earlier in this chapter you learned that resonances between laser light and cavity optics can select modes and wavelengths in a laser.




Figure 3-14. Turning a prism can tune laser wavelength.

Cavity optics can also be designed to select wavelengths in other ways.

One way is by limiting the wavelengths the optics transmit or reflect. Although we usually think of metal mirrors as reflecting uniformly across the spectrum, copper and gold absorb blue light, so they look reddish or yellow. Mirrors made by depositing many layers of thin films can reflect or transmit narrow bands of light, as you will learn in Section 5.3, which offer other possibilities.

The optical properties of a laser cavity and the laser wavelength can be tuned dynamically by placing a prism or diffraction grating in the laser cavity, so the colors spread out to show a spectrum. Turning the grating or prism can change the wavelength that oscillates in the cavity, as shown in Figure 3-14, with shorter and longer wavelengths being lost off the sides.

3.5.4 Changing Wavelengths: Harmonics and Other Nonlinear Tricks

So far we have discussed selecting wavelengths, but it is also possible to convert light to other wavelengths by using optical devices you will learn about in Chapter 6. These devices depend on light interacting with other materials and other photons in various ways to produce a number of interactions:

- *Harmonic generation* multiplies the frequency of light waves, shifting the light to shorter wavelengths. This is a particularly valuable tool because the best solid-state lasers emit in the near

infrared, and frequency doubling those wavelengths can produce visible light not easily generated by other lasers, particularly in the green and yellow. Higher harmonics can also be generated.

- *Sum and difference frequency mixing* mixes two wavelengths to generate a third, either the sum of the two frequencies (corresponding to a shorter wavelength) or the difference frequency (a longer wavelength). This is used in some of the laser-like sources described in Chapter 11.
- *Wavelength shifting* changes the wavelength by absorbing a photon at one wavelength and emitting a photon at a different wavelength. Usually, the atom that absorbs the light keeps a bit of it, thus shifting the light to a longer wavelength, but sometimes it may add a bit of energy, reducing the wavelength. Raman scattering and fluorescence are important examples.

3.6 LASER EXCITATION AND EFFICIENCY

Our general description of laser physics has given you an idea of how lasers work, but we have glossed over how atoms or molecules are excited to produce a population inversion. There are three major approaches, plus some others used only rarely. We will cover the basic ideas here and defer the specifics to later chapters describing specific laser types.

3.6.1 Choice of Excitation Methods

Two factors dominate the choice of laser excitation techniques—energy efficiency and the physical nature of the laser medium.

Input power generally comes in the form of electricity, so electrical excitation seems a logical approach. It is standard for most gas lasers and for most semiconductor lasers because both of them conduct electricity into the interior of the laser material. However, solid-state lasers in which the light-emitting atoms are embedded in a transparent glass or crystal are nonconductive, leaving optical excitation the only option.

Energy efficiency can be measured in two ways. The most meaningful for laser users usually is *wall-plug efficiency*, meaning the fraction of the input electrical energy emerging in the laser beam. The alternative is *conversion efficiency* or *pump efficiency*, which measures how much of the energy input to the laser is

converted to laser output. That neglects losses in converting energy into a form required by the laser, such as converting electricity into pump light for an optically pumped laser, which can be substantial.

Efficiency is not merely an environmental concern; for most applications, efficiency affects costs and cooling requirements. Power costs money, and any power not converted to light winds up as heat that must be removed from the laser. The amount of heat can be substantial. Some gas lasers have wall-plug efficiency lower than 0.1%, so producing a watt of laser light can generate a kilowatt or more of heat. Wall-plug efficiencies have improved greatly over the years, reaching 80% for some near-infrared semiconductor diode lasers covered in Chapter 10. The best efficiencies for high-power industrial lasers are fiber lasers (covered in Chapter 9) which can exceed 40%, so that the laser produces about 1.5 watts of heat for every watt of laser output.

3.6.2 Optical Pumping

Optical pumping uses light to excite atoms or molecules. As you saw in Figure 3-13, the excitation energy raises the laser species to a high-energy level, from which it naturally drops to the upper laser level. Stimulated emission then drops the atom or molecule to the lower laser level. It is conceptually simple and is standard for most solid-state lasers because light is the best way to deliver energy to atoms inside a transparent solid that does not conduct electricity.

Two types of light sources are used for optical pumping: intense lamps emitting white light, usually in pulses, and other lasers. It might seem odd to use lasers to excite other lasers, but it can be the most efficient case if the pump laser generates light efficiently, its emission wavelength matches a strong absorption line of the excited laser, and pump geometry allows efficient light absorption.

The wall-plug efficiency of optical pumping is the product of how efficiently the pump source converts electricity to pump light, and how efficiently the laser converts the pump light into laser energy.

Early solid-state lasers were pumped by flash lamps, which generated bright flashes of white light. The wall-plug efficiency of that approach was limited to the 1% range both by how much electrical power was converted into flashlamp light, and how much flashlamp energy could be converted to laser light. That *optical-to-optical conversion* was hampered by poor absorption of some parts of the spectrum by the laser species.

The development of high-power semiconductor diode lasers in the 1990s opened another possibility—laser pumping at a narrow range of wavelengths strongly absorbed by the solid-state laser material, greatly increasing the optical-to-optical conversion efficiency. Most semiconductor diode lasers have poor beam quality, but can convert much more of the input electrical energy into light than flashlamps. You will learn more about this process, called *diode pumping*, in Chapters 8 and 9, which cover solid-state and fiber lasers.

Flashlamps remain in use because they are less expensive sources of bright pulses of pump light than semiconductor diode lasers, but diode pumping is preferred for most applications because of its high efficiency.

As you will learn in Chapter 11, optical pumping is also used for some special-purpose lasers and laser-like devices.

3.6.3 Electrical Excitation of Gas Lasers

Most gas lasers are excited by an electrical discharge passing through the gas. As in fluorescent lamps, a high voltage is applied across gas inside a nonconductive tube, making current flow through the gas and transferring energy to atoms in the gas.

In a conventional fluorescent lamp, the electrons excite mercury atoms which spontaneously emit ultraviolet light that excites phosphors on the inside of a glass tube to emit visible light. In a typical gas laser, the electrons excite atoms or molecules, which either drop to the upper laser level or transfer the extra energy to another species in the gas so it can drop to the upper laser level. The resulting population inversion generates light, as described in more detail in Chapter 7.

Many electrically excited lasers produce continuous beams, but some generate pulses at regular intervals.

3.6.4 Electrical Excitation of Semiconductor Diode Lasers

Electrical currents also power semiconductor lasers, but the process occurs in a solid and is very different from excitation of gas lasers. Most semiconductor lasers are two-terminal electrical devices called diodes, in which current flows between materials with different compositions, exciting a population inversion in a thin *junction layer*, which is the boundary between regions where the semiconductor has different compositions.

Strictly speaking, the current in semiconductor diodes is carried by electrons and vacancies in electron shells called (appropriately enough) *holes* because they can accommodate electrons. Both electrons and holes move through the semiconductor and *recombine* in the junction layer, producing a metastable state which eventually releases its energy by spontaneous emission or stimulated emission. Chapter 10 explains the process in much more detail.

3.6.5 Other Excitation Mechanisms

A few types of lasers can be excited in other ways. Chemical reactions in flowing gases can produce population inversions, and military experiments have reached powers in the megawatt range, described in Section 13.10. Beams of electrons accelerated to high energies can produce intense ultrashort pulses at wavelengths from X-rays to the infrared in free-electron lasers, as described in Section 11.6, which are being used in scientific research.

3.6.6 Limits on Pump Efficiency

Wall-plug efficiency is inherently limited by factors including the nature of the laser transition, the deposition of pump energy, and the extraction of laser light. Some are inherent to the material, but others depend on design.

The nature of the laser transition sets a ceiling on maximum efficiency, as you saw in the simple four-level laser diagram of Figure 3-11. The pump energy raises the laser species from the ground state to a higher energy level, but only part of that input energy is released on the laser transition. That diagram makes it look like more than half of the energy is released on the laser transition, and in some cases that is true. Optical pumping always requires photons with more energy than the laser transition, but sometimes the difference is not large. The difference is less than 10% for an ytterbium-doped fiber laser emitting at 1030 nm and pumped by a diode laser emitting at 975 nm, but actual optical-to-optical conversion is not that efficient.

Other types of pumping can be quite inefficient. For example, the laser transition in argon-ion gas lasers is between two levels far above the ground state, so many times the transition energy is needed to excite the argon atoms to the upper laser level, contributing to a wall-plug efficiency of less than 0.1%.

You will learn more about laser efficiency in Section 4.5.

3.6.7 Confining Laser Energy

Laser excitation is most efficient when the excitation energy is confined to a small volume. The more concentrated the excitation energy is, the more atoms are likely to be excited, and the stronger the population inversion. It is like heating an object; the smaller the volume being heated, the more its temperature will increase.

Optical confinement of stimulated emission works similarly, by concentrating stimulated emission in the zone where a population inversion exists to amplify it further. The overall result is that confinement of both excitation energy and stimulated emission tends to improve laser operation and efficiency. This is critical for some types of lasers, such as ruby, in which more than half the chromium atoms that produce the laser light must be excited to produce a population inversion.

Confinement plays an important role in enhancing the efficiency of two important classes of lasers—fiber and semiconductor lasers. In fiber lasers, both stimulated emission and optical pump light are confined to the central core region, which guides light along the fiber. In semiconductor lasers, both current flow and light transmission are confined to certain parts of the laser. You will learn more about confinement in fiber and semiconductor lasers in Chapters 9 and 10, respectively.

3.7 WHAT HAVE WE LEARNED?

- A population inversion is needed for stimulated emission to dominate on a transition between an upper laser level and a lower laser level.
- Atoms must be selectively excited to the upper laser level to create a population inversion.
- Upper laser levels are metastable states with unusually long lifetimes.
- In three-level lasers, the laser transition usually is between the metastable upper laser level and the ground state, which is the lower laser level.
- In four-level lasers, the laser transition is between a metastable upper laser level and a lower laser level above the ground state, so a population inversion is easier to produce than in a three-level laser.
- Most continuous-wave lasers are four-level lasers.

- Laser oscillation starts with amplification of a spontaneously emitted photon.
- Resonant cavities help build up stimulated emission in a laser oscillator.
- The output mirror transmits a fraction of the power oscillating in a laser cavity to produce the laser beam.
- The simplest resonant cavity is a pair of flat parallel mirrors called a Fabry–Perot resonator on opposite ends of the laser. There are many other types.
- Stable resonators use curved mirrors to keep light rays bouncing between them without light leaking away to the sides.
- Unstable resonators allow light to leak away to the side of the output mirror, forming a donut-shaped beam.
- Laser amplification can increase laser power exponentially at first, but as power increases gain eventually saturates.
- The gain in laser power as the light oscillates in a resonant cavity equals internal losses plus the power transmitted in the output beam.
- Beam diameter is defined by the point at which power drops to $1/e^2$ of the central intensity.
- An integral number of waves at the oscillation wavelength must equal the round-trip distance in a laser oscillator.
- Lasers may oscillate simultaneously on two or more wavelengths within the gain bandwidth of the laser.
- Each resonant wavelength is a longitudinal mode along the length of the laser.
- Transverse modes define the intensity pattern across the laser beam.
- A TEM_{00} mode is the simplest transverse mode, with a single peak in the center.
- The laser cavity and the laser medium combine to determine what wavelengths a laser emits.
- The strongest laser line dominates emission unless the cavity optics select against it.
- Optical pumping uses light to excite atoms or molecules in solids, gases, or liquids.
- An electrical discharge can excite a gas laser.
- An electric current can excite a semiconductor laser.
- Wall-plug efficiency measures the fraction of input energy that emerges in the laser beam.
- Efficiency varies widely among lasers.

- Wavelengths can be changed by harmonic generation, sum and difference frequency mixing, or wavelength shifting.

WHAT'S NEXT?

Now that we have learned how lasers work, we will look at the important characteristics of lasers in Chapter 4.

QUIZ FOR CHAPTER 3

1. What is the major advantage of a four-level laser over a three-level laser?
 - a. More levels to excite atoms to.
 - b. The lower laser level is not the ground state.
 - c. More metastable states.
 - d. No advantage.
2. What is a metastable state?
 - a. A long-lived energy state that can serve as the upper laser level.
 - b. An energy state that is short-lived and unstable.
 - c. An energy state of a radioactive atom.
 - d. The ground state of a laser transition.
 - e. A state in which an atom can easily be excited.
3. How does stimulated emission increase laser power?
 - a. Linearly, each photon produced by stimulated-emission produces one additional photon.
 - b. Geometrically, each stimulated-emission photon produces exactly two additional photons.
 - c. Exponentially, each stimulated-emission photon can stimulate the emission of a cascade of additional photons.
 - d. Stimulated emission triggers spontaneous emission, which produces more photons.
4. The gain per unit length of a laser medium is 0.01 per centimeter. What is the amplification factor after 20 centimeters?
 - a. 0.02
 - b. 0.20
 - c. 1.001
 - d. 1.020
 - e. 1.221

5. The round-trip loss of helium–neon laser light in a cavity is 2%. The output mirror lets 1% of the light escape in the beam. When the laser is operating in a steady state, what is the round-trip amplification in the laser, measured as a percentage?
 - a. 1%
 - b. 2%
 - c. 3%
 - d. 5%
 - e. 99%
6. The helium–cadmium laser has a wavelength of 442 nm. What is the approximate round-trip length of a 30-cm long cavity measured in wavelengths?
 - a. 13,260
 - b. 135,700
 - c. 474,000
 - d. 679,000
 - e. 1,357,000
7. The 632.8-nm line of the helium–neon laser covers a wide enough range of wavelengths to have multiple resonances in a 15-cm laser cavity. What is the difference between the wavelengths of two adjacent longitudinal modes in that cavity? (Hint: Think of the wavelengths for which the cavity is N and $N + 1$ wavelengths long.)
 - a. 0.001 nm
 - b. 0.011 nm
 - c. 0.055 nm
 - d. 0.110 nm
 - e. 0.111 nm
8. How many internal minimum-intensity points are there in a TEM_{01} mode beam?
 - a. None
 - b. 1
 - c. 2
 - d. 3
 - e. 6
9. Which of the following excitation mechanisms cannot produce a population inversion?
 - a. Optical pumping with an external laser.
 - b. Passing an electrical current through a semiconductor.
 - c. Heating a gas.

- d. Passing an electrical discharge through a gas.
 - e. Optical pumping with light from a flashlamp.
10. Which of the following factors does not harm power-conversion efficiency in a laser?
- a. Atmospheric absorption.
 - b. Excitation energy not absorbed.
 - c. Problems in depopulating the lower laser level.
 - d. Inefficiency in populating the upper laser level.
11. What is the wall-plug efficiency of a diode-pumped fiber laser if the power supply converts 95% of the input electricity to DC to drive the diode laser, the diode laser converts 50% of the electricity supplied to it into light, and the fiber laser converts 60% of the diode pump light into laser output?
- a. 16.8%
 - b. 28.5%
 - c. 30%
 - d. 47.5%
 - e. 65%
12. The quantum defect is the fraction of energy in the pump energy that is inevitably lost because of the energy difference between the energy of a photon needed to excite a laser transition and the energy of the laser transition itself. What is the quantum defect in percent if you use a diode laser emitting at 975 nm to excite a fiber laser transition emitting at 1080 nm? (Hint: Remember that photon energy is inversely proportional to the wavelength.)
- a. 9.03%
 - b. 9.75%
 - c. 10.0%
 - d. 28.5%
 - e. Impossible to calculate because the excitation energy is smaller than the emitted photon energy.

LASER CHARACTERISTICS

ABOUT THIS CHAPTER

Now that you have learned how lasers work, it is time to look at the properties of lasers and laser light that affect how they are used. Properties such as the coherence, directionality, and single-wavelength nature of laser light set lasers apart from other light sources. Other properties are important in measuring laser performance, such as beam divergence, peak power, pulse duration, and efficiency. In this chapter, you will learn what these properties are and how they are important.

4.1 COHERENCE

Coherence may be the best-known property of laser light. Light waves are coherent if they are in phase with each other, that is, if their peaks and valleys line up, as shown in Figure 4-1. Laser light starts out coherent because stimulated emission has the same phase, wavelength, and direction as the stimulating light. The stimulated photon, in turn, can stimulate the emission of other photons, which are coherent with both it and the original wave. To maintain coherence, the waves that start in phase must remain so as they travel, which requires the waves all have the same wavelength.




Figure 4-1. Coherent light waves.

4.1.1 Coherence of Laser Light

There is no such thing as “perfect” coherence. Although our simple description may make it seem that one photon could trigger the whole cascade of stimulated emission in a laser beam, things are more complicated. Even if a single photon starts the cascade, more spontaneous emission can occur at any time, triggering another cascade of stimulated emission. Although such extra cascades of stimulated emission are on the same transition, they are not in phase with the initial light, and their wavelengths may not be identical.

In addition, photons produced by stimulated emission are not perfectly identical to the originals; their wavelengths can differ slightly, so the photons in the beam can drift out of phase over long distances. Tiny fluctuations within the laser, such as thermal gradients or vibrations, can also degrade coherence.

The degree of coherence differs widely among lasers. The more times light bounces back and forth through the laser cavity, the more coherent the light. That means high-gain lasers are less coherent, because their photons make only a few round trips of the laser cavity. High-gain lasers also have higher spontaneous emission in the background. Low-gain lasers tend to emit a narrower range of wavelengths, making them more coherent. Lasers oscillating in a single longitudinal mode have a narrower wavelength range than those oscillating in multiple modes, so they tend to be more coherent.

Coherence also depends on the duration of laser emission. The uncertainty principle causes the range of wavelengths in a pulse to increase as the pulse length decreases, so the shortest pulses tend to have the broadest wavelength ranges and be the least coherent. On the other hand, lasers that emit continuously can have the most narrow emission bands and the greatest coherence. Add this all together and you find that the most coherent beams come from continuous-wave, low-gain lasers operating in the lowest-order TEM_{00} mode.

4.1.2 Coherence Types and Measurement

There are two kinds of coherence: temporal and spatial. *Temporal coherence* measures how long in time light waves remain in phase. The term *temporal* is used because coherence is compared at different times, but temporal coherence also measures how far light travels while remaining temporally coherent. All light has temporal coherence over a certain *coherence length*, but that can differ widely. The temporal coherence of an incandescent light bulb is less than a micrometer, but that of a narrow-line laser can be many meters. Long coherence length is important for applications such as recording holograms of large three-dimensional objects.

Coherence length can be calculated from the light's nominal wavelength λ and the range of wavelengths $\Delta\lambda$ it contains. Equivalently, it can be calculated from the speed of light in vacuum c and the light source's bandwidth in frequency units, $\Delta\nu$:

$$\text{Coherence length} = \left(\frac{\lambda^2}{2\Delta\lambda} \right) = \frac{c}{2\Delta\nu} \quad (4-1)$$

A few simple calculations show what this means for some representative sources. Incandescent bulbs emit light from about 400 to 1000 nm, with an average wavelength of 700 nm. Plugging in those numbers gives a coherence length of 400 nm. An inexpensive semiconductor laser emitting at 800 nm with a 1-nm bandwidth has a coherence length of 0.3 mm. An ordinary helium–neon laser has a much narrower line width of about 0.002 nm at its 632.8-nm wavelength, corresponding to a coherence length of 10 centimeters. Stabilizing a helium–neon laser so it emits in a single longitudinal mode limits the line width to about 0.000002 nm, which extends the coherence length to about 100 meters.

Spatial coherence, on the other hand, measures the area over which light is coherent. Strictly speaking, it is independent of temporal coherence. If a laser emits a single transverse mode, its emission is spatially coherent across the diameter of the beam, at least over reasonable propagation distances. Spatial coherence is essential for a laser beam to be highly directional.

4.1.3 Interference and Coherence Effects

The interference effects described in Chapter 2 require some degree of coherence. If you superimpose many incoherent light waves of

different wavelengths, interference effects average out. This is why we do not see interference effects in places illuminated by light bulbs or the sun. Adding light waves from a few sources with some degree of coherence produces observable interference effects, as in the two-slit example of Figure 2-3.

These interference effects are desirable in many measurement applications, because they let us measure distance by counting in units of wavelength. Because light waves are so small, this lets us measure distance very precisely, a practice called interferometry. Interference effects also manifest themselves in other ways.

Coherence is not always desirable, however. Small-scale turbulence in air causes slight random shifts in the paths of coherent light passing through it, producing shifting patterns called *speckle*, which you can see when you look closely at a laser spot on the wall. Those patterns are an array of light and dark zones produced by interference among coherent light waves that take slightly different paths through the air. Speckles contain information on the quality of the air and the surface, but that information is essentially noise when you see speckles on the wall, which makes lasers poor for lighting.

4.2 LASER WAVELENGTHS

Laser light is *monochromatic*, meaning single-colored, but as you learned in Chapter 3, stimulated emission generates a range of wavelengths that depends on both the nature of the transition and the optics of the laser cavity. To understand what that means, we need to look more closely at what happens inside lasers and how it affects laser output.

4.2.1 Gain Bandwidth and Amplification

The bandwidth of a laser transition can be measured in two different ways: as how the gain on the laser transition varies with wavelength, and as the range of the wavelengths in the laser beam. That may seem a small difference, but it really is not.

The wide, low curve in Figure 4-2 shows the *gain bandwidth*, the range of wavelengths over which stimulated emission produces gain in the laser medium. The curve shows the probability of stimulated emission at a particular wavelength, which is highest in the




Figure 4-2. Gain bandwidth and laser emission on a laser transition, drawn to fit on the same scale. The laser peak is much higher.

middle of the gain band and lower at the edges. It is essentially small-signal gain.

The sharp, narrow curve in Figure 4-2 shows the *laser bandwidth*, the relative power across the range of wavelengths in the beam emerging from the laser. The peak is not to scale; if it were it would not fit on the page.

The difference between the two arises from the difference in amplification across the gain bandwidth. Suppose that the gain is 0.1/cm at the center of the transition and 0.01/cm near the side. If we use Equation 3-1 to calculate the amplification factor in 100 cm of laser medium, we find that the gain on the higher-gain line is 8100 times more than the gain on the weaker line. So essentially that is a large-signal gain within the laser rather than the small-signal gain of a short length of amplification. The sharp peak of the laser bandwidth is what happens when the broader gain curve builds up laser power inside the laser.

The gain bandwidth depends on the laser material and operating conditions. The laser bandwidth is also affected by the cavity optics. If the optical properties of the cavity are uniform across the gain band, the laser bandwidth will peak sharply at the peak of the gain band. However, if the cavity optics are designed to select a particular wavelength, as described in Section 4.2.3, the laser bandwidth will peak at the selected wavelengths. Laser oscillation can occur at any wavelength where gain is large enough to overcome losses in the laser cavity, so the cavity optics can select wavelengths across the gain band.

4.2.2 Line Broadening

Several effects can broaden the line width of a laser transition.

The quantum world is inherently uncertain. The likelihood of a transition is a quantum-mechanical probability function, which peaks at the nominal transition wavelength, and is lower, but not zero, at nearby wavelengths.

The continual random motion of atoms and molecules in a gas shifts wavelengths by a process called the Doppler effect. If an atom moves toward you, the atom comes close in the interval between wave peaks, making the wave shorter, an effect called blue-shifting. If the atom moves away, the wavelength is stretched, or red-shifted, toward longer wavelengths. The atoms in a gas laser are moving randomly, so their motion spreads the range of emitted wavelengths in both directions, increasing the line width.

Line broadening also comes from interactions between excited atoms and their neighbors. In gases, collisions between atoms and molecules cause broadening, which increases with gas pressure because that squeezes atoms and molecules. Most gas lasers normally operate at low pressures, but some gas lasers operate at much higher pressures, at which collisional broadening can merge separate laser emission lines into a continuous spectrum, as in the carbon dioxide lasers in Section 7.6. In solids, atoms can transfer vibrational energy to close neighbors, which also affects emission line width.

4.2.3 Wavelength Selection by a Cavity

The laser cavity combines with the gain bandwidth to determine the range of wavelengths at which a laser oscillates. Recall that in Section 3.4 you learned that a laser cavity of length L resonates at many wavelengths, described by Equation 3-5. If the laser cavity is thousands of wavelengths long, more than one resonant wavelength may fall within the gain bandwidth of the laser transition, as shown in Figure 4-3. Each resonant wavelength is a distinct longitudinal mode.

The actual resonant wavelength depends on another factor not included in Equation 3-5—the refractive index of the material in the laser cavity. That is acceptable for gas lasers because, in most cases, the refractive index of the laser gas, like that of air, is close to one. However, that oversimplification becomes a problem when the laser




Figure 4-3. Multiple longitudinal modes fall within the gain bandwidth of a gas laser.

medium is a solid. As you saw in Table 2-4, solids and semiconductors typically have refractive indexes from 1.4 to 3.5, which reduce wavelengths inside the laser medium. To correct Equation 3-5 and make the equation for cavity resonance exact, we should include the refractive index n , along with wavelength λ , cavity length L , and an integer N :

$$2nL = N\lambda \quad (4-2)$$

The optics of the laser cavity can be designed to affect the wavelength generated by tailoring how much light they absorb, transmit, and reflect at different wavelengths. This can be done by applying coatings to the surface of lenses and mirrors, or inserting wavelength-selecting filters, as described in Section 5-3. Laser cavity optics can also be adjusted to change the laser wavelength, as described in Section 11.1. You will learn more about such tunable lasers later in the book.

The actual profile of wavelengths emitted by a laser is the product of the gain curve in the laser cavity, the envelope of longitudinal oscillation modes, and the optical properties of the laser cavity. The profile in Figure 4-3 is based on typical emission from a gas laser, which can oscillate in a few longitudinal modes.

Adding components to the laser cavity can limit the laser to oscillating in a single longitudinal mode, which has a line width of about one megahertz, which is equivalent (for helium–neon lasers) to 0.000001 nm. The most common technique is to insert

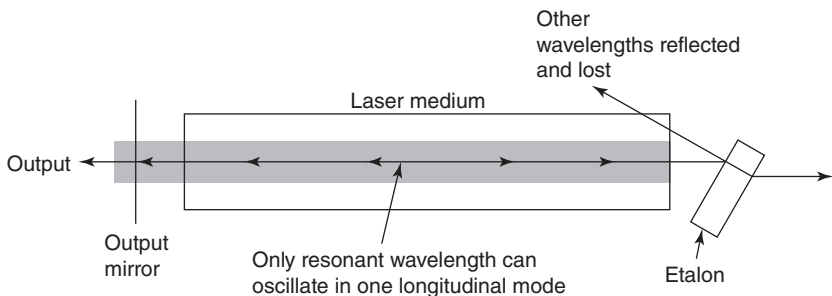


Figure 4-4. A Fabry–Perot etalon limits laser oscillation to one longitudinal mode.

a Fabry–Perot etalon into the resonant cavity. As shown in Figure 4-4, an etalon is a structure with a pair of parallel reflective surfaces, which is tilted at an angle to the axis of the laser cavity. At most wavelengths, the etalon reflects some light away from the cavity mirrors, increasing loss and suppressing laser oscillation, but in a narrow band interference effects eliminate the reflection, suppressing the loss and allowing laser oscillation. This limits the laser to operating in a single longitudinal mode, the wavelength of which can be varied by adjusting the spacing between the Fabry–Perot reflectors, or by turning the device in the laser cavity.

4.2.4 Single- and Multiwavelength Operation

The line narrowing caused by laser amplification, described earlier, does not prevent lasers from emitting simultaneously on two or more transitions.

A few lasers can emit on a family of closely spaced transitions at the same time. In the carbon dioxide laser, the main transition occurs when the molecule shifts between two vibrational modes. (Carbon dioxide emits on two such transitions one centered near $10.5\ \mu\text{m}$, the other near $9.4\ \mu\text{m}$.) As it shifts vibrational states, the CO_2 molecule also changes rotational energy levels, which are spaced much more closely than vibrational levels. Figure 4-5 shows the family of lines emitted by these lasers, each one a transition between different rotational sublevels during a vibrational transition. At high pressures, they blur together to form a continuum.

A few lasers emit simultaneously on two or more discrete transitions, typically involving different upper and/or lower laser




Figure 4-5. The many wavelengths emitted by a carbon dioxide laser near $10\ \mu\text{m}$.

levels. The argon- and krypton-ion lasers described in Chapter 7 are important examples.

4.3 PROPERTIES OF LASER BEAMS

We think of laser beams as tightly focused and straight, but their behavior is more complex than that. In general, beams *diverge* or spread with distance, but their precise behavior depends on the distance from the laser and on the size of the output port.

4.3.1 Beam Divergence

The spreading angle of a laser beam is called its *divergence*, and usually is measured in milliradians, thousandths of a radian, a unit of circular measurement. A full circle equals 2π radians, and one milliradian (mrad) equals 0.0572958° , or about 3.5 arcminutes. Radians are convenient because the tangent of a small angle roughly equals its size in radians, making calculations easy. Most continuous-wave gas lasers have divergence of around 1 mrad, but divergence usually is larger for pulsed lasers, and much larger for many semiconductor lasers which lack external beam-focusing optics. (Rapidly diverging beams often are measured in degrees.)

Beam divergence depends on the nature of the resonator, the width of the emitting area, and diffraction, which you learned in Chapter 2 arises from the wave nature of light. The theoretical lower limit on beam divergence, called the *diffraction limit*, arises from diffraction at the edge of the emitting aperture or output mirror. It is lowest for TEM_{00} beams. The diffraction limit roughly equals wavelength λ divided by the diameter D of the output beam.

$$\text{Diffraction limit (radians)} = \lambda/D \quad (4-3)$$

Focusing optics can change beam divergence. One example is the use of lenses to focus the rapidly diverging beam from a red semiconductor laser into the narrow beam required for laser pointers.

4.3.2 Near-Field Conditions

In the near field close to the laser, a narrow beam normally does not spread much, behaving like a bundle of parallel light rays. The distance over which the light rays remain parallel is called the *Rayleigh range* and depends on the beam diameter exiting the laser, D , and the wavelength λ :

$$\text{Rayleigh range} = D^2 / \lambda \quad (4-4)$$

For a visible beam of 500-nm light, the near-field range is 2 m for a 1-mm beam, and 50 meters for a 5-mm beam.

4.3.3 Far-Field Beam Divergence

In the far field beyond the Rayleigh range, beam divergence becomes the critical parameter in calculating beam diameter, as shown in Figure 4-6. Divergence angle is normally measured from the center of the beam to the edge. Beam intensity drops off gradually at the sides, so the edge of the beam is usually defined as where intensity drops to $1/e^2$ of the maximum value.

Calculating beam diameter D is a matter of trigonometry. Multiplying distance d from the laser by the tangent of the divergence angle θ gives the beam radius, which must be doubled to get the diameter:

$$D = 2d \times \tan \theta \quad (4-5)$$




Figure 4-6. Divergence of a laser beam, exaggerated to make it visible.

If divergence is under about 0.1 radian (6 degrees), you do not have to bother calculating the tangent; the angle in radians is a good enough approximation. That is the case for most lasers except semiconductor types.

4.3.4 Beam Waist and Divergence

Things get more complicated if the resonator mirrors are curved, so they reflect light back and forth through the center of a cylindrical laser medium. This results in the power distribution shown in Figure 4-7, with the power concentrated in a narrow *beam waist* in the middle of the laser rod or tube, and diverges toward the ends of the cavity, as shown by the curved lines on the sides. The presence of a beam waist affects the far-field divergence angle θ causing it to depend on the width of the beam waist W rather than the emitting aperture. In the far field, the divergence angle is approximately

$$\theta = \frac{\lambda}{\pi W} \quad (4-6)$$

The beam waist diameter depends on laser wavelength, cavity length, and design. For a simple confocal resonator (with two mirrors both having their focal points at the midpoint of the cavity as in Figure 4-7) of length L , the beam waist diameter W is

$$W = \sqrt{\frac{L\lambda}{2\pi}} \quad (4-7)$$

or about 0.17 millimeter for a 30-centimeter (one-foot) long helium–neon laser emitting at 632.8 nanometers. Note that this is smaller than the output spot size, typically about 1 mm for such lasers.




Figure 4-7. Beam waist in a confocal resonator.

4.3.5 Output Port Diameter and Beam Divergence

The diffraction limit mentioned in Section 4.3.1 is a fundamental rule of optics, which affects how an optical system can focus light. The minimum possible divergence angle θ roughly equals the wavelength divided by the diameter D of the output port through which the beam emerges. The formula for divergence angle also includes a constant K , which depends on the distribution of light across the output port:

$$\theta = \frac{K\lambda}{D} \quad (4-8)$$

The constant K equals 1.22 if light is equally bright across the entire opening. It is closer to one if the intensity is highest in the center and drops off smoothly toward the edges, which is the case for most laser beams. Equation 4-8 applies across the whole electromagnetic spectrum and for any emitting aperture, so it works for laser output mirrors, focusing lenses, telescopes, or radio antennas.

This formula tells us that the larger the output optics, the smaller the beam divergence. Thus, you need large optics to focus light onto a small spot. Conversely, the shorter the wavelength, the smaller the spot size. In practice, this means that you cannot have both tiny spot diameter and narrow divergence. If you start with a 1-mm beam from a helium–neon laser, for example, its minimum divergence is about 0.6 mrad.

Divergence is larger than these formulas indicate if the beam contains multiple transverse modes, which spread out more rapidly than a single-mode TEM_{00} beam. Far-field divergence increases roughly with the square root of the number of transverse modes ($N^{1/2}$) the laser emits, and the area of the illuminated spot is thus roughly proportional to the number of transverse modes.

Semiconductor diode lasers are a special case. As you will learn in Section 10.6, edge-emitting diode lasers emit light from rectangular apertures that typically are less than one micrometer wide, causing strong diffraction and high beam divergence in that direction. The apertures are much wider than they are high, so diffraction and beam divergence are less in the other direction, as covered in Section 10.8.1. Other semiconductor lasers differ in more subtle ways.

4.3.6 Focusing Laser Beams

The diffraction limit also affects how tightly lenses can focus light and how finely optical systems can resolve detail on distant objects. The minimum diameter of a focal spot S formed by a lens of diameter D and focal length f with light of wavelength λ is roughly

$$S = \frac{f\lambda}{D} \quad (4-9)$$

In practice, the minimum size of a focal spot roughly equals the wavelength of the light.

4.3.7 Laser Modes

Discussions of laser modes can be confusing because lasers have two types of modes, longitudinal and transverse, but laser operation is often specified only as “single-mode” or “multimode.” In most cases, descriptions of beam modes refer to transverse modes. Lasers oscillating in a single longitudinal mode often are called single-wavelength or single-frequency lasers and have narrower line widths than lasers oscillating in multiple longitudinal modes.

A laser operating in a single transverse mode can oscillate in two or more longitudinal modes. Many gas lasers emit a single TEM_{00} mode beam that contains two or more longitudinal modes, although special optics can restrict them to oscillating on a single longitudinal mode.

4.3.8 Beam Quality

Beam quality is an important topic for laser users, but “quality” can be a hard idea to pin down and measure quantitatively because requirements vary among applications. For some measurement applications, the key factor is the *beam profile*, the pattern of power variation across the beam. For laser machining, it often is “power in the bucket”—the amount of power delivered within a limited area, such as the zone being welded. Section 5.8.4 will describe important ways to measure beam quality, how they work and where they are applied.

4.4 LASER POWER

Power is a vital laser characteristic, but quantifying it can be tricky because it can be measured in different ways. This section explains the different ways to look at laser power.

4.4.1 Intracavity and Beam Power

Normally, laser power refers to the output power in the laser beam. However, the *intracavity* power that circulates within the laser cavity can also be important.

Intracavity power is the stimulated emission generated inside the laser that bounces back and forth between the laser mirrors. The output mirror transmits part of that power as the laser beam and reflects the rest back into the cavity for further amplification. The fraction of light exiting the laser cavity determines the intracavity power. For example, if the output mirror transmits 1% of the incident light, the output beam power is roughly 1% of the circulating power inside the cavity. If that laser has a continuous output power of one milliwatt, that means the intracavity power is about 100 mW. The ratio is not as simple to calculate in a pulsed laser, but as long as the output mirror reflects some light back into the cavity for further amplification, the intracavity power is higher—often much higher—than in the laser beam.

Some optical devices that work better at higher power can be placed inside the laser cavity to take advantage of this high intracavity power, as shown in Figure 4-8. For example, nonlinear devices like the harmonic generators described in Section 5.6.1 are more efficient at higher power, so they benefit from being placed inside the laser cavity.




Figure 4-8. Circulating power inside a laser cavity is higher than beam power, so devices that need high power may be placed between the mirrors and the laser medium.

4.4.2 Laser Output Power and Scaling

The output power in the laser beam depends on a number of factors including the gain, the nature of the laser cavity, and the inherent properties of the laser medium.

Many lasers are inherently feeble because their internal physics makes it impossible to reach high powers. For example, the red helium–neon laser can emit powers from a fraction of a milliwatt to about 40 mW, but becomes impractical at larger sizes.

Other lasers can be scaled over a much wider range, although design details change with power level. Commercial ytterbium-fiber lasers can generate continuous beams from a fraction of a watt to 100 kilowatts, but designs differ greatly. Small size, low cost, and easy use are important at low power; at higher power, the dissipation of waste heat becomes crucial. You will learn more about power ranges in the chapters describing individual laser types.

4.4.3 Power Measurement in Time and Space

Power in a laser beam normally is measured in watts. One watt equals one joule of energy flowing through the laser each second. This measurement is straightforward if the beam is continuous, although special equipment may be needed at the highest powers.

Power is an instantaneous measurement, so things get more complex if laser power varies with time. The output of pulsed lasers may be averaged over a number of pulses during a second, or an instantaneous peak power may be measured in an individual pulse. Pulses may also be characterized by the energy they contain. You will learn more about this in Section 4.6.

Another way to look at power is as *power density*, the number of watts delivered either per unit area or per unit solid angle rather than total power in the beam. Power density is important in applications such as laser machining, which requires certain power densities for particular purposes, such as melting metals to weld a joint. Beam divergence causes power density to decrease with distance from the laser, but focusing optics close to the target can concentrate the beam to higher power density.

The details can be complex, and as you will learn in Section 5.8.2 and Table 5-3, terms such as power, intensity, and irradiance have formal definitions for optical measurements that are more specific than in ordinary usage.

4.4.4 Modulation

Many laser applications require modulation of laser output power. For example, if a laser beam is to carry a signal in a fiber-optic communication system, the beam must be switched on and off to carry digital ones and zeros. Imaging and some other applications require modulating laser power continuously, by transmitting from 0% to 100% of the laser beam. You will learn more about modulators and modulation in Section 5.7.

Lasers can be modulated internally (also called directly), by changing the amount of power delivered to the laser, or externally, by passing light through a device called a modulator that transmits varying fractions of the light. The ease and effectiveness of direct modulation depends on the type of laser; it works much better for semiconductor lasers than for gas or solid-state lasers. External modulation can be used with any laser.

4.5 LASER EFFICIENCY

Lasers do not produce energy; they convert other forms of energy, including light, into a laser beam. Power-conversion efficiency affects the amount of waste heat that must be removed from the laser, and low efficiency can limit how the maximum output power a laser can generate.

Laser efficiency can be measured in several ways. As mentioned in Section 3.6, the most useful measurement generally is the overall or *wall-plug efficiency*, the fraction of the energy delivered to the laser (through the wall plug) that emerges in the laser beam. That overall efficiency, in turn, depends on a number of other efficiency terms. We will start by looking at overall efficiency, then turn to its components, and finally mention a few cases in which other measures of efficiency are more important.

4.5.1 Elements of Overall Efficiency

To show what contributes to wall-plug efficiency, we will look at a generic electrically powered laser, but the same principles apply to all lasers. Electrical power enters the laser through a power supply, which converts much (but not all) of the input energy into drive current that is passed through the laser material. Much (but not all) of that electrical energy is deposited in the laser material

(energy-deposition efficiency). Much of the deposited energy excites atoms or molecules to produce a population inversion. The laser converts much of the energy in the population inversion into light by stimulated emission. Much of the stimulated emission emerges in the laser beam.

The central point is that each step loses some energy, and the overall wall-plug efficiency is the product of the efficiencies of those processes. For our generic laser, these are:

Power-supply efficiency

× Energy-deposition efficiency

× Excitation efficiency

× Fraction of excitation energy available on laser transition

× Fraction of excited atoms that emit laser light

× Cavity output-coupling efficiency

= Overall efficiency

Plug in some numbers and you can see why lasers have low wall-plug efficiencies. If the power supply, energy deposition, and excitation each are 80% efficient, the laser transition represents 80% of the excitation energy, 80% the atoms emit laser light, and 80% of the laser energy is coupled out of the cavity in the beam (all of which are optimistic assumptions for most gas lasers), the wall-plug efficiency is $0.8 \times 0.8 \times 0.8 \times 0.8 \times 0.8 \times 0.8 = 0.2621$, or 26%.

Similar considerations apply to solid-state and fiber lasers, which are optically pumped and thus have extra steps, although many are more efficient than most gas lasers.

Let us look at the major elements in the efficiency equation.

4.5.2 Power-Supply Efficiency

The power supply converts 120-volt alternating current (or higher voltages used for industrial lasers) to the form needed to excite the laser. This is straightforward for most lasers, but no power supply is 100% efficient, so some energy is lost. More energy is lost when the electrical input has to be converted to high-voltage pulses to drive pulsed gas lasers.

Power requirements and power supplies differ among lasers. Gas lasers generally require high voltage to produce a discharge. The power supply for a semiconductor laser must convert wall

current into the low-voltage drive current. The “power supply” in an optically pumped solid-state or fiber laser actually involves two stages, one of which is a laser or lamp which generates pump light, and the other of which delivers the pump light into the solid laser material. This means the overall or wall-plug efficiency is the product of the efficiency of converting electrical power to pump light, times the efficiency of converting that optical pump power into the output laser beam. As you learned in Section 3.6.2, pumping with diode lasers can be much more efficient than pumping with lamps.

4.5.3 Excitation Efficiency

All the excitation energy that reaches the laser medium does not excite atoms to the upper laser level. Some energy may pass through the laser medium without being absorbed. Some absorbed light may not excite the light-emitting species into the energy levels needed to produce laser action; for example, it might be absorbed by some non-emitting gas in a gas laser.

4.5.4 Laser Transition Efficiency

Some losses are inherent in the laser transitions. Exciting a four-level laser raises an atom to an excited level above the upper laser level. After stimulated emission drops the atom to the lower laser level, it then drops further to the ground state. Figure 4-9 shows this process for an atom with energy levels indicated in energy units—electron volts, the energy an electron needs to move across a one-volt potential. Initially, absorption of 4 eV raises the atom to an excited state, from which it drops to the upper laser level at 3 eV. The laser transition takes the atom to the lower laser level at 2 eV, from which the atom drops to the ground state at 0 eV. It takes 4 eV to produce a 1 eV laser photon, so only 25% of the pump energy emerges in a laser photon.

As you will learn in later chapters, laser transition efficiencies differ widely among lasers depending on factors including how far the lower laser level is above the ground state. For optical pumping, where photons excite the laser species, this fraction is called *quantum efficiency*, meaning the fraction of the pump photon energy emerging in a laser photon. The fraction of energy lost is called the *photon defect*.

Quantum efficiency, like the transition energy shown in Figure 4-9, is a theoretical upper limit. One reason actual laser




Figure 4-9. Only part of the energy that excites a four-level laser winds up in the laser beam. In this example, 4 eV of excitation produces only 1 eV of laser energy.

efficiency is lower is that the *threshold* for laser action is not reached until the population inversion is large enough for stimulated emission to offset other losses. Figure 4-10 compares the input of pump energy compared with the output laser power. Look closely at the threshold, and you can see that output initially rises slowly, then faster. Well above threshold, the output power increases smoothly with a *slope efficiency* (measured as the output power per unit input power) that is uniform over a wide power range. High slope efficiency is crucial for high laser efficiency at high power. As output power increases, total laser efficiency increases toward the limit of the slope efficiency, but at a higher power not shown the slope efficiency declines as the laser approaches its output limit.

4.5.5 Energy Extraction Efficiency

Energy is also lost in extracting light from the laser cavity. Figure 4-11 shows how the mirrors of a concentric resonator collect stimulated emission from only part of the laser medium. Although the whole laser medium is excited, the cavity resonator only collects light energy from the shaded region between the two mirrors.




Figure 4-10. After laser turns on above threshold, its efficiency rises toward the slope value of 10%. Leveling off at higher power not shown.

In general, an unstable resonator like that in Figure 3-12 can collect light from more of the laser cavity than stable resonators.

4.5.6 Waste Heat and Wall-Plug Efficiency

The energy lost inside the laser does not vanish; most of it heats the laser. Small and efficient lasers, such as semiconductor laser




Figure 4-11. Electrical discharge passes through a laser tube, but the resonator extracts energy only from the gas in part of the tube.

Table 4-1. Wall-plug efficiency ranges of commercial lasers

Type	Wall-plug efficiency
Argon ion	0.001%–0.1%
Carbon dioxide	5%–20%
Excimer (rare-gas halide)	1.5%–2%
Helium–cadmium	0.002%–0.014%
Helium–neon	0.01%–0.1%
Neodymium solid state	0.1%–20%
Semiconductor diode	1%–70%
Ytterbium-fiber laser	30%–40% at high power

pointers, can easily dissipate the heat by convection cooling into a room. But as shown in Table 4-1, many lasers have wall-plug efficiency less than 1%, and the higher the power, the more cooling needed. Forced-air cooling with a fan may suffice at low power, but higher powers require closed-cycle refrigerators or flowing water. A 20-watt argon laser that is 0.1% efficient produces 20 kilowatts of heat and needs a fire hose of cooling water flowing through to keep it from frying. A few lasers require cooling below room temperature, but they are rarely used outside the laboratory.

An important efficiency measure for optically pumped lasers is the *optical-to-optical efficiency*, which compares the input pump light to the laser output. This number can exceed 80% for laser pumping of some ytterbium-doped fiber lasers. Optical pumping with semiconductor diode lasers, called *diode pumping*, yields high wall-plug efficiency because much of the electrical power into the diode laser is converted into light.

4.6 PULSE CHARACTERISTICS

Like wavelength and power level, the duration of laser emission is critical for laser applications. Lasers fall into two basic categories: pulsed or continuous wave. Pulse durations range from milliseconds to femtoseconds (10^{-15} second), and pulses may be repeated at a steady rate or fired one at a time. How a particular laser operates depends on the nature of the laser, its design, on its control and power management systems, and on accessories used with the laser.

Some types of lasers are inherently limited to pulsed operation because their internal physics cannot sustain a steady population

inversion. In the three-level ruby laser, the lower laser level is the ground state, which quickly fills up as stimulated emission drops atoms to the ground state. Some four-level lasers cannot sustain a population inversion because atoms drop out of the lower laser level slowly, causing it to fill up and halt laser operation. Both types can be pulsed repeatedly.

Lasers may also be limited to pulsed operation by their need for powerful electric discharges to reach threshold, or by the need for their optics and laser media to cool down between pulses.

In most cases, pulsing is controlled by a modulator or electronic input signal, by optical components inside the laser, or by a power management system. This section will introduce key concepts of pulsed laser characteristics, measurement, and operation. You will learn more about pulsed lasers in general in Section 6-4 and for specific lasers in later chapters describing those types.

4.6.1 Pulse Characteristics and Measurement

Figure 4-12 shows how the power of a laser pulse changes over time. Typically, it rises slowly at first, then rapidly increases to its *peak power*, before dropping off symmetrically. Pulses have characteristic *rise and fall times*, which are measured at the points where laser power is between 10% and 90% of the maximum value (not




Figure 4-12. Power variation as a function of time (dashed lines show prepulse and long tail).




Figure 4-13. The interval between repetitive laser pulses is much longer than the pulse duration. Output is zero between pulses.

from zero to peak power). Some lasers may emit a small prepulse (before the main pulse), and/or have a long tail dropping off slowly after the main pulse, which are shown as dashed lines in Figure 4-12. *Pulse duration* typically is defined as “full width at half maximum,” the interval from the time the rising pulse passes 50% of the peak power to the time the power drops below that 50% point.

The *repetition rate*, or the number of pulses per second, is measured as a frequency in hertz, like electromagnetic waves, but the pulses normally are separated by dark intervals, as shown in Figure 4-13.

The energy in a laser pulse can be measured by collecting the light and recording its effects during the pulse, such as heating a sensor. Alternatively, pulse energy can be measured by calculating the area under the curve showing laser power over time in Figure 4-12. Mathematically, this is called integrating power over time and involves calculus. However, a useful approximation for most pulse shapes is to multiply the full-width, half-maximum power P by the pulse duration Δt to get pulse energy E :

$$E = P \times \Delta t \quad (4-10)$$

Using this formula, we find that a one-megawatt pulse lasting one microsecond delivers 1 J. If you know pulse energy and duration, you can flip the formula to calculate the half-maximum pulse power. Multiplying by 2 gives peak power.

$$P = \frac{E}{\Delta t} \quad (4-11)$$

If a series of identical pulses are spaced uniformly in time, you can modify this formula to calculate the *average power* emitted

during a series of pulses. In that case, you sum energy of the pulses emitted during an interval and divide by the duration of the interval. Thus, if each pulse has an energy E and the laser emits N pulses per second, the average power (in joules per second, which equals watts) is:

$$\text{Average power} = NE \quad (4-12)$$

If the laser emits 100 pulses per second, each with energy of 1 J, the average power is 100 W.

Another useful number is the *duty cycle*, the fraction of the time the laser emits light. If a laser switches from “on” to “off” once every microsecond, the duty cycle is 50% because it is on half the time. If the laser emits N pulses per second, each lasting a time Δt seconds, the duty cycle is a dimensionless number, the product of repetition rate times pulse length in seconds:

$$\text{Duty cycle} = N \times \Delta t \quad (4-13)$$

If the pulse length is 100 microseconds and the repetition rate is 1000 pulses per second, the duty cycle is 0.1 or 10%. That is roughly the duty cycle of the pulses shown in Figure 4-13.

4.6.2 Limits on Pulse Duration and Bandwidth

The Heisenberg Uncertainty Principle says that the product of pulse duration and the range of wavelengths a laser can emit during a pulse must be at least a minimum number. That means you need a very wide bandwidth to produce a very short pulse, or an almost continuous beam to generate a very narrow bandwidth. In other words, very short pulses cannot have a very narrow bandwidth because you cannot precisely measure both the location and energy of a photon at the same time. This is called the *transform limit*, because it is derived from the Fourier transform between time and frequency domains.

Numerically, the product of the pulse length Δt (in seconds) times the bandwidth $\Delta \nu$ (in hertz) must equal at least 0.441 for pulses with the Gaussian shape of Figure 4-12. This means that the bandwidth $\Delta \nu$ of a laser pulse can be no smaller than

$$\Delta \nu = \frac{0.441}{\Delta t} \quad (4-14)$$

The bandwidth does not become infinitely narrow for a continuous beam. Turn the formula around, and you see that a laser must have a broad bandwidth to generate short pulses.

$$\Delta t = \frac{0.441}{\Delta \nu} \quad (4-15)$$

For example, suppose you want to know what spectral bandwidth is needed to generate pulses lasting only 5 femtoseconds (5×10^{-15} second), which is only a few wavelengths of visible light. Reversing the equation to find $\Delta \nu$ in frequency terms yields a bandwidth of 88.2×10^{12} Hz or 88.2 terahertz (THz). For comparison, the visible spectrum of 400 to 700 nm is equivalent to frequencies of 430 to 750 THz, a bandwidth of 320 THz. That means a 5-femtosecond pulse would span more than a quarter of the visible spectrum.

4.6.3 Pulse Control in Lasers

Some lasers have a “natural” pulse length because the internal physics switch off the population inversion after a certain time. Examples include the solid-state ruby laser described in Section 8.4.1 and excimer gas lasers, covered in Section 7.7. But pulsing in most lasers is controlled externally, either by switching the power off and on, or by adding accessories that control how light circulates inside the laser cavity. These components are crucially important for many applications that require short pulses. Section 6.4 describes them in detail, but the basic concepts deserve a brief explanation here.

Most of these accessories are inserted into the laser cavity and create pulses by changing how photons circulate within the laser cavity.

- *Q-Switching* switches the *quality factor* or *Q* of the laser cavity between two states by changing the reflection or transmission of a component called a *Q switch* inside the cavity. Between pulses, the *Q* switch keeps photons bottled up inside the laser cavity, building up a large population inversion but emitting essentially no light through the output mirror because the laser remains below its oscillation threshold. Then it switches its internal state to put the laser cavity above the oscillation threshold, producing a burst of stimulated emission that extracts the energy stored in the population inversion in a giant pulse. *Q-switching* works best

for solid-state laser with long upper-state lifetimes that let them store large amounts of energy.

- *Cavity dumping* is somewhat similar, but is based on inserting a mirror into a high- Q resonant cavity that dumps the accumulated energy out through a separate output port. This dumps light already circulating in the laser cavity in a pulse lasting the cavity round trip.
- *Modelocking* is a technique that makes a laser oscillate simultaneously on many different longitudinal modes (described in Sections 4.2.3 and 4.3.7) to produce pulses lasting picoseconds to femtoseconds. You can also visualize modelocking as creating a clump of photons of various wavelengths that bounce back and forth between two laser mirrors, with a laser pulse emerging every time the clump reaches the output mirror. The more modes locked together and the wider their range of wavelengths, the shorter a pulse can be.

4.7 POLARIZATION

Polarization is an important property of electromagnetic waves, which consist of oscillating electric and magnetic fields that are perpendicular to each other, as you saw in Figure 2.1. Light is polarized when all the waves have their electric and magnetic fields aligned with each other.

The simplest type is *linear polarization*. The polarization direction is normally defined by the electric field, so linearly polarized waves have their electric fields aligned in the same plane.

Normally, most lasers emit unpolarized light, like the sun or light bulbs. However, a laser can generate linearly polarized light when an optical component is added that separates the two perpendicular polarizations. The usual choice is a Brewster window, a surface which is aligned at a critical angle θ_B , defined in air as the arctangent of the refractive index n of the window material:

$$\theta_B = \arctan n \quad (4-16)$$

For optical glass with refractive index 1.5, Brewster's angle is about 57 degrees.

As shown in Figure 4-14, the Brewster window transmits all light and reflects none if the polarization is in the plane of the page.




Figure 4-14. Operation of a Brewster window in a laser.

However, it reflects about 15% of the light if the plane of polarization is perpendicular to the page.

Placing a Brewster window between the laser medium and one mirror in the cavity increases losses for the plane of polarization perpendicular to the page. That means that light polarized in the plane of the page is amplified more and only light polarized in that plane will appear in the output beam. This produces a linearly polarized output beam with the same power that an unpolarized beam would have had. (Suppressing the other polarization concentrates laser power in the plane of the page, but it does not “throw away” power.)

4.8 WHAT HAVE WE LEARNED?

- Laser light nominally is coherent, with all waves having the same wavelength and their peaks and valleys all lined up in phase. But the coherence is not perfect, and the degree of coherence varies among lasers.
- Temporal coherence measures how long light waves remain in phase.
- Spatial coherence is the area over which light is coherent.

- Coherence length equals the speed of light divided by twice the bandwidth.
- Coherence is needed to produce interference.
- Coherent effects generate grainy noise patterns called speckle.
- Laser light is called monochromatic, but actually spans a range of wavelengths.
- Amplification narrows the width of laser lines so they are narrower than the gain bandwidth.
- The laser cavity combines with the gain bandwidth of the transition to determine the range of wavelengths at which the laser cavity oscillates.
- Multiple longitudinal modes may fall within the gain bandwidth of the laser transition in many laser cavities.
- Lasers can emit on single wavelengths or multiple transitions, depending on the laser cavity and the gain medium.
- Beam divergence decreases as the diameter of the output optics increases, and increases as wavelength increases.
- The diffraction limit is the lower limit on beam divergence.
- Focusing optics can change beam divergence.
- A narrow laser beam behaves like parallel rays and does not spread in the near field, to a distance called the Rayleigh range.
- Lasers oscillating in a single longitudinal mode are often called single-wavelength or single-frequency lasers.
- Lasers emitting a single transverse mode may oscillate in multiple longitudinal modes.
- Power in watts measures the amount of energy in joules flowing past a point per second.
- Intracavity power circulating in the laser resonator is higher than output power.
- Power density is the power per unit area or per unit solid angle.
- Wall-plug efficiency measures how much input electrical power a laser needs to generate light. It is the product of the efficiencies of several processes.
- Efficiency varies widely between types of lasers.
- Semiconductor diode lasers have the highest wall-plug efficiency.
- Threshold power is the input power at which the laser medium begins oscillation.
- Slope efficiency is the increase in laser output per unit increase of input power when the laser is operated above threshold. It is higher than total efficiency because the laser must reach threshold before it starts emitting light.

- Waste energy not converted to light becomes heat, which must be dissipated.
- Pump conversion efficiency measures how much input pump energy is converted into laser emission. It can be very high for diode-pumped fiber lasers.
- Some lasers are limited to pulsed operation by internal physics or by excitation techniques.
- Laser pulses have a characteristic rise time, pulse duration, and fall time.
- Pulse duration normally is measured as full width at half maximum of power.
- The repetition rate is the number of pulses per second.
- Energy is the power received during an interval.
- Average power is the total energy delivered divided by the time.
- The uncertainty principle limits how precisely we can measure wavelength and pulse duration. Very short pulses have very wide bandwidths; very narrow bandwidths require continuous emission over a long time.
- Polarization is the alignment of the electric field in light waves.
- A Brewster window can polarize laser output.

WHAT'S NEXT?

In Chapter 5, we will learn about major laser accessories and how they are used.

QUIZ FOR CHAPTER 4

1. A semiconductor laser emits 820-nm light with a bandwidth of 2 nm. What is its coherence length?
 - a. 822 nm
 - b. 168 μm
 - c. 8.4 mm
 - d. 840 mm
 - e. 12.8 m
2. A single-frequency laser has bandwidth of 100 kilohertz (10^5 Hz) and wavelength of 600 nm. What is its coherence length?
 - a. 168 μm
 - b. 8.4 mm

- c. 840 mm
 - d. 12.8 m
 - e. 1.5 km
3. Which of the following contributes to the broadening of laser emission bandwidth?
- a. Doppler shift of moving atoms and molecules
 - b. Amplification within the laser medium
 - c. Coherence of the laser light
 - d. Optical pumping of the laser transition
 - e. None of the above.
4. How many longitudinal modes can fall within a laser's gain bandwidth?
- a. 1 only
 - b. 2
 - c. 3
 - d. 10
 - e. No fixed limit; depends on bandwidth and mode spacing
5. The Rayleigh range is the near-field distance over which light from a laser does not spread significantly. What is this distance for a helium–neon laser with 1-mm beam diameter and 632.8-nm wavelength?
- a. No such distance; beam starts diverging immediately
 - b. 0.23 m
 - c. 1.58 m
 - d. 6 m
 - e. 1.5 km
6. A helium–neon laser beam has divergence of one milliradian. What would its diameter be at the 38,000-km altitude of geosynchronous orbit?
- a. 1 m
 - b. 38 m
 - c. 76 m
 - d. 1 km
 - e. 76 km
7. If we direct a helium–neon laser of problem 6 through a one-meter telescope to limit its divergence, what is its beam diameter at the 38,000-km altitude of geosynchronous orbit? (Assume that the telescope is diffraction limited to a divergence of λ/D .)
- a. 24 m
 - b. 48 m

- c. 76 m
 - d. 94 m
 - e. 3.8 km
8. A semiconductor diode laser has output wavelength of 650 nm and beam divergence of 20 degrees without external optics. If you pointed it at a wall 2 m away, how big would the laser spot be?
- a. 1 mm
 - b. 14.6 mm
 - c. 14.6 cm
 - d. 1.46 m
 - e. 2 m
9. You want to focus the red diode laser in Problem 8 onto a 1-mm spot on a wall 2 meters away. What diameter of a focusing lens would you need?
- a. 650 μm
 - b. 1 mm
 - c. 1.3 mm
 - d. 2.6 mm
 - e. 650 mm
10. What is the maximum wall-plug efficiency for a laser with a power supply that is 70% efficient with an energy-deposition efficiency of 25%, in which the laser transition represents 60% of the excitation energy and half the excited atoms emit laser light?
- a. 1%
 - b. 5.25%
 - c. 10%
 - d. 25%
 - e. None of the above
11. A diode laser emitting at 980 nm pumps an ytterbium-doped fiber laser emitting at 1080 nm. Assuming all the pump light is absorbed, all the emitted light is collected, and everything else goes perfectly, what is the maximum possible conversion efficiency?
- a. 10%
 - b. 50%
 - c. 89.6%
 - d. 90.7%
 - e. 98.0%

12. A pulsed laser generates 500-kilowatt pulses with full width at half maximum of 10 nanoseconds at a repetition rate of 200 hertz. Making simplifying assumptions about pulse power, what is the average power?
- a. 1 watt
 - b. 2 watts
 - c. 5 watts
 - d. 10 watts
 - e. 20 watts

CHAPTER 5

OPTICS, LASER ACCESSORIES, AND MEASUREMENTS

ABOUT THIS CHAPTER

Lasers are used with other equipment which laser users need to understand. This chapter starts with an overview of optical components, materials, and coatings which are vital for laser applications. Then we describe optical systems used for beam delivery, and the mechanical systems used to mount and position lasers and optics. The final section introduces specialized measurement units, terminology, and equipment used with optics and lasers, which are also important in understanding lasers.

5.1 CLASSICAL OPTICAL DEVICES

In Chapter 2, you learned the basics of how lenses and mirrors refract and reflect light. That is the foundation of classical optics, static devices used to manipulate light. This section focuses on their use with lasers.

We will pick up where we left off in Chapter 2. Our descriptions largely consider laser beams as bundles of parallel light rays traveling in straight lines as if they were emitted by an object an infinite distance away.

5.1.1 Simple Lenses

Simple lenses are a single piece of glass or other transparent material, with one or both sides curved to focus light. A single positive lens focuses parallel light rays in a laser beam to a spot at its focal point, as you saw in Figure 2-11. The minimum spot size for a lens with focal length f and diameter D at a laser wavelength λ is the same as the minimum spot size for a laser beam (Equation 4-9):

$$\text{Spot size} = \frac{f\lambda}{D} \quad (5-1)$$

Lenses focus light rays the same way, no matter which direction they travel, so the same optics that focus parallel light rays to a point can also collimate a diverging beam of light to form a bundle of parallel rays.

5.1.2 Compound Lenses

Compound lenses are assembled from two or more simple single-element lenses. Full-size camera and projector lenses are good examples. Compound lenses are more complex and expensive than single-element lenses, but they do a better job of focusing light because careful design can compensate for the inherent imperfections of simple lenses over a wide field of view, long distances, or a large depth of focus. Some compound lenses have fixed elements, but others, such as photographic zoom lenses, have elements that move back and forth to bring light into focus. Compound lenses are more important for visual optics that require focusing all wavelengths to the same point; laser systems using only a single wavelength can often use simple lenses.

5.1.3 Collimators, Beam Expanders, and Telescopes

Though it may seem amazing today, nearly two millennia passed between the discovery that lenses or mirrors could focus light and the discovery of the telescope. The ancient Greeks knew about burning glasses, but the telescope was invented about 1600 AD. It seems unlikely that either discovery was the outcome of a systematic investigation. Legend says that children playing with lenses in the shop of Dutch spectacle maker Hans Lippershey discovered




Figure 5-1. A refracting telescope serves as a beam expander or collimator for a laser beam.

the Galilean telescope, which is made of a large positive lens with a smaller negative lens that serves as the eyepiece. Replacing the negative lens with a positive lens makes a more powerful telescope, but turns the image upside down.

Telescopes do more than make distant objects look closer. As shown in Figure 5-1, a simple refracting telescope made from a pair of positive lenses can also expand the bundle of parallel light rays in a laser beam, enlarging the beam diameter. A telescope that serves this purpose is called a *beam expander* or *collimator*. If you reverse the direction of the light, the collimator optics can shrink the laser beam.

5.1.4 Optical Aberrations

Simple optics do not bring light to a perfect focus because they suffer from *aberrations* that blur light even if the lenses are polished perfectly to the ideal shape specified by the optical designer. These aberrations arise from limitations inherent in the nature and design of simple lenses.

One important type is *chromatic aberration*, which occurs because the refractive index of transparent materials depends on the wavelength. In Chapter 2, you saw that the focal length of a positive lens with radii of curvature of its surfaces R_1 and R_2 has a focal length f that also depends on refractive index of the lens material n :

$$f = \frac{1}{(n - 1) \left[\frac{1}{R_1} + \frac{1}{R_2} \right]} \quad (5-2)$$

Refractive index does not change much with wavelength, but it changes enough to affect the focusing of visible light. In one

common type of optical glass, n is 1.533 at 434 nm in the violet and 1.517 at 656 nm in the red. This means that the focal length in the red is 1.03 times that in the violet, or 3% longer. The effect is most obvious near the edge of a lens or optical system, where it may give white objects a red fringe on one side and a blue fringe on the other. It is evident in cheap binoculars or in strong plastic corrective lenses. This effect is also called *dispersion*, because it spreads out light according to wavelength; it produces the rainbow.

Chromatic aberration is not a big problem with lasers that emit a single wavelength. However, compensation may be needed if the laser emits multiple wavelengths, or if the same optical system is used by the human eye or other lasers. The simplest way to compensate for chromatic aberration is to glue together lenses made of glasses with refractive indexes that change in different ways with wavelength. Compound lenses designed for this purpose are called *achromatic lenses*. They can greatly reduce chromatic aberration, but the correction is not perfect.

Most simple lenses have surfaces shaped like parts of a perfect sphere. That shape is easy to manufacture precisely, but it suffers from an effect called *spherical aberration*, because light passing through the outer part of the lens is focused to a focal length slightly shorter than light passing through the central zone. Using only the central part of the lens or using only lenses with long focal length lessens the impact of spherical aberration. Spherical aberration can also be controlled by fabricating lenses with *aspheric* (nonspherical) shapes that focus all light rays in the system to the desired distance. Aspheric lenses require careful design and are hard to produce by conventional optical polishing of glass, but they can be molded from plastics at reasonable costs. The progressive lenses that are an alternative to bifocals for people with poor visual depth accommodation are aspheric lenses.

5.1.5 Cylindrical Optics

Most lenses are radially symmetric, with spherically curved surfaces, so they refract the same way if you turn them around their centers without tilting them. Light reaching such a lens at the same angle and the same distance from its center is always bent in the same way.

Spherical lenses are fine for most purposes. However, many semiconductor diode lasers emit light from narrow stripes on their




Figure 5-2. A cylindrical lens focuses light in one direction, but not in the perpendicular direction.

edges, causing their beams to diverge more strongly in one direction than in the other, as you will learn in Section 10.8.1. Those beams can be made symmetrical by passing them through a *cylindrical lens* that bends light in only one direction, as shown in Figure 5-2. The name comes from the fact that the surfaces are slices of a cylinder, not a sphere, and, as shown in the figure, it focuses light vertically but not side to side. Passing light from such semiconductor diode lasers through a cylindrical lens evens out the beam divergence, but an additional spherical lens must be added to form a pencil-like beam.

5.1.6 Mirrors and Retroreflectors

Mirrors can serve two functions in laser systems. As you learned in Chapter 2, curved mirrors can focus light like simple lenses. Flat mirrors also redirect light, either reflecting light at the ends of a laser cavity or redirecting a laser beam outside the cavity.

Laser mirrors differ in one crucial aspect from household mirrors. The reflective layer in a household mirror is covered by a layer of glass to protect it from the accumulation of dust and dirt, but that limits its optical quality. Lasers and other precision optics use *front-surface mirrors* in which the reflective layer is on the front surface, which can expose the reflective surface to damage but avoids double reflection when a rear-surface bathroom mirror reflects light both from its front and its back surfaces.

Although many laser applications use standard flat or curved mirrors, some use a special type of mirror called a *retroreflector* or *corner cube*. This consists of three flat mirrors at right angles to




Figure 5-3. A retroreflector returns light precisely back in the direction from which it came.

one another, like the corners of a cube, as shown in Figure 5-3. Some retroreflectors are “hollow,” made up of three front-surface mirrors. Others are prisms with rear surfaces that are totally reflective because the light strikes the boundary with a low-index material at a steep angle.

What is special about retroreflectors? A flat mirror reflects an incident beam at an angle equal to the angle of reflection, so it goes off in another direction unless the laser beam is perfectly perpendicular to the beam. A retroreflector bounces light directly back to the source, as shown in Figure 5-3. Apollo 11 astronauts Neil Armstrong and Buzz Aldrin put an array of retroreflectors on the lunar surface to return strong laser signals to earth, enabling the first highly accurate measurements of the distance to the Moon. The retroreflectors remain in use for precision measurements.

5.1.7 Dispersive Optics

Many laser applications require spreading out or dispersing light according to its wavelength. The two most important types of




Figure 5-4. A glass prism forms a spectrum by bending different wavelengths at different angles.

dispersive optics, *prisms* and *diffraction gratings*, rely on different principles to do similar tasks.

A prism is a slab of glass with flat polished faces. Prisms come in many shapes, and have many uses that involve only redirecting light, as in the prisms used to fold the path of light inside binoculars to show the image right side up.

Dispersive prisms are triangular in cross-section, often equilateral, and their operation depends on how the refractive index of the prism varies with wavelength. Light entering the prism is refracted at an angle that depends on the prism's refractive index. Glass has a higher index at short blue wavelengths than at longer red wavelengths, so the blue wavelengths are bent the most, as shown in Figure 5-4. For a prism to spread out a spectrum, its output face must be at an angle to the input face. If the two faces were parallel, refraction at the exit face would cancel refraction at the entrance, and all wavelengths would emerge at the same angle, although the light waves of different colors would be separated in space.

A diffraction grating is a flat piece of glass or plastic with a series of parallel grooves in its surface. The grooves diffract (or scatter) light at all wavelengths from the surface across a wide range of angles. The diffracted light waves interfere with each other, so the light of one wavelength adds constructively at only one angle but adds destructively at other angles, spreading out a rainbow

spectrum like a prism. You do not need a special diffraction grating to see the effect; the tightly wound spiral of data bits on a CD, CD-ROM, or DVD acts in the same way, diffracting light and spreading out a rainbow, which you can see if you hold the disk in the light and look for its reflection. The width of the spectrum depends on the spacing of the lines for both a standard diffraction grating and an optical disk. Look carefully and you can see that a grating also produces fainter secondary order rainbows.

Some important optical instruments contain diffraction gratings or prisms to spread out a spectrum that separates light by its wavelength. The *spectroscope* is an instrument that spreads out the spectrum of light from an external source, such as a laser or fluorescent lamp.

A more elaborate instrument called the *monochromator* spreads out the spectrum from an internal light source, then directs the light through a narrow slit that transmits a narrow slice of the spectrum to the outside world. The monochromator was invented before the laser, and its output is often called “monochromatic,” although most lasers emit a narrower range of wavelengths.

5.1.8 Polarization and Polarizing Optics

As you learned in Chapter 4, light can be divided into two orthogonal polarizations, with their electric fields perpendicular to each other. In the simplest case, called *linear polarization*, the electric field is oriented in a fixed direction. As shown in Figure 5-5, this polarization can be considered as a vector going from the origin (0, 0) to a point in the X–Y plane. You can break this vector down into horizontal and vertical components, as shown, or you can turn the coordinates so one direction matches the direction of polarization.

Light can be polarized in different ways. Light is *linearly polarized* if its electric field remains continually in the same plane. *Circular polarization* occurs when the plane of polarization rotates at a constant speed with constant amplitude, so the polarization vector draws a circle in the X–Y plane. If the polarization vector changes both in direction and amplitude, it is *elliptically polarized* because it draws an ellipse in the X–Y plane. Light is called *unpolarized* if it is the sum of light of all different polarizations. (*Random polarization* is different, meaning that the direction of polarization changes randomly in time.) Laser light may be polarized or unpolarized.

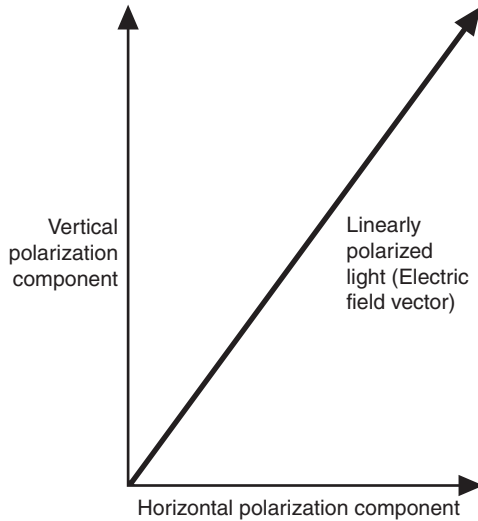


Figure 5-5. Direction of polarization as a vector.

A number of optical devices can manipulate light polarization.

Polarizers select or separate light of particular polarizations. Simple dichroic film polarizers (like those used for polarizing sunglasses) absorb light polarized along one axis and transmit light linearly polarized along the other axis. For example, a polarizer might transmit light polarized along the Y axis in Figure 5-5 and absorb light polarized along the X axis. Higher-performance polarizers normally are made from *birefringent* materials such as calcite, which separate light of different polarizations so it goes in different directions instead of one being absorbed.

Birefringent materials have refractive indexes that depend on polarization, so they can also retard or rotate polarization. Polarization retardation occurs because light of different polarizations has different refractive indexes, so the slower polarization falls behind the faster one.

How much the polarization is retarded depends on the thickness of the birefringent material. Typically, polarization retarders separate the two polarizations by one-quarter or one-half of a wavelength. Such retarders are called *quarter-wave* or *half-wave plates*. One application is to rotate the phase of a linearly polarized input beam. If the input polarization is at an angle θ to the principal plane of the retarder, the polarization of the output beam is rotated by 2θ . Another application is to convert an input linearly polarized beam




Figure 5-6. A beam splitter transmits part of the input beam and reflects the rest.

to circular polarization by passing it through a quarter-wave plate with its axis 45° from the input beam polarization.

5.1.9 Beam Splitters

A key optical component used with lasers is the *beam splitter*. It does what its name says, splitting a beam into two parts. Normally, the beam splitter is at an angle to the incident beam, so it transmits some light and reflects the remaining light at an angle, as shown in Figure 5-6. (You need not align the beam splitter so the reflected beam is at a right angle to the incident and transmitted beam, but it is common and is a convenient way to draw the beams.)

Beam splitters come in many types. Some reflect one linear polarization and transmit the orthogonal polarization. Others are insensitive to polarization. The reflectivity and transmission of some are very sensitive to wavelength; others vary little with wavelength. Two important things to note are that the beams do not have to be reflected and transmitted at right angles, and that the light is not always divided equally between transmitted and reflected beams.

5.2 OPTICAL MATERIALS

Section 2.3 described the properties of optical materials, but so far we have given few details about the materials themselves. Now it is time to look more closely at optical materials and the crucial roles




Figure 5-7. Optical transmission of a 2-mm sheet of borosilicate glass. (Courtesy of Thorlabs).

they play in laser optics. This section focuses on *passive* optical materials, which reflect, refract, transmit, or absorb light but do not otherwise alter it.

How materials interact with light depends on their internal composition and the wavelength of light. You learned in Section 2.3.6 that the refractive index changes with wavelength, with Table 2-6 showing the effect in silicon dioxide glass. In fact, the reflection, transmission, scattering, and absorption in optical materials also vary with wavelength, often in complex ways, depending on the composition of the material. We do not notice this much in everyday life because the glasses, crystals, and plastics we use as windows and lenses are picked because they transmit visible light. However, most common window and optical glasses absorb strongly at wavelengths shorter than about 300 nm in the ultraviolet and longer than about 3 μm in the infrared, as shown in Figure 5-7.

What this means in practice is that a wide variety of optical materials are used in different parts of the spectrum. Figure 5-8 shows a sampling of materials used in transmissive laser optics and the wavelength ranges where they transmit light best. All optical properties vary with wavelength, including refractive index and reflectivity as well as transmission. Let us look briefly at the types of materials used for optics.

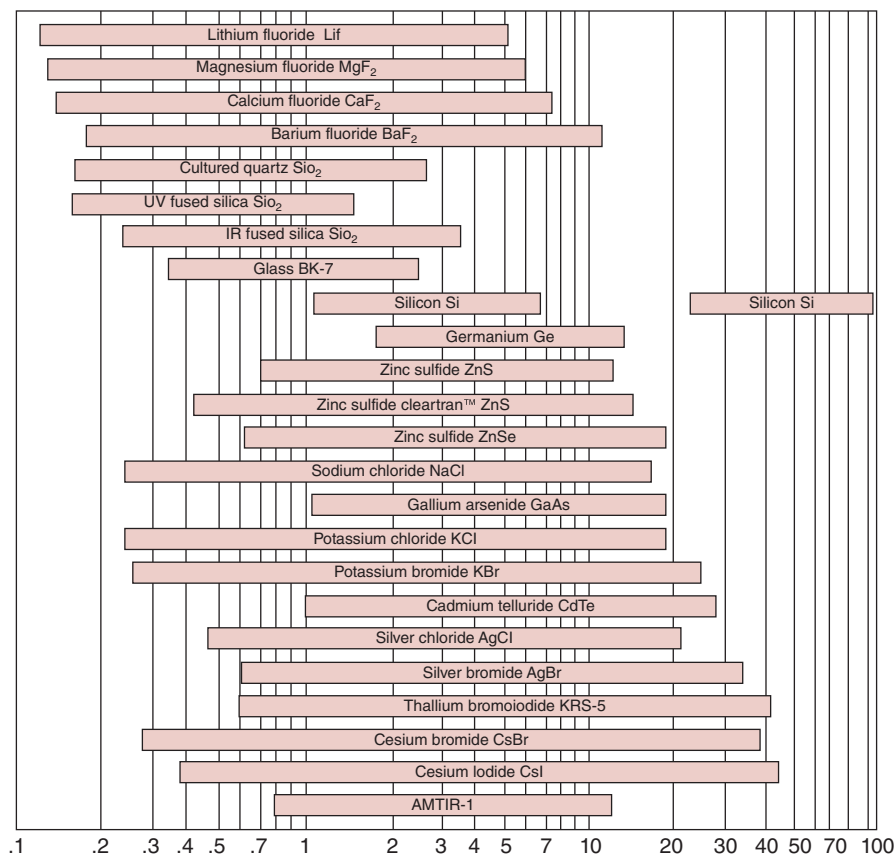


Figure 5-8. Important optical materials and their transmission ranges. (Courtesy of Janos Technology Inc.)

5.2.1 Transparent Optical Materials

The choice of materials for transparent optics depends on their mechanical and chemical properties as well as their optical properties. The durability of the glasses used for visual optics spoils us. Many materials transparent at other wavelengths cannot withstand normal use without sagging, scratching, wearing away, or dissolving into puddles on a humid day. Sodium chloride, which we use as table salt, is transparent to infrared light, and has been used in infrared laser optics, but it has to be isolated from moisture. Other materials may be too expensive, or too hard or soft to shape into lenses.

Transparency inherently depends on the chemical composition of a material, and it differs widely among materials. The high

transparency of silica-based glasses at visible wavelengths combines with their durability, ease of fabrication, and low cost to make them the most widely used optical materials. However, such materials are essentially opaque at infrared wavelengths longer than 3 μm , so different materials are needed for that part of the spectrum. Figure 5-8 shows the wavelengths where important commercial materials are transparent enough for optical applications.

Transparent optical materials fall into four broad categories, based on their internal structure: glasses, crystals, ceramics, and polymers. Where in the spectrum they are transparent depends more on the chemical composition than on their internal structures.

What the word *glass* means depends on the context. To a physicist specializing in materials science, a glass is a noncrystalline solid—that is, a material in which the atoms or molecules are arranged randomly, as in a liquid, but frozen in place. Around the house, and often in optics, glass means the stuff of windows and drinking vessels, a transparent solid made by melting sand with other oxide compounds. Optical glasses usually contain a large fraction of silicon dioxide mixed with other oxide compounds to give the glasses desirable values of refractive index and a very uniform optical quality. Some glass types are identified by important components, such as lead, phosphorous, or boron. Some optical glasses are made from *quartz*, a natural crystalline form of silica. Optical fibers are made from an extremely pure form of silica known as *fused silica*. These materials are also called *oxide glasses* because chemically they are made from oxide compounds. Oxide glasses usually are transparent from around 300 to 3000 nm, but light-absorbing impurities may restrict that range.

Some glasses not made from oxide compounds are transparent at infrared wavelength. *Fluoride glasses* based on fluoride compounds are transparent at wavelengths longer than about 3 μm in the infrared. Another family of glasses transparent in the infrared are *chalcogenide glasses*, containing compounds of sulfur, selenium, and tellurium, such as zinc sulfide. Many of these materials are not as rugged as oxide glasses, but transmit light in parts of the spectrum where oxide glasses are not available.

Crystals are materials in which atoms are arranged in periodic structures with regular spacing between atoms. Some materials, such as silicon and germanium, are opaque at visible wavelengths but transparent in the infrared. Many of the materials listed in Figure 5-8 are crystals. Single crystals in which all the atoms are

lined up in the same structure are used in lasers and for some optics because their high optical quality.

Ceramics are polycrystalline materials, often made from ground-up crystals. Ceramics in your kitchen are not transparent at visible wavelengths, but with suitable processing ceramics can be transparent at some wavelengths, and may be used as optics or laser materials.

Polymers are organic compounds often called plastics, some of which can be made transparent and used as optics. They are cheaper and lighter than glass, can be molded to aspheric and other shapes, and now dominate the market for spectacles. They are best developed for visible wavelengths.

5.2.2 Reflective Materials

Reflection is a surface effect, so it depends on the surface exposed to light and the environment. Some materials are inherently reflective at optical wavelengths, but other reflective surfaces are just thin films deposited on nonreflective materials such as glass.

Most metals naturally reflect strongly at visible and infrared wavelengths, but absorb light at shorter wavelengths. The three most reflective metals in the visible are aluminum, silver, and gold, as shown in Figure 5-9. Silver has the highest reflectivity in the visible, but absorbs strongly at ultraviolet wavelengths shorter than about 320 nm, and tarnishes when exposed to air. Gold reflects red and orange light strongly, but absorb blue light, giving gold its distinctive color. Aluminum reflects somewhat less light at visible wavelengths, but remains reflective deeper into the ultraviolet. These reflective metals often are coated onto smoothly polished glass surfaces.

Highly reflective surfaces can also be made by depositing a series of alternating thin layers of two materials with different refractive indexes, as described in Section 5.3.

5.2.3 Bulk and Micro- or Nanostructured Materials

So far we have talked about solid materials that are uniform in composition and easily visible or handled. These are called *bulk* materials, and they have long been the standard materials for lenses and optical components.

New technology has created new classes of materials that are likely to play increasingly important roles in laser optics. We




Figure 5-9. Reflectivity of aluminum, silver, and gold in the ultraviolet, visible, and near infrared. (Courtesy Dr. Bob at the English language Wikipedia, Creative Commons License.)

cannot go into detail here, but these are the emerging trends to watch for.

- *Micro-optics*: tiny components on a millimeter to micrometer scale. Physically, these work like bulk optics, but their small size makes them difficult to handle.
- *Nano-optics*: nanometer-scale devices, which because of their scale operate differently than bulk optics.
- *Micro- or nanostructured devices*, including photonic crystal devices, with internal structures that give them novel properties. These may be made of a single material or two or more different materials.
- *Metamaterials*: nanostructured devices made of elements smaller than the wavelength of light, sometimes containing two or more different materials, which can be designed to have properties not available from ordinary materials.

5.3 OPTICAL COATINGS AND FILTERS

Reflection and refraction both occur at optical surfaces, so coating the surface of an optical device can change its optical properties in useful ways. Coatings play a very important role in laser optics.

5.3.1 Metal Mirror Coatings

A thin layer of reflective metal can convert a surface into a mirror. Bathroom mirrors typically are sheet glass coated on their backs with metal, so the glass protects the exposed surface, but this causes undesirable effects, including “ghost images” when viewed from an angle. Optical mirrors typically are coated on the front surface for higher performance, but must be handled carefully. A transparent coating can be applied to cover the metal surface and protect it from damage.

Although metal mirrors look highly reflective, the metal surface typically absorbs some light, as shown in Figure 5-9. Aluminum is widely used because of its low cost and high reflectivity in the visible. Silver is more reflective, but tarnishes easily. Gold is highly reflective at longer visible wavelengths. The metal coatings may need additional coatings to enhance performance, particularly in the ultraviolet.

5.3.2 Antireflection Coatings

Some reflection is inevitable any time light passes from air into glass or, in general, from a low-index transparent material into a higher-index transparent material. The amount of reflectance depends on the refractive indexes of the materials, the polarization of light, and the angle. If the light is incident perpendicular to the surface, the reflectance is

$$\text{Reflectance} = \frac{\left(\frac{n_{\text{high}}}{n_{\text{low}}} - 1\right)^2}{\left(\frac{n_{\text{high}}}{n_{\text{low}}} + 1\right)^2} \quad (5-3)$$

where n_{high} and n_{low} are the refractive indexes of the high- and low-index materials, respectively. For light in air, this yields a reflection loss of about 4%.

At different angles, the reflectance also depends on the polarization direction and the angle, but the overall dependence on refractive index difference remains. That reflective loss can be reduced by coating the glass with a thin layer intermediate in refractive index between the air and the glass. This creates two interfaces: between the air and the coating, and between the

coating and the glass. However, the differences are smaller if the coating has an index of 1.25, halfway between glass and air. In that case, loss is 1.23% at the air–coating interface and 0.83% at the coating–glass interface. Total loss is about 2%, calculated by multiplying the fractions of light transmitted through the surfaces, not by adding reflective losses.

5.3.3 Multilayer Interference Coatings

More complex effects are possible when light is directed through a series of many thin alternating layers of two materials which act as what are variously called *multilayer*, *interference*, and/or *dichroic* coatings. The two materials have different refractive indexes, so as the light passes from one layer to the next, a fraction of the light is reflected upward and the rest transmitted to the next layer. The more layers, the more light is scattered. Figure 5-10 shows how a




Figure 5-10. Reflections at the boundaries between two different materials scatters some light upward and transmits some downward at each junction. The more layers in the structure, the stronger the interference selecting wavelengths emerge exactly in phase with other light in the layers, and thus the narrower bandwidth reflected up or transmitted down, depending on the design.

little light is scattered upward at each interface between the two different layers, as other light is transmitted downward.

Careful selection of the layer thicknesses can make the scattered light from successive layers add either constructively or destructively, depending on the wavelength. If, for example, the reflected light waves add constructively at a wavelength of 500 nm, a stack of layers can become a highly selective mirror that reflects 500-nm light very strongly but transmits other wavelengths. This effect can be used as a filter to selectively reflect a narrow range of wavelengths, while transmitting other wavelengths. Such highly selective filters are very useful in laser systems because they can either transmit or block the narrow range of wavelengths in a laser beam.

Proper selection of the layer materials and thicknesses lets coating designers meet demanding specifications. For example, they can make coatings that reflect or transmit only a narrow range of wavelengths. Such coatings can transmit (or block) a narrow range of wavelengths at a laser line. Thus, they can block white light from overhead lights while transmitting the red line of the laser in the laser scanners used for automated checkout in supermarkets.

Multilayer coatings can give very high peak reflection or transmission at a particular wavelength, but this makes them very sensitive to wavelength. They are also very sensitive to the angle of incidence, because changing the angle changes the thickness of material the light has to pass through, and thus changes the wavelength at which interference effects occur. Tilt an interference coating that strongly reflects one color, and its apparent color will change with the viewing angle.

Coatings deposited on flat glass are used as filters to select particular ranges of wavelengths. Generally, most light is transmitted or reflected rather than absorbed. Table 5-1 lists some examples.

5.3.4 Color Filters

Wavelengths can also be selected by passing them through *color filters* made of colored transparent glasses or plastics. Such color filters are often used for photography. Unlike thin-film filters, color filters absorb the light they do not transmit, which is fine for photography, but can cause heating if they are used to attenuate high-power laser beams.

Table 5-1. Some types of multilayer coatings

Mirror coatings	Strongly reflect selected wavelengths; used in laser cavities
Narrow-band filters	Transmit a narrow band of wavelengths, e.g., a laser line
Wide-band filters	Transmit a wide band of wavelengths, e.g., red light
Band-pass filters	Transmit a range of wavelengths: narrow, medium, or wide
Band-rejection filters	Blocks a band of wavelengths, e.g., a laser line

5.3.5 Neutral-Density Filters

Neutral-density filters reduce or attenuate light intensity uniformly at all wavelengths covered. Metallic neutral-density filters are coatings that attenuate light by reflecting the unwanted portion. Other neutral-density filters absorb the unwanted portion. Because their absorption is uniform across the visible spectrum, they look gray to the eye.

Attenuation of a neutral-density filter is normally measured as *optical density*, defined as

$$\text{Optical density} = -\log_{10} \left(\frac{\text{Output}}{\text{Input}} \right) \quad (5-4)$$

Thus, an optical filter that transmits 1% (0.01) of incident light has an optical density of 2. Reversing the formula gives the fraction of transmitted light:

$$\text{Transmission} = 10^{(-\text{optical density})} \quad (5-5)$$

5.3.6 Spatial Filters

A spatial filter transmits one part of a laser beam and blocks the rest. A common type of spatial filter is a pinhole in a sheet of black-painted metal. It transmits the center of the laser beam, but not the outer portions. Often, spatial filters are used to pick out the central, most uniform, part of a laser beam. Slits can also serve as spatial filters.

5.4 BEAM DELIVERY, DIRECTION, AND PROPAGATION

Chapter 4 described how laser beams diverge and propagate through free space. Many laser applications require more beam

direction than that to do their job. This section will look at four different approaches for delivering laser beams to desired locations, beam scanners, articulated arms, fiber optics, and optical waveguides.

5.4.1 Beam Scanners

Laser beam scanners uses simple mechanical or optical movement to sweep a reflected beam across a field of view. Applications include construction alignment systems, laser light shows, and automated checkout systems.

One common approach is to spin a polygonal mirror around its central axis. Each reflective face on the spinning mirror sweeps the laser beam across a range of angles as the mirror turns past the fixed beam, following the rule that the angle of incidence equals the angle of reflection. When a new facet turns toward the laser, it repeats that scan. More complex scanning patterns are possible if the simple mirror is replaced by a *holographic optical element*, an optical element fabricated by holography and designed to scatter a beam of light so it scans patterns otherwise hard to produce.

Resonant and galvanometer scanners also rely on moving mirrors, but they twist flat mirrors back and forth across a limited angle, and the twisting mirror sweeps the laser beam across a line.

An alternative approach that does not require mechanical motion is the solid-state *acousto-optic beam deflector*. Passing sound waves through a suitable transparent material changes its refractive index, deflecting a laser beam passing through the material. The angle over which the beam is deflected is proportional to the ratio of the acoustic wavelength to the light wavelength. The amount of beam deflection depends on the acoustic power. The absence of moving parts can be an advantage for some applications.

5.4.2 Articulated Mechanical Arms

Laser machining requires moving a laser beam across the surface of objects. When it is not convenient to move the laser or the object, the beam can be redirected by using an articulated arm that looks like the arm on a dental drill. The beam passes through a series of tubes that are connected mechanically so that lenses or mirrors can guide the laser beam through them to the object. Typically, the arm

is robotic or robot-controlled, so it can perform the same operation repeatedly.

5.4.3 Fiber-optic Beam Delivery

Fiber-optic beam delivery is more attractive for most laser machining applications than articulated arms. Fibers are light, thin and flexible, making them much simpler and easier to move and much less expensive than articulated mechanical arms. Section 2.5 described the various types of optical fibers and how they guide light.

For beam delivery, the laser output is coupled into one end of an optical fiber. Usually the transmitting fiber is a step-index fiber with pure silica core 50 to 1000 μm in diameter, large enough to handle high laser powers without damage to the fiber ends or interior. Losses are not as low as telecommunication fibers, but the distances are meters rather than kilometers. The beam profile delivered is important for machining applications. For many, the preferred shape is a *top-hat profile*, with uniform intensity across the whole core, dropping off steeply at the cladding boundary.

Beam delivery fibers usually are moved by precision robotic arms, which are much simpler and less expensive than the articulated optical arms described in Section 5.4.2. The high attenuation of silica-based fibers at wavelengths outside their primary range of 350 nm to 2 μm limits their use in the short ultraviolet or mid- to far-infrared. Special ultraviolet grades of fused silica or quartz can be drawn into fibers that transmit wavelengths to about 200 nm. Chalcogenide and fluoride glasses listed in Figure 5-8 can be drawn into fibers for longer wavelengths. Another alternative is the hollow optical waveguides described in the next section.

5.4.4 Optical Waveguides and Integrated Optics

Waveguides are devices for guiding electromagnetic or sound waves. If your background is in electronics, you may have seen hollow metal waveguides that transmit microwaves. Optical fibers are another type of waveguide, which guide light through transparent solids, but they are not available at all wavelengths.

Hollow optical waveguides have been developed for carbon-dioxide lasers, which emit at 10 μm in the infrared, where few good

fiber materials are available. Typically, the hollow waveguides are made of silica or sometimes from sapphire (aluminum oxide). The beam travels through air in the core, so infrared absorption by the waveguide material has little impact on attenuation. The hollow core may be coated with other materials to enhance reflectivity and reduce losses. For example, hollow silica waveguides with cores from 300 to 1000 μm are coated on the inside with silver halides that reflect between 2.9 and 20 μm for use with medical CO_2 lasers at 10 μm . They can also be used with erbium-doped solid-state lasers emitting at 2.94 μm .

Solid optical waveguides can also be made. The most important are *planar optical waveguides* fabricated from thin layers of material deposited on flat substrates. They have a rectangular cross-section and confine light by total internal reflection, because they are surrounded partly by air and partly by other material with lower refractive index than the waveguide. Planar waveguides are used in some semiconductor lasers and in integrated optics used in processing optical signals for telecommunications and other applications.

5.5 MOUNTING AND POSITIONING EQUIPMENT

Walk into any well-equipped laser laboratory, and the first thing you are likely to notice is a massive table, usually painted flat black, with screw holes regularly spaced across its metal surface. This behemoth is called an *optical table* and is shown in Figure 5-11. Mounted on it you will find an array of special mounting equipment holding lasers, lenses, and other optical components.

Optical tables and mounts are not optics per se, but they are used with optics and lasers. Optical tables and massive linear rails called optical benches serve as foundations. They often are mounted on shock-absorbing legs that isolate the optics from local vibrations that can shake precision optics and holography. Vibrations that move optics only a fraction of a wavelength can disrupt sensitive optical measurements, and people moving in a room, elevators in a building, or trucks on a nearby street can cause such movement.

Optical mounts are designed fit into holes drilled in optical tables so they can hold lasers and optical components firmly in place. Many mounts also allow precision movement of components. Optical mounts and positioning equipment sound like mundane




Figure 5-11. An optical table on vibration-isolation legs being set up a laser laboratory at the University of California at San Diego (Courtesy of Newport Corp.)

technology compared with lasers, but they are important to the success of precision measurements and experiments.

5.6 NONLINEAR OPTICS

So far, we have described only simple optical devices that are passive or only have a linear effect on light passing through them. Linear in this context means that the output is directly proportional to the input:

$$\text{Output} = A \times (\text{Input}) \quad (5-6)$$

This means that light emerging from a linear system may be fainter or brighter than the input light, but is otherwise unchanged, still in the same proportions and still the same color or wavelength.

It is the optical equivalent of listening to music played on a good sound system—the music is the same but it is fainter or louder, not distorted.

Other types of devices are based on *nonlinear optics*, in which the output is a more complex function of the input:

$$\text{Output} = A(\text{Input}) + B(\text{Input})^2 + C(\text{Input})^3 + \dots \quad (5-7)$$

What this equation means is that the interaction of light with the nonlinear material converts some of the input light energy to other wavelengths, although from a mathematical viewpoint the noise actually is at harmonics (multiples) of the light frequency. It is like listening to music played too loud on a poor sound system, so part of the sound in a pure tone could wind up at double the frequency, and part at triple the frequency. Even more complex combinations are possible. For light, it means that infrared light at 1064 nm from a neodymium solid-state laser is halved in wavelength to 532 nm in the green (equivalent to doubling the frequency). That is what happens inside a green laser pointer.

The theory of nonlinear optics relies on the complex mathematics of differential equations three pages long and PhD theses, so we will not go into much detail. What is important to know at this stage in your learning about lasers is what nonlinear optics can do and how they can be useful. The subsections that follow give a basic introduction to important processes including: generating harmonics of laser light, generating sum and difference frequencies, and parametric processes that can produce tunable frequencies.

Like distortion in a sound system occurs at high audio power, nonlinear optical effects occur at high light intensities. Laser light is particularly efficient at producing nonlinear effects because it can be focused to very high intensity, and because it includes a narrow range of wavelengths. Physicists had long known that it should be possible to generate harmonics of light, as it was for radio waves. But it took the laser to actually generate the second harmonic of light, turning some of the bright red pulse from a ruby laser into a faint ultraviolet pulse recorded in a photo. Since then, researchers have generated third, fourth, fifth, and much higher harmonics from laser light.

Nonlinear optical effects increase in strength more rapidly than linear effects because they depend on the square or some higher power of the input. Double the input power, and the output of the

second harmonic increases by a factor of four, the square of the input increase. Raise the input a factor of five, and the output jumps by a factor of 25. So once you start producing harmonics, turning up the power converts even more of the input light into the harmonic.

Nonlinear optical devices are mainly used to alter the wavelength and other properties of laser beams. Thus, they are covered in Section 6.5, part of the chapter on types of lasers, because they effectively create different types of lasers.

5.6.1 Harmonic Generation

The simplest nonlinear effect to visualize is the second-harmonic generation, which doubles the frequency of a wave. If you are familiar with music, you think of harmonics as overtones, a higher tone with a frequency a multiple of that of a lower or fundamental tone. The first overtone is double the frequency of the original tone.

Optical harmonics work the same way. The second harmonic is at twice the frequency of the fundamental wave or, equivalently, at half its wavelength, as shown in Figure 5-12. In the laser world, the frequencies are very high, and light waves are normally identified




Figure 5-12. In second-harmonic generation, a high power beam at the fundamental frequency is doubled to the second harmonic by passing through a nonlinear material. The harmonic output power inevitably is lower than the input, and some residual input at the fundamental frequency light remains unless it is filtered out.

by wavelength, even though the process is still called harmonic generation. Thus, laser specialists say they frequency-double the output of a 1064-nm neodymium laser to produce 532-nm light, but they never mention the frequencies involved (2.83×10^{14} and 5.66×10^{14} hertz, respectively).

As in music, high-order harmonics are integral multiples of the original frequency. The second harmonic is generated by the (Input)² term in the formula for nonlinear output (Equation 5-7). The third harmonic is three times the fundamental frequency, corresponding to the (Input)³ term. The fourth harmonic and higher-order harmonics follow the same rules, each corresponding to a higher power of the (Input) term.

In practice, the second harmonic is the easiest to produce, so normally optical harmonic power is highest at the second harmonic, with lower powers at higher harmonics. The strength of nonlinear effects depends on the internal structure of the material. Nonlinear effects are very weak in most materials such as ordinary glass. However, some other materials have much stronger nonlinear coefficients, and those strongly nonlinear materials can produce much higher harmonic powers. Important nonlinear optical materials are listed in Section 5.6.3.

High power density increases the output power of harmonic generation. That can be produced by using a high-power laser and by focusing the beam onto a small spot on nonlinear materials. Very high power can also be achieved by concentrating the energy from the laser into very short pulses with very high peak power, so the highest optical harmonic power comes in short pulses.

Optical harmonic generation is not as simple as this brief description may make it seem. Efficient second harmonic generation requires using highly nonlinear materials in which the refractive indexes can be matched at the first and second harmonics so the two waves travel at the same speed through the material and their phases match to aid harmonic generation.

Second-harmonic generation is generally the most efficient and most important type of harmonic generation. The third harmonic can be generated directly, but it is often more practical to first generate the second harmonic, then mix it with the fundamental beam to produce the third harmonic as the sum of the fundamental frequency and the second harmonic. Likewise, the fourth harmonic is generally easiest to produce by doubling the second harmonic than by directly producing the fourth harmonic. The fifth harmonic can be produced by generating the sum of the second and third

harmonics. Higher-order harmonics tend to be even less efficient, but in some cases, they are the most efficient way to produce very short wavelengths, as described in Section 11.5.2.

5.6.2 Sum- and Difference-Frequency Generation

Nonlinear optics can also generate light waves at frequencies that are the sum or difference of the frequencies of two light waves with different wavelengths (or frequencies) passing through a nonlinear material. The strength of the sum and difference frequencies is proportional to the product of the input intensities of the two waves. Note that harmonic generation produces the sum frequency of two identical waves.

Sum-frequency generation is described by the formula

$$\nu_1 + \nu_2 = \nu_3 \quad (5-8)$$

where ν_1 and ν_2 are frequencies, and ν_3 is larger than the other two. Similarly, *difference-frequency generation* is described by the formula

$$\nu_1 - \nu_2 = \nu_3 \quad (5-9)$$

where in this case ν_1 is the highest frequency.

Like harmonic generation, sum- and difference-frequency generation is proportional to the square of the power or the second term in the nonlinear formula (Equation 5-7). By convention, these effects may be called *third-order nonlinearities* because they involve three photons, although they rely on the second-order nonlinearity. In harmonic generation, two photons from the input beam with the same frequency combine to generate a third photon with twice the input energy, so $\nu_1 = \nu_2$. In sum- and difference-frequency generation, the two photons that combine to generate the output frequency are different.

Sum- and difference-frequency generation offers another way to generate new wavelengths, in this case by mixing two different beams to generate a third frequency.

5.6.3 Parametric Gain

Another type of process involving three photons is called *parametric gain* and it is a bit more complex because it essentially uses difference-frequency generation in a nonlinear material to amplify a

weak beam of light. In this case, two beams enter a nonlinear crystal, one a stronger pump beam and the other a weaker “signal” beam. Within the nonlinear medium, a pump photon interacts with a signal photon to produce an additional signal photon plus a photon at the difference frequency, called the “idler.” The initial signal photon survives, so the process essentially amplifies the “signal” frequency and produces one photon at the idler frequency. You can think of the result as similar to a chemical equation, with the input photons on the left and the output photons on the right.

$$\nu_{\text{pump}} + \nu_{\text{signal}} = 2\nu_{\text{signal}} + \nu_{\text{idler}} \quad (5-10)$$

Both sides of the process have the same amount of energy, but it is in different forms, one pump and one signal photon as the input and two signal photons and an idler photon on the output (right).

The result is an amplification process, where input pump photons are split into signal and idler photons, depleting the population of pump photons and increasing the number of signal and idler photons. What you actually get out depends on the optics of the system, which can be designed either to collect the signal photons and discard the idler photons or to collect both. Section 11.2 describes two types of light sources based on parametric gain. Optical parametric amplifiers combine the pump and signal frequencies in a single pass through a nonlinear medium to amplify the signal frequency, with the idler normally lost. Optical parametric oscillators receive only pump photons from an external laser sources and rely on quantum fluctuations to start the parametric process by effectively splitting photons into signal and idler photons selected by resonances in a resonant cavity.

A final note: other nonlinear processes exist that involve more complex interactions, but we will not discuss them here because they are rarely used to intentionally convert laser wavelengths, and because the terminology and concepts are more likely to confuse you than to enlighten you at this stage of your learning about lasers.

5.6.4 Raman Scattering

Another important nonlinear interaction in many materials is a change in photon energy called Raman scattering, named after its discoverer, Indian physicist C. V. Raman. In normal light scattering, an atom or molecule absorbs a photon and almost instantly re-emits

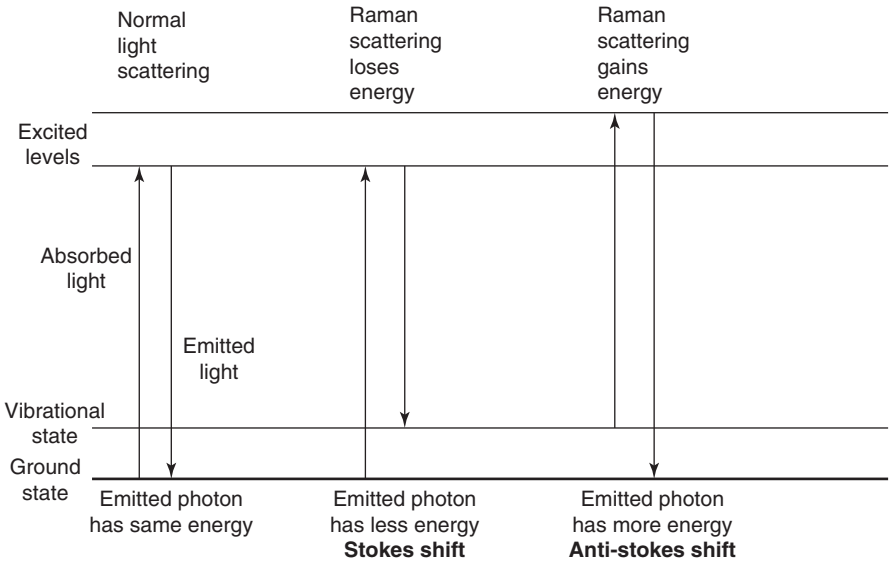


Figure 5-13. Raman scattering changes the energy of the scattered photon, unlike normal scattering, which does not change energy. Raman scattering that keeps a bit of the absorbed energy to move into an excited vibrational state is called Stokes scattering. If the photon is scattered from an atom in an excited vibrational state, the scattered photon may carry away that vibrational energy in what is called anti-Stokes scattering.

a photon of the same energy in another direction. In Raman scattering, an atom or molecule absorbs a photon but adds or subtracts a bit of energy before it scatters the light, so the emitted photon has more or less energy than the initial photon, as shown in Figure 5-13.

The change comes for one of two reasons. If the atom or molecule was initially in an energy level above the ground state, that extra bit of energy may be added to the emitted photon, so it has a higher energy or shorter wavelength. If the atom or molecule was initially in the ground state, it may keep some energy from the input photon to increase the energy level of an electron in its outer shell, leaving the emitted photon with less energy or a longer wavelength. It is more common for Raman scattering to leave the scattered photon with less energy (a longer wavelength) than with more energy (a shorter wavelength).

The Raman shift (measured in frequency or energy change of the photon) is an inherent characteristic of the material's vibrational or rotational state, so for any particularly material, such as silica

glass, the energy shift is the same for any illumination wavelength. Raman scattering is a valuable tool for shifting laser wavelength to otherwise unavailable lines in the visible, near infrared, or near ultraviolet. As you will learn in Section 9.6, an effect called stimulated Raman scattering can serve as the basis for optical amplification and laser action in optical fibers.

Fluorescence is similar to Raman scattering, but in fluorescence the atom or molecule absorbs a photon, making a transition to a higher-energy level, then drops to a different state and emits a photon. In Raman scattering, the input photon is not absorbed by the atom or molecule.

5.6.5 Nonlinear Materials

Nonlinearity is an inherent property of materials that can be measured as a nonlinearity coefficient. Generally, nonlinearity arises from strain within the lattice of the crystal. Most materials have little strain and are only weakly nonlinear, but some are highly strained and thus are highly nonlinear. Several materials with high levels of optical linearity are used with lasers:

- Beta barium borate (β -BaB₂O₄ or BBO)
- Lithium iodate (LiIO₃)
- Lithium niobate (LiNbO₃)
- Lithium triborate (LiB₃O₅ or LBO)
- Potassium dihydrogen phosphate (KH₂PO₄ or KDP)
- Potassium niobate (KNbO₃)
- Potassium titanyl phosphate (KTiOPQ₄ or KTP)

The use of these materials is limited by their optical transparency range, which must include both fundamental and harmonic wavelengths. The material must also meet phase-matching requirements. Fabricating the crystals in a series of layers of alternating orientation, a technique called *periodic poling*, eases restrictions on phase matching and improves the performance of nonlinear materials.

5.7 BEAM MODULATION AND OUTPUT CONTROL

Many laser applications require control over the laser's output power over time. This can be done within the laser by changing the

drive power, or externally by modulating the fraction of the laser beam transmitted by an external modulator. The choice generally depends on the type of laser and the application requirements.

5.7.1 Internal Power Control and Direct Modulation

The simplest kind of *direct modulation* of laser output is by controlling the input power to the laser. You can do this easily with a laser pointer; push the button and the beam turns on, release the button and the power goes off.

Direct modulation works very well with semiconductor diode lasers. Like a laser pointer, you can turn them on by pushing a button. However, you can also modulate the light electronically. In a simple fiber-optic data link, an input electrical signal switches the laser on and off to transmit data bits. Direct modulation of diode lasers can be very fast because they are very small devices, allowing data rates over a billion bits per second. Diode lasers are also reasonably linear over much of their operating range, so they can be modulated with continuously varying analog signals.

Direct modulation is harder for many other types of lasers. The larger the laser and the lower the gain, the longer it takes to turn on. For example, light takes about one nanosecond to travel 30 cm, so it would take two nanoseconds to make a round trip through a 30-cm helium–neon laser tube. In fact, the emission would take much longer to build and stabilize, because the gain of a helium–neon laser is so low that many round trips would be needed. A few lasers are even worse because they require heating or cooling to run properly and need time to stabilize.

Some lasers have a type of modulation built in by their internal physics. The excimer lasers described in Section 7.7 are powered by intense electrical discharges through the laser gas, but the discharges cannot last longer than about 100 nanoseconds, after which the laser emission stops.

The operation of optically pumped lasers is limited by the light sources that power them. Semiconductor diode lasers are easy to modulate directly, so modulating the pump diodes can directly modulate a diode-pumped solid-state laser. Direct modulation of other lasers pumped by other types of lasers or by flashlamps similarly is limited by the light sources pumping them.

Many lasers are built with internal devices that control their output, usually to generate a series of repetitive pulses that depend

on the type of device and its design. Although the purchase of these internal devices may be optional, they play such an important role in control of the laser's operation that we consider them as determining the type of laser and cover them in Section 6.4.

5.7.2 External Modulation and Modulators

Alternatively, a modulator placed outside the laser cavity can modulate the beam by changing the fraction of light it transmits. Some external modulators can be computer-controlled or driven by data signals to modulate the beam in almost any way desired. Others are limited to producing a series of regularly spaced pulses. There are several approaches to external modulation.

Mechanical modulation is typically done with a shutter or beam chopper. *Shutters*, like those on cameras, open to let the beam through and close to block light. Shutters serve as safety cutoffs on many lasers, to prevent inadvertent exposure to the beam. Shutters are controlled manually or automatically and operate in a fraction of a second.

Beam choppers are mechanical devices that interrupt the beam repeatedly and are typically faster than shutters. One example is a rotating disk with holes or slots that transmit the beam for brief intervals, as shown in Figure 5-14. When you align a laser beam with the holes in the disk, spinning the disk switches the beam off and on regularly. Another example is a mask placed on the arm of a vibrating tuning fork, which swings in and out of the beam to block or transmit it. Both types of beam choppers generate a series of uniformly spaced pulses.

Optical modulators are computer controllable devices that vary the transparency of a material in the beam path. Three types of optical modulators are common:

- *Electro-optic phase modulators.* Application of an electric field across materials such as lithium niobate (LiNbO_3) changes the refractive index in proportion to the voltage. This shifts the phase of the light and can be used to modulate light amplitude if the phase-shifted light is combined to interfere with a portion of the light that has not been phase-shifted.
- *Electro-optic polarization rotators.* Application of an electric field across similar materials changes refractive index differently for




Figure 5-14. A beam chopper is a spinning disk with holes or slots that transmit a laser beam during intervals in which the solid part of the disk does not block the light. The faster the disk spins, the faster the chopper switches the light off and on.

different polarizations of light, effectively rotating the polarization of the light. Combining a polarization rotator with a suitably oriented polarizer can modulate amplitude because the polarizer's transmission changes with the polarization rotation.

- *Acousto-optic modulators.* When an acoustic wave passes through a solid, it alternately raises and lowers the pressure, which changes the refractive index. The stronger the acoustic wave, the more light is scattered from the main beam, modulating its amplitude. (The same effect can be used to deflect part of the laser beam at another angle, as you saw in Section 5.4.1.)

5.8 MEASUREMENTS IN OPTICS

Optics has its own specialized measurement practices and terminology that are used along with common measurements such as power and energy when working with lasers. This section will introduce you to important optical measurements.

5.8.1 Light Detection

Today, most light measurement is done electronically. A device called a *photodetector* senses the light and generates an electronic signal proportional to the input light power or energy. Then, the electronic signal is converted into a measurement of the optical power using calibration scales that account for the sensitivity of the light detector.

Most detectors are semiconductor devices, which fall into two main categories: *photoconductive* and *photovoltaic*. Their mechanisms differ subtly. Photoconductive devices become more conductive when illuminated by light, so they measure light by measuring a change in resistance. Photovoltaic detectors, often called *photodiodes*, are two-terminal semiconductor devices called *diodes*, which when illuminated generate an electric current proportional to the light that illuminates them. The electrical responses of those semiconductor detectors can be converted into measurements of optical power detected.

The wavelength response of a photovoltaic detector depends on its composition because a photon must deliver enough energy to free an electron from a valence bond. This means the sensitivity of the device varies with wavelength. Table 5-2 lists the wavelength ranges of some important types. Some photovoltaic detectors, called *avalanche photodiodes*, have internal amplification stages, which increases their sensitivity to light, so they can detect weaker signals.

You may also encounter an older generation of vacuum-tube light detectors called *vacuum photodiodes* in which light frees

Table 5-2. Wavelength ranges of important visible and infrared detectors

Type	Material	Wavelengths (nm)
Photoemissive	Potassium-cesium-antimony	200-600
Photovoltaic (semiconductor)	Silicon	400-1000
PV semiconductor	Germanium	600-1600
PV semiconductor	Gallium arsenide	800-1000
PV semiconductor	Indium gallium arsenide	1000-1700
PV semiconductor	Indium arsenide	1500-3000
PV semiconductor	Lead sulfide	1500-3300
PV semiconductor	Lead selenide	1500-6000
PV semiconductor	Mercury cadmium telluride	2700-10,600
PV semiconductor	Indium arsenide antimonide	1000-11,000
Thermoelectromotive	Thermopile	1000-25,000

electrons from a *photoemissive* metal surface, and the electrons flow through the vacuum to a positive electrode. Electrons from a vacuum photodiode can be passed through a cascade of amplifying stages in the vacuum to boost the signal strength, making a *photomultiplier tube* that has very high sensitivity, and is still used in research to measure very weak optical signals. There are also *solid-state photomultipliers* based on semiconductors that detect light and then pass the photocurrent through amplification stages.

Special detectors are needed for wavelengths longer than about 1700 nm. Some are semiconductor devices which respond to the lower energy of longer-wavelength infrared photons, notably mercury cadmium telluride (HgCdTe) and indium arsenide antimonide (InAsSb) in the mid-infrared.

Thermopile detectors are arrays of thermocouples that convert heat energy into electrical energy, with their response depending on temperature differential rather than absolute temperature. Their big advantage is a uniform response across a broad range in the infrared, extending between about 1 and 25 μm .

Most practical optical measurements are made with power and energy meters or other instruments specifically designed and calibrated for the task. Many come with a single detector, but others come with different detectors for use in different spectral ranges.

5.8.2 Radiometry, Photometry, and Light Measurement

Light measurement is often divided into two categories—radiometry and photometry—although the latter term is often misused. *Radiometry* is the measurement of power and energy contained in electromagnetic radiation, regardless of wavelength, including the infrared and ultraviolet. Strictly speaking, *photometry* is the measurement only of light visible to the human eye at 400 to 700 nm, taking into account how eye sensitivity varies across that range. Proper photometry makes measurements in units such as lumens. However, many companies offer “photometers” that measure light in radiometric units. You can also find instruments called *spectrophotometers* which measure the power of a series of narrow bands of wavelengths across of the spectrum in radiometric units.

Virtually all laser-related measurements use radiometric units. Table 5-3 summarizes these units and lists common symbols for

Table 5-3. Radiometric units

Quantity and symbol	Meaning	Units
Energy (Q)	Amount of light energy	joules (J)
Power (P or ϕ)	Light energy flowing past a point during an interval ($\Delta Q/\Delta t$)	watts (W)
Intensity (I)	Power per unit solid angle	watts/steradian
Irradiance (E)	Power incident per unit area	watts/cm ²
Radiance (L)	Power per unit angle per unit projected area	W/steradian-m ²

them, with brief descriptions of their meanings. As the table shows, these units are related to each other. *Power* is the rate of energy flow past a point in a given time; it is measured in watts. *Irradiance* is power per unit area; it is measured in W/cm². *Intensity* is power per unit solid angle, W/steradian. Similar units are defined for photometric power visible to the human eye.

Many instruments also measure power in decibels, a logarithmic ratio of two power levels, P_{in} and P_{out} :

$$dB = 10 \times \log_{10} \left(\frac{P_{\text{out}}}{P_{\text{in}}} \right) \quad (5-11)$$

The value is positive when the output power is higher than the input, and negative when the output power is lower, that is, when some power has been lost. Power can also be measured in decibels relative to a predefined level, typically one milliwatt (1 mW) or one microwatt (1 μ W). A positive number indicates that the measured power level is above the comparison, a negative number indicates it is lower.

If you are familiar with electronic measurements, you may have seen the logarithm of the voltage or current ratio multiplied by 20 rather than 10. This is because decibels are a power ratio, and power is proportional to the square of voltage or current, so the factor of 2 must be added. Optical measurements are in power units, so the proper ratio is 10.

5.8.3 Laser Safety and Light Measurement

Light measurement is important in assessing the safety of lasers, which is described briefly in Appendix A. The eye is the organ most

sensitive to light, but invisible infrared light at wavelengths as long as $1.4\ \mu\text{m}$ can also penetrate the eye and cause retinal damage, so safety measurements must include those wavelengths.

A recent study by the National Institute of Standards and Technology pointed out the importance of measuring all potentially hazardous wavelengths. Scientists used carefully calibrated instruments to measure light emitted by green laser pointers, readily available inexpensively on the Internet. As you will learn in Section 8.4.2, those devices generate three wavelengths internally. A semiconductor diode laser emits infrared light at 808 nm, which pumps a solid-state neodymium laser emitting at 1064 nm in the infrared, and a nonlinear second-harmonic generator converts the 1064-nm light to green light at 532 nm. Properly made green lasers should include filters to block the infrared light, but the NIST measurements found that many of the green lasers emitted enough infrared light to damage the retina if the beam was pointed into a person's eye.

5.8.4 Beam Quality and Characterization

Beam quality is an important metric for laser measurements, but it can be a hard one to pin down because the definition of quality depends on the application. In general, the main concern is the distribution of laser power in the beam at a distance from the laser, which are important for laser uses including industrial materials working and military weapons.

Much beam profiling involves imaging a cross-section of the beam to see where power is concentrated and how that compares to application requirements. Beam quality can also be measured on three related scales:

- *Beam Parameter Product (BPP)*, which equals half the divergence angle in milliradians times the beam waist radius in millimeters. The smallest possible BPP is for a diffraction-limited Gaussian beam, which at $1.064\ \mu\text{m}$ comes to $0.339\ \text{mm-mrad}$.
- *Beam Propagation Factor (M^2)*, which is the product of π times beam radius at the waist (w_0) times beam divergence (θ) divided by wavelength (λ), or

$$M^2 = \frac{\pi w_0 \theta}{\lambda} \quad (5-12)$$

The minimum is 1 for a diffraction-limited Gaussian beam.

- *Power in the Bucket*, which measures how much power is deposited within the target area and ignores anything outside that region. It is the simplest criterion and can be measured directly.

5.9 WHAT HAVE WE LEARNED?

- A simple lens is made of a single piece of transparent material, with one or both sides curved to focus light.
- A lens can focus parallel light rays to a point or light from a point source into a parallel beam.
- A collimator is a simple telescope that aligns light so it is all parallel; it may expand or shrink a laser beam.
- Chromatic aberration is the focusing of light to different points depending on its wavelength.
- Achromatic lenses use two materials to compensate for chromatic aberration.
- Cylindrical optics focus light to a line rather than to a point.
- A retroreflector bounces light directly back to its source.
- Prisms and diffraction gratings are dispersive optics that spread light out to show a spectrum.
- Light can be linearly, circularly, or elliptically polarized.
- A beam splitter divides an input beam into two parts.
- Silica glass is widely used for transparent optics at visible wavelengths. Other materials are used in the infrared and ultraviolet.
- Metal coatings on glass are used as mirrors.
- A low-index coating can reduce surface reflection by glass.
- Multilayer coatings use interference effects to control light transmission and reflection as a function of wavelength.
- A neutral-density filter absorbs the same fraction of light at all wavelengths in its operating range.
- Beam scanners deflect beams to sweep across a field.
- Optical fibers manipulated by robotic arms direct laser beams for machining.
- Optical tables provide a stable base for lasers and optical equipment.
- Nonlinear optical effects are proportional to the square or higher power of light intensity.

- Harmonic generation uses nonlinear effects to produce a second- or higher-order harmonic of the input laser frequency.
- Adding or subtracting the frequencies of two input waves can produce new wavelengths at the sum or difference of the two frequencies. This can produce wavelength-tunable light.
- Nonlinear effects require special nonlinear materials.
- Raman shifting changes wavelength by adding or subtracting vibrational energy.
- Direct modulation changes laser output by modulating input power.
- External modulation deflects or absorbs part of the beam to change output power.
- Power is the flow of light energy past a point over a period of time.
- Semiconductor photodiodes are widely used to detect light.
- Radiometry measures electromagnetic power at all wavelengths.
- Photometry measures only light visible to the human eye.

WHAT'S NEXT?

In Chapter 6, we will describe the major types of lasers and how they differ.

QUIZ FOR CHAPTER 5

1. A double-convex lens with two equal radii of curvature of 20 cm has refractive index of 1.60 at 400 nm and 1.50 at 700 nm. What is the difference between the lens' focal lengths at those two wavelengths?
 - a. 400 nm focal length is 3.33 cm shorter
 - b. 400 nm focal length is 0.33 cm shorter
 - c. 700 nm focal length is 0.033 cm shorter
 - d. 700 nm focal length is 1 cm shorter
 - e. No difference
2. Silicon has a refractive index of 3.42 at 6 μm in the infrared. What is the reflective loss for 6- μm light incident from air (with refractive index of 1.0) normal to the surface?
 - a. 5%
 - b. 10%
 - c. 20%

- d. 30%
 - e. 40%
3. For the silicon sample in problem 2, what is the reflective loss if a coating with refractive index of 2 is applied on the surface? (Remember that you cannot just add the reflective losses. You must multiply the fractions of transmitted light to get the total transmitted light to assess overall reflective loss.)
- a. 5%
 - b. 10%
 - c. 15.4%
 - d. 17.2%
 - e. 20.5%
4. Which of the following optical materials is transparent in the entire optical spectrum between 0.4 and 0.7 μm ? (Use Figure 5-8.)
- a. Gallium arsenide
 - b. Silicon
 - c. Magnesium fluoride
 - d. Zinc selenide
 - e. Silver chloride
5. You need an optical material that is transparent between 0.9 and 1.0 μm in the infrared. Which of the following materials is not suitable? (Use Figure 5-8.)
- a. Sodium chloride
 - b. Silicon
 - c. Zinc selenide
 - d. Magnesium fluoride
 - e. Silver chloride
6. What type of filter would you use to block light from a narrow-line laser, but transmit other light?
- a. Neutral-density filter
 - b. Interference filter
 - c. Color filter
 - d. Spatial filter
 - e. Any of the above
7. What is the wavelength of the fourth harmonic of the ruby laser (694 nm)?
- a. 2776 nm
 - b. 1060 nm
 - c. 698 nm
 - d. 347 nm
 - e. 173.5 nm

8. What is the wavelength of the difference frequency generated by mixing light from a ruby laser at 694 nm with a neodymium laser emitting at 1064 nm?
 - a. 370 nm
 - b. 532 nm
 - c. 694 nm
 - d. 1758 nm
 - e. 1996 nm
9. Raman shifting does what to input light?
 - a. Modulates its intensity
 - b. Doubles its frequency
 - c. Changes its wavelength by a modest amount
 - d. A and B
 - e. Nothing
10. Direct modulation by changing the drive current works best for which type of laser?
 - a. Semiconductor diode
 - b. Ruby
 - c. Helium–neon
 - d. Neodymium–YAG
 - e. All are equally difficult
11. Which detector material can sense green light at 525 nm?
 - a. Silicon
 - b. Germanium
 - c. Gallium arsenide
 - d. Indium arsenide
 - e. All of the above.
12. What do decibels measure?
 - a. Amount of light energy passing a point per second.
 - b. Logarithmic ratio of two power levels.
 - c. Nonlinearity of optical materials.
 - d. Ratio of visible light in lumens to all light in watts.
 - e. Power incident per unit area.

LASER TYPES, FEATURES, AND ENHANCEMENTS

ABOUT THIS CHAPTER

Now that we have described the basics of lasers and associated optics, it is time to look at the wide variety of lasers. They are made from different materials, powered in different ways, and store energy differently. Lasers may emit continuously or in pulses of various durations, produced in different ways. Their emission bandwidth ranges from fixed and extremely narrow to tunable across a wide range. The output wavelength can be changed to meet application requirements. They may be single-stage oscillators or include both oscillator and amplifier stages. There are also laser-like sources that act like lasers in many ways but do not meet all the standard criteria to be called lasers. This chapter explains these differences and prepares you to understand the following chapters that cover specific laser types.

6.1 PERSPECTIVES ON LASER TYPES

What is most important about lasers is a matter of perspective. Traditionally, laser developers thought first about the material used in the laser and how that laser medium was excited to produce a laser beam. On the other hand, laser users care more about the light

emerging from the laser and its properties, such as wavelength, output power, and whether the laser emits light continuously or in a series of pulses.

Both these perspectives have advantages, but it is hard to organize a book in a way that reflects both perspectives simultaneously. For convenience, I have chosen to group lasers by the type of material used because similar materials usually work in similar ways. Chapters 7 through 10 group cover four major classes of materials in widespread practical use. Chapter 11 covers other types of lasers and laser-like devices that do not fit into earlier chapters.

Chapter 7 covers gas lasers, which share many properties because most are powered by electric discharges passing through a gas. However, the properties of the gases differ widely, as does their output, which ranges from intense nanosecond pulses in the ultraviolet to continuous output from milliwatts to watts in the visible and to kilowatts in the infrared.

Chapters 8 and 9 cover two types often called *solid-state lasers*, a term with a special meaning in the laser world, as described in Box 6-1. These are transparent nonconductive solids like glass that contain small fractions—usually 1% or less—of atoms that absorb light energy and emit on laser transitions. All are excited in the same way by light from an external source, but differ in the shape of the light-emitting material. Chapter 8 covers lasers in which the solid is a rod, slab, or disk. Chapter 9 covers lasers in which the solid is a long, thin optical fiber with a core-cladding structure that guides light along the fiber. That distinction may seem small, but it makes a big difference in how the lasers work, so the industry usually considers them different types.

Chapter 10 covers lasers made of semiconductors, which are solids that conduct electric current reasonably well. They are called semiconductors because electrons can move about more freely than in insulators such as glass, but not as freely as in conductors such as copper and other metals. Semiconductors are considered solid-state devices in electronics and physics. But the laser world distinguishes conductive semiconductor materials from the nonconductive solid-state materials used in solid-state lasers. As you will learn later in this chapter, that is an important practical decision.

As you will learn in Chapter 11, many lasers do not fit neatly into any of those four categories, but few of them are in wide use.

BOX 6.1 LASER-SPECIFIC JARGON: WHAT IS A “SOLID-STATE” LASER

In electronics, the term solid-state often means made of semiconductors. In the laser world, solid-state usually means *not* semiconductor. It is a quirk of terminology that exists for a very good reason. Semiconductor lasers generate light in a way fundamentally different from other solid-state lasers.

The term *solid-state laser* specifically covers lasers made from glass, crystals, ceramics, or other solids described in Section 6.2.4. These solids do not conduct electricity, so they cannot be powered electrically. They do transmit light, so the light-emitting atoms they contain can be excited by light at a wavelength transmitted by the material.

In contrast, semiconductors conduct current, although not as well as conductors like copper. This allows electrons flowing through the semiconductor to deliver energy to light-emitting atoms in the semiconductor, which has important advantages, so most lasers made of semiconductors are powered electrically. Such semiconductor lasers are two-terminal electrical devices called diodes and are usually called *diode lasers* or *laser diodes*. Chapter 10 explains the details. Just to complicate things, some semiconductor lasers are powered by light, which has advantages in certain cases. Those *optically pumped semiconductor lasers* or *OPSLs* are covered in Section 8.5 in the chapter on solid-state lasers.

6.2 LASER MEDIA

In principle, lasers can be made from many materials. All you need are a laser cavity and a population inversion. In 1970, future Nobel Laureates Arthur Schawlow and Theodor Hänsch made a “Jell-O” laser by adding a fluorescent dye to gelatin and illuminating it. Other laser pioneers joked that even a telephone pole would “lase” if zapped with enough energy, but it probably would have caught fire instead.

The properties of laser materials differ widely, and a few make much better lasers than most others. Key factors include how well

the material absorbs pump energy, how long a population inversion lasts, the wavelength of light generated, and how well extra heat can escape from the material. In practice, only a tiny fraction of the many thousands of lasers demonstrated in the laboratory ever become commercially viable products. The most successful lasers are those that fulfill a particular need, such as cutting metals easily at high powers, or generating red or green light inexpensively. This section focuses on lasers that have been used in research, industry, or medicine.

6.2.1 Atomic Gases and Ions

Atomic gases are simple systems for physicists to study, so they were among the first materials used to make lasers. Atoms can be treated as isolated except when they collide with other atoms, or when photons or electrons transfer energy to or from them. Typically, these gases are sealed in long thin tubes and are excited by electric current flowing through the gas along the length of the tube, with electrons transferring energy to gas atoms.

The most important laser transitions are those of electrons in the outer shells of neutral atoms or of ions, which are atoms that have lost one or more electrons. The neutral atoms most important in lasers are rare gases such as neon, argon, krypton, or neon, which emit in the near infrared or visible. Ions which have lost one or two electrons emit at shorter visible wavelengths and into the ultraviolet. Those lasers are described in Sections 7.3 and 7.4. Some atomic gas lasers are based on vaporized metals, including cadmium and copper, described in Section 7.5. These lasers must be heated for the laser to operate.

Many atomic gas lasers contain two types of atoms, with the electric discharge transferring energy to one species, and that species then transferring the energy to a second species to produce a population inversion. For example, in the helium–neon laser the more abundant helium atoms absorb energy from current flowing through the gas and then transfer energy to the neon atoms, exciting an outer-shell electron to emit laser light.

Atomic gas lasers are a mature technology, but they retain some important applications, such as coherent light sources for holography. Research continues on some new types, notably the alkali vapor lasers described in Section 7.10.3.

6.2.2 Molecular Gases

Chapter 7 also covers gas lasers in which molecules make transitions between vibrational and rotational energy levels as well as on electronic transitions. Electronic transitions of molecules are generally at wavelengths shorter than $1\ \mu\text{m}$. Vibrational transitions occur at wavelengths between about 3 and $12\ \mu\text{m}$, with rotational transitions at lower frequencies, corresponding to longer wavelengths. Molecules usually make rotational transitions at the same time as vibrational transitions, so the energy released in a photon is the sum of rotational and vibrational transitions and can vary over a range of wavelengths. Longer-wavelength transitions usually occur between purely rotational energy levels.

Molecular laser transitions tend to involve energy levels relatively close to the ground state, so they convert pump energy into light more efficiently than shorter-wavelength atomic gas lasers. This higher efficiency produces less waste heat and allows higher operating powers. That efficiency is an important advantage of the carbon-dioxide laser emitting at 9 to $11\ \mu\text{m}$, which is the most widely used molecular gas laser, described in Section 7.6.

Most molecular gas lasers are excited by electric discharges, but some are excited by light from another laser, often CO_2 . Strong chemical reactions that release energy can also power molecular gas lasers when they produce excited molecules, as described in Section 7.9, but they are not widely used.

6.2.3 Short-Pulsed Molecular Gas Lasers

Strong pulsed electrical discharges can produce strong ultraviolet pulses from some diatomic molecules. The simplest of these is the nitrogen laser, described in Section 7.8, which has a small niche in the laser market.

The most commercially important short-pulsed molecular gas lasers are based on a family of short-lived diatomic molecules called *excimers*, which give the lasers their name. They are covered in Section 7.7. The pulsed electrical discharge triggers a reaction between a rare gas such as krypton and a halogen such as fluorine to produce a rare-gas-halide molecule such as krypton fluoride (KrF). These molecules are inherently unstable, falling apart when they emit their excess energy in the form of an ultraviolet photon. (The rare

gases have complete outer electron shells, so they normally do not react with other atoms, so they have also been called *inert* or *noble gases*.)

Excimer lasers are complex and expensive devices that use some hazardous gases, but their unusual physics can generate short and powerful pulses of ultraviolet light not available from other sources. They have found important applications in refractive eye surgery and producing semiconductor electronics.

6.2.4 Dielectric Solids and Solid-State Lasers

The solids used in solid-state lasers, described in Chapter 8, are called *dielectric* materials, meaning they are electrically nonconductive or insulating, such as glass, and have electrons that are tightly bound to atoms so they cannot conduct electricity. In practice, the bulk of a solid-state laser usually is a glass or crystalline *host* material that must be transparent to allow optical pumping from an external source. A small fraction of the atoms are an *active material* or *laser species* that absorbs pump light and emits laser light. In the ruby laser, the host is sapphire, a crystalline form of aluminum oxide (Al_2O_3), and the active material is chromium, which replaces a few aluminum atoms and emits the red light in the laser beam.

Normally, the laser species makes up 1% or less of the mass of the solid. They are essentially dopants added to a glass or crystal, but they largely determine the optical features of a laser, such as its emission wavelength, pump bands, and gain. The host material is chosen for its durability, stability, ease of fabrication, heat dissipation, interaction with the active material, and transparency at pump wavelengths.

The pump source is a vital element of the laser that must produce a population inversion by exciting the laser species to a high-energy state, from which it drops to the upper laser level and then produces stimulated emission, as shown in Figure 6-1. Bright lamps were used to pump early solid-state lasers. Now semiconductor diode lasers are widely used because they allow much more efficient pumping, as explained in Section 8.2.3.3.

Early solid-state lasers were simply long cylindrical rods, but other shapes have come into use, including slabs and thin disks, as described in Sections 8.3.2 and 8.3.3. Removing excess heat from the rod or other laser medium remains an important issue.




Figure 6-1. Energy levels of the light-emitting species in a generic four-level solid-state laser. The solid host material is transparent at these wavelengths.

6.2.5 Fiber Lasers and Amplifiers

You can think of a fiber laser as a very thin and very long laser rod. However, optical fibers also have an internal structure that guides light along the length of the fiber. These light-guiding structures concentrate the pump light onto an inner core where the active laser material is concentrated. Concentrating pump power onto a small volume increases the efficiency of laser pumping and the laser's overall efficiency in converting input energy into laser light. As you will learn in Chapter 9, this makes fiber lasers the most energy-efficient solid-state lasers when they are pumped with semiconductor lasers.

The high efficiency of fiber lasers is very attractive for use in industrial machining such as cutting and welding, which can require steady powers of kilowatts or more. The technology for fiber lasers also has another important application in amplifying light in fiber-optic communication systems. You will learn more about optical amplifiers and how they differ from laser oscillators in Section 6.5.

6.2.6 Semiconductor Diode Lasers

Semiconductor lasers work quite differently than the solid-state lasers described above, which is why laser specialists do not classify semiconductors as solid-state. As you read in Box 6.1,

semiconductors can conduct electricity, unlike the dielectric materials in solid-state lasers, so semiconductor lasers can be powered efficiently by electrical current passing through them. In addition, almost all semiconductor lasers have internal structures that make them electrical diodes, so they can be called *diode lasers* or *laser diodes*.

Diodes are two-terminal electronic devices that conduct current readily in one direction but poorly or not at all in the other. In electrical terms, they have high resistance in one direction and low resistance in the other. What makes that happen in semiconductors is adding materials that change the material's behavior to selected parts of the semiconductor. You will learn more about this in Chapter 10.

Semiconductor diodes can control current flow in electronics. Diodes of some materials, notably silicon, do not emit light. However, diodes made of some other semiconductors convert some of the electrical energy passing through them into light, making them *light-emitting diodes* or *LEDs*, which are all around us in lighting and displays. The semiconductor materials that are best in producing light are compounds such as gallium arsenide or gallium nitride.

Increasing the current density (the number of electrons flowing through an area) and adding a laser cavity can push an LED over the threshold for laser oscillation. The resulting diode laser differs in some important ways from solid-state or gas lasers. Building structures into the diode laser improves its efficiency in generating light, so it is emitting very brightly from a very small area. Diode lasers can convert over 60% the input electric energy into light, an incredible efficiency by laser standards.

The bad news is that the light is emitted from a small area, so diffraction causes it to spread out rapidly with distance. However, suitable optics can collect that light and direct it into a fiber laser, where most of it can be converted into a nice, tightly focused laser beam. So diode lasers have become very important, as you will learn more about in Chapter 10.

We should note that not all semiconductors are diodes. Thin semiconductor disks can be optically pumped to make lasers that behave more like solid-state lasers, as described in Section 8.5.

6.2.7 Organic Dyes and Liquid Lasers

Laser action can occur in liquids as well as in gases and solids. A number of liquid lasers have been demonstrated in the laboratory,

but only one type has found significant practical use, the tunable dye laser.

Like the solid-state laser, the dye laser consists of a light-emitting species dissolved in a host material. In this case, the light emitter is an organic dye and the host usually is a liquid solvent in which the dye is dissolved. (Solid plastics can also serve as dye hosts, but they are not standard products.) The dyes are fluorescent materials that absorb light, store the energy briefly, then drop to a lower energy level and emit light at a longer wavelength. The dyes are also complex molecules that vibrate as well as undergo electronic transitions, so their light-emitting transitions are spread across a range of wavelengths. That allows their wavelength to be tuned by adjusting the laser cavity to resonate at different wavelengths. Dye lasers are valuable for applications which require specific wavelengths not available from other lasers, or which require tuning across a range of wavelengths. You will learn more about them in Section 11.1.

6.2.8 Free-Electron Lasers

Stimulated emission can also extract energy from unbound or *free electrons* in a beam of high-energy electrons passing through an array of magnets with their fields oriented in alternating directions. The interaction of the electrons with the magnetic field causes them to emit electromagnetic energy, which can stimulate the emission of more electromagnetic radiation from other electrons. The emitted wavelength can be tuned by adjusting the electron velocity or the structure of the magnetic field, making the laser tunable.

The free-electron laser is large and complex, but it is tunable across an exceptionally broad range of wavelengths, making it attractive for research. It has become particularly important as a narrow-line source of X-rays for research. Section 11.4 describes it in more detail.

6.3 PUMPING AND ENERGY STORAGE

Transferring energy into a laser medium is called *pumping* or *exciting* the laser. Most lasers are pumped either by light from an outside source or by an electric current passing through the laser medium, but a few other approaches have been demonstrated in the laboratory. Energy storage within a laser depends on the physics of the laser medium as well as the pumping, as explained below.

6.3.1 Optical Pumping

Optical pumping shines light from an external source into the laser medium, which must transmit both the pump light and light emitted on the laser transition. In practice, this means that optical pumping can be used with almost all types of lasers, including gases, solids, semiconductors, and even liquids. The media cannot be perfectly transparent because something in the medium must absorb the pump energy. Usually, the absorber is dispersed within the medium, like chromium atoms in a ruby crystal, dyes in a liquid solvent, or atoms dispersed in a gas laser.

The two main external light sources used for laser pumping are bright lamps and other lasers.

Pulsed flashlamps were common light sources early in laser development because they were readily available and less costly than most laser sources. However, lamp emission is generally much broader than the strongest absorption lines in most materials, so most lamp emission is not absorbed, reducing laser efficiency. The main use of lamp pumping today is where high pulse energies are needed at moderate pulse rates and budgets are limited.

Laser pumping can be much more efficient if the laser emission is well matched to a strong absorption line of the laser medium, and if the pump laser has high wall-plug efficiency. The most common choice of a pump laser is a semiconductor diode laser, which has a high wall-plug efficiency, so the cost per pump photon is low. In addition, *diode pumping* can convert a poor-quality diode laser beam into a much higher-quality beam from a solid-state or fiber laser. Diode pumping took time to catch on because the diodes were expensive, but it has transformed solid-state lasers. Diode pumping has been crucial for the development of new lasers. Diode beams are easy to couple into optical fibers, slabs, thin disks, and other shapes of laser media.

You can also look at diode pumping as a kind of laser *wavelength conversion*, because the light generated is at a different wavelength than the pump diode. Diode-pumped fiber lasers can also be used as pump lasers if pump diodes are not available with the proper wavelength. In that case, optical pumping can become a multistage process, with pump diodes serving as the first stage to pump a fiber laser, which then pumps another laser. Such multistage pumping was not practical until diode pumping became available, but now it may be the most practical way to produce a desired

wavelength. For example, an 808-nm diode laser may pump a solid-state neodymium laser to generate $1.06\ \mu\text{m}$ in the near infrared, and that beam could be frequency-doubled to produce the 532-nm second harmonic to pump an organic dye laser emitting yellow or orange light.

6.3.2 Electrical Pumping

Electrical pumping can be more efficient than optical pumping in some cases because it requires only a single step to convert electric power into optical power. However, it only works in materials that can conduct electricity, limiting its use in gas and semiconductor diode lasers. In fact, as you will learn in Section 8.5, certain semiconductor compounds are optically pumped because it is hard to fabricate electrical connections inside them.

Electrical pumping is done differently in gases and semiconductor diodes because they conduct current differently.

Application of a small voltage in the right direction across a semiconductor diode is enough to get current flowing through the device. About 1.5 volts is enough to turn on a gallium arsenide LED or laser emitting near 800 nm. That low voltage allows input electronic signals to switch the laser or LED on and off or modulate its brightness.

Gases require a much higher voltage to start current flowing because the gas atoms normally do not carry a charge. They must be ionized, which usually requires applying hundreds or thousands of volts across the gas to pull electrons off the atoms, ionizing the gas and starting an electric *discharge*. Electrons accelerated by the high voltage transfer their energy directly or indirectly to the laser species, either exciting the neutral atom to a higher energy level or ionizing the atom, putting it an even higher energy level, as you will learn in Chapter 7.

Some discharge-excited gas lasers produce a series of short repetitive pulses. Pulse length may be limited by the gas pressure; discharges are not stable through a laser tube at atmospheric pressure. Pulse length may also be limited by the internal physics of the laser, which may terminate the population inversion after a certain time. An example is the excimer laser described in Section 7.7.

Other discharge-excited lasers can emit continuous beams if they operate at low pressures where electric discharges are more stable. Examples include the helium–neon and ion lasers in

Sections 7.3 and 7.4. It takes appreciable time on an atomic scale to stabilize such discharges, so those lasers are not modulated directly, but their beams can be externally modulated with a modulator placed in the beam path.

The carbon-dioxide lasers described in Section 7.6 are different because they operate on low-level transitions of the CO_2 molecule, so they do not require high energy to populate the upper laser level. This lets them be much more efficient than atomic gas lasers, which require more energy to reach the laser level.

6.3.3 Energy Storage in Lasers

Lasers store energy in two ways, as light from stimulated emission circulating back and forth between the cavity mirrors and as the laser species excited to the upper laser level, ready to be stimulated to emit more light. How much energy is stored in each way depends on a number of factors, including the concentration of the laser species, the lifetime of the upper laser level, the cross-section for stimulated emission, cavity length, gain in the laser medium, and the feedback within the laser cavity. Some of these dependences are simple, such as the increase in energy storage in the laser medium with the lifetime of the upper laser level and the volume of the laser cavity. Others are more complex, such as the lower energy storage in high-gain lasers because high gain means energy is released rapidly.

If the laser generates a continuous beam, these processes balance out to maintain a steady level of stored energy. That is, the amount of circulating power and the population of the laser species in the upper laser level will remain constant when the input energy offsets the energy lost to emission of the laser beam and other factors. As long as everything remains in balance, the amount of energy stored in a continuous-wave laser remains constant and is not very important.

Changes in conditions inside the laser cavity make energy storage more important because those changes affect how much energy can be extracted. In fact, it is possible to change conditions in the cavity to make the laser briefly emit a pulse of light with peak power much higher than the laser would emit if it was emitting continuously.

If you simply modulate the output of a continuous laser by turning the power off and on, the power will drop from the normal

continuous value to zero, then return to the original continuous level when you switch the power on. That is what happens when you directly modulate the current flowing through an otherwise continuous semiconductor diode laser, which stores very little energy because its cavity is very small and its upper-state lifetime is short.

However, altering a laser cavity which does store energy has a different effect. Instead of turning the pump power off, suppose you put a shutter inside the laser cavity that reduces feedback to below the level needed for the laser to oscillate. Pump energy keeps coming into the laser cavity and exciting a population inversion, but all the spontaneous and stimulated emission inside the cavity is absorbed by the shutter rather than reflected back into the cavity. With the shutter in place, the population inversion keeps growing until it reaches a saturation level much higher than if the laser was running continuously without the shutter in place. However, no power would circulate in the laser cavity with the shutter closed.

Opening or removing the shutter lets the cavity mirror reflect light back into the cavity, increasing feedback and bringing the laser above its oscillation threshold. With light once again able to oscillate in the cavity, spontaneous emission triggers a cascade of stimulated emission, extracting the energy stored in the population inversion in a pulse that depletes the population inversion within a few cavity round trips to end the pulse. The power emitted during the pulse peaks at a level many times higher than the continuous output, and the laser emits much more energy during the pulse than it otherwise could have produced in that much time. Once the pulse is over, the shutter is put back into place to generate another pulse.

This process is called *Q switching* because it changes the *quality* of the laser cavity to produce a “giant pulse” that is very large compared with the laser’s normal output. A typical Q-switched pulse is measured in nanoseconds, and Q switching can be done with lasers that are normally pumped by pulses of light lasting for milliseconds. As with continuous lasers, Q switching millisecond lasers produces giant pulses. But in either case, Q switching has an important limit—it can only extract energy stored for relatively long times on an atomic scale in a long-lived upper laser level. It does not work for laser media with short-lived upper laser levels.

A second process called *cavity dumping* gets around that limitation by switching an optical element to redirect power circulating

in a resonant cavity with a pair of totally reflective mirrors instead of one total reflector and one output mirror. In this case, the laser is operating above threshold, building up energy which is stored by trapping it between two total reflectors in the laser cavity. Redirecting that light out of the cavity switches the circulating power right out of the laser, so all the light leaves in a single round-trip time of the cavity. That generates high-power pulses even if the laser medium has a short-lived upper laser level.

6.4 LASER PULSE CHARACTERISTICS

We almost instinctively realize that a laser's output power is one of its crucial features. It is not as obvious that the duration of emission is another crucial feature of lasers. Continuous beams work best for applications that require steady illumination, from defining a straight line for a construction project to welding together two sheets of metal. Pulsed beams work best for applications that require brief but intense pulses, from measuring distances in laser radars to drilling holes through brittle materials.

Pulse duration and repetition rate can vary widely, and the specifics can be very important for laser applications. A series of short but powerful pulses can drill holes, with pulses separated long enough for the hot plasma from one pulse to clear away so it does not block the light from the next pulse. Intense picosecond pulses can ablate material so fast that their heat does not damage the layer just underneath, making them useful for cutting glass. Let us look at the various modes of laser operation in time.

6.4.1 Continuous-wave Laser Emission

As the word implies, a continuous-wave laser emits a beam continually, at least on an atomic time scale. That means the laser may emit for only seconds at a time, but it should be long enough for its output to stabilize at a steady level, which could take much longer for some lasers, notably the metal vapor lasers described in Section 7.5. Many continuous-wave lasers can emit stably for hours, days, or longer.

Many lasers can emit continuously, but some cannot because of their internal physics. Examples include nitrogen and excimer gas lasers described in Chapter 7.

6.4.2 Lasers Pulsed by Power Sources

Pulsing of some lasers is controlled by their power sources. One example is driving a semiconductor diode laser with series of electrical pulses to transmit a digital signal through a fiber-optic data link.

Other lasers are optically pumped by pulsed lamps or pulsed lasers. Many solid-state lasers are pumped by flashlamps which fire pulses that last milliseconds. Without other controls, the solid-state laser would fire millisecond pulses lasting a bit less than the duration of the lamp pulses because it takes time for the pump energy to bring the laser to threshold.

Gain within the laser cavity can be increased by applying a strong pulse of pump energy. Rapid amplification within the laser medium can produce pulses in laser output that are shorter than the pump pulse in suitable gas, semiconductor diode, and optically pumped solid-state lasers.

6.4.3 Quasi-continuous-wave Laser Operation

Some lasers are pulsed so often that their operation may seem nearly continuous-wave. These are called *quasi-continuous-wave*, and they may be operated that way to meet application requirements or to overcome thermal limitations.

Quasi-continuous-wave operation allows arrays of diode lasers to briefly emit at higher powers than they could emit continuously. Figure 6-2 shows how this works, with an array that emits high-power pulses that last about one-tenth of the time, then is off the remaining nine-tenths of the time. That makes the average power about one-tenth of the peak power. Concentrating that much energy into short intervals is an advantage for laser materials working



Figure 6-2. A quasi-continuous laser emits for short intervals than switches off for lower intervals to dissipate waste heat. Typical diode arrays emit individual pulses lasting 0.05 to 50 milliseconds.

because the brighter pulses ablate material more efficiently. The long off intervals help the array dissipate waste energy more efficiently than it could during continuous emission.

Such off-on operation is called a *duty cycle*, with the laser typically emitting 1% to 10% of the time, and off the rest. In diode lasers, it is produced by turning the electric power supply on and off. Pumping with quasi-continuous-wave diode lasers can generate quasi-continuous-wave output from solid-state lasers. Generally, individual pulses last 10 μs to 50 ms.

Much shorter pulses can be used at much higher rates to produce quasi-continuous-wave output for other applications. For example, a frequency-tripled solid-state laser emitting at 355 nanometers in the ultraviolet fires 100 million pulses per second, each lasting less than 10 ps, corresponding to a duty cycle of 0.1%. It uses modelocking, a process described in Section 6.4.6, and the modelocked pulses are so close together that they blur together when used to measure response of a slow system.

6.4.4 Q Switches and Nanosecond Pulses

In Section 6.3.3, you learned how a process called Q switching produces pulses by switching the *quality factor* or *Q* of a laser cavity switching between two states, one that stores energy in a laser cavity and another that extracts energy. Q switching collects energy over a relatively long time and then releases it quickly, concentrating laser energy to produce pulse with peak power much higher than available from a continuous beam. The ability of Q switching to produce such powerful pulses lasting from fractions of nanoseconds to hundreds of nanoseconds has made it a powerful tool for laser applications. Let us look at the process more carefully.

The central idea of Q switching is to make a laser store its energy over time rather than releasing it continually, like trying to build up an electric charge to release in one big spark. This requires keeping the laser operating below threshold to prevent oscillation. The easiest way to do that is to reduce feedback from the cavity mirrors by putting something in the way so light does not reach one of the mirrors. That something is a Q switch, which can be moved out of the way or made transparent to allow laser oscillation to start and the accumulated light energy to escape in a giant pulse through the output mirror. Generally, it sits between the totally reflecting back

mirror and the laser rod, with the output mirror on the other end of the rod.

Q switching only works with laser media that can store energy in the laser cavity when oscillation is suppressed. That means that atoms excited to the upper laser level must have a long spontaneous emission lifetime so they remain in that state a long time on an atomic time scale. In practice, Q switching works best with optically pumped solid-state lasers, many of which have spontaneous emission lifetimes of microseconds.

Technically, what the Q switch does is to change the Q factor or quality factor of the resonant cavity, which is low when energy is being stored below laser threshold, and high when the Q switch changes state to produce a pulse. The Q factor for a low-loss resonator depends on the product of the laser frequency ν and the round-trip time t_{rt} divided by the fraction of light lost per round trip L , or

$$Q = \frac{2\pi\nu t_{rt}}{L} \quad (6-1)$$

In practice, the laser frequency and the cavity round-trip time are fixed for any given laser, so the easiest way to alter cavity Q is by changing the round-trip loss L , which is the job of a Q switch. Note that because the laser frequency is a high number (300 terahertz or 3×10^{14} Hz for a 1- μm laser), the typical Q of a laser cavity also tends to be high—about 5.6 billion for a round-trip loss of 1% and time of 30 ns (as for an 8-cm rod of neodymium-doped crystal). A Q switch essentially blocks all the light, dropping the Q by a very large factor.

To show you how Q switching works, Figure 6-3 steps through the generation of a pulse from an optically pumped solid-state laser. It starts by turning on the pump light with the Q switch set to block light. This excites atoms which drop down into the upper laser level, where they will remain around a millisecond if left undisturbed as shown at the top. Suppose it takes 500 μs for the upper-state population to reach the desired level. Then the Q switch turns on and the pulse begins to grow, as shown in the next step. The pulse peaks and begins to decline as more of the stored energy is depleted. Finally, after essentially all the energy is extracted, output drops to zero, and the Q switch turns off to begin building up energy for the next pulse, as shown at the bottom.

A typical pulse lasts longer than a cavity round-trip time because it takes time to build up stimulated emission from all the




Figure 6-3. Q switch releases energy accumulated in a laser cavity, producing pulses lasting from less than a nanosecond to tens of nanoseconds. The intervals between pulses are much longer than the output pulse, typically ranging from 10 μ s to 1 ms.

atoms in the upper laser level. Pulse durations range from under a nanosecond to tens of nanoseconds, with repetition rates typically 1 to 100 kilohertz. The duration of a Q-switched pulse depends on output-mirror reflectivity R and on the time t that laser light takes to make a round trip of the laser cavity, and is roughly

$$\text{Pulse duration} = \frac{t}{(1 - R)} \tag{6-2}$$

If you insert the value of cavity round-trip time—twice the cavity length L divided by the speed of light in the laser medium (c/n)—you have a more useful estimate of Q-switched pulse duration:

$$\text{Pulse duration} = \frac{2Ln}{c(1-R)} \quad (6-3)$$

A Q switch can turn on once to produce a single shot, but usually it switches on repeatedly to generate a series of regularly spaced pulses.

6.4.4.1 Active Q Switches

So far, we have not stated anything about how actual Q switches work. That is because Q switching can be done in several ways. Three major types of Q switches are called *active* because they require active control, and input energy, in order to change the Q in the laser resonator. The fourth does not require power and is described in Section 6.4.4.2

The simplest type of active Q switch to understand is a rotating hexagonal mirror, shown in Figure 6-4. The figure looks down on the mirror spinning on its axis. When one face of the spinning

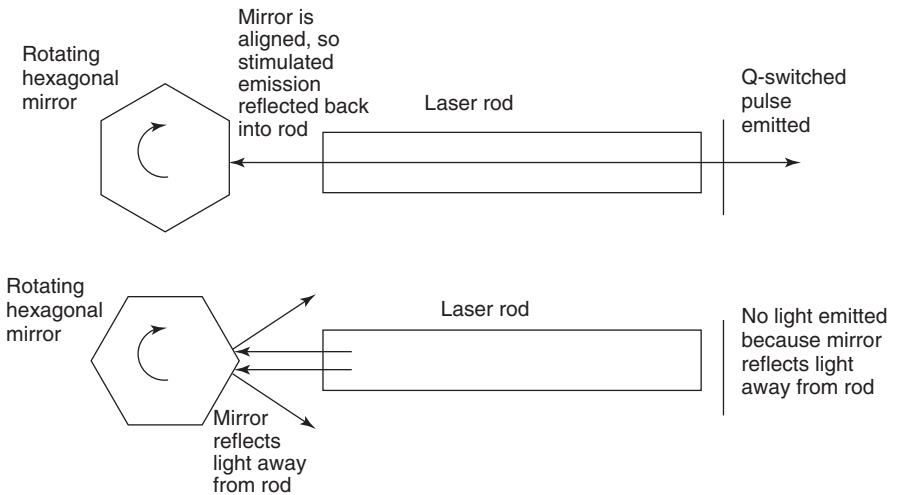


Figure 6-4. Mechanical Q switch releases energy accumulated in a laser cavity when the spinning hexagonal mirror lines up one face so it reflects light directly back along the length of the laser rod, producing feedback that generates a Q-switched pulse.

mirror aligns perpendicular to the end of the laser rod, as shown at top, it reflects light back into the rod, creating feedback into the laser cavity, which stimulates emission of light from the laser species and triggers a pulse that lasts as long as the feedback is high. After the mirror turns away, it starts reflecting light from the laser rod off to the side, reducing feedback to zero and stopping the Q-switched pulse, as shown at the bottom. Such mechanical Q switches are not used much today, but they illustrate the idea of Q switching well.

Other types of Q switches are based on the acousto-optic and electro-optic modulators described in Section 5.7.2. An acousto-optic Q modulator is transparent when no signal is applied to it, but applying an acoustic signal diffracts light that otherwise would pass straight through the modulator. Thus, an acousto-optic modulator can work as a Q switch if it is switched between the low-loss and high-loss states, making the cavity Q high in the high-transparency state to produce a pulse, and reducing the cavity Q when a signal diffracts the laser beam to increase cavity losses. Acousto-optic Q switches are widely used because of their low cost. However, they have important performance limitations. They do not diffract all the light from the beam, so they may not shut down laser emission totally, and they require relatively high power because they need to be diffracting light during the long intervals between Q-switched pulses.

Electro-optic Q switches are much more expensive and require much higher drive voltage than acousto-optic modulators. Their advantages are speed and performance, based on changing the polarization of input light to control how much of it is transmitted. Although often considered higher-performance Q switches, they are not used as widely as acousto-optic Q switches.

6.4.4.2 *Passive Q Switches*

Passive Q switches take advantage of an effect called *optical saturation* to avoid the need for expensive active control systems or power supplies to drive them. Instead, they rely on saturation, sometimes called *bleaching*, to make a normally dark material transparent when exposed to intense light.

Under normal light, the saturable material looks dark and strongly absorbs light. However, the absorption of light excites the dark light-absorbing material into a higher-energy state that does not absorb light at the same wavelength, so the material becomes transparent, at least at that wavelength. It is the inverse of photochromic

sunglasses, which are transparent indoors, but when exposed to bright ultraviolet light outdoors absorb the ultraviolet light and become dark enough to serve as sunglasses. Various types of saturable absorbers may be used in passive Q switches.

Certain organic dyes can be used as passive Q-switch materials because their dark states easily saturate make them transparent. The passive Q switch starts dark, so when the laser medium is pumped and begins producing spontaneous and stimulated emission, the dye soaks it up and remains dark. But as the light level increases, all the dye molecules get excited to the higher-energy transparent state, leaving the material clear and opening the way for the energy built up in the laser medium to stimulate more emission and start oscillating. Once the stored energy escapes, the light is no longer intense enough to keep bleaching the dye, so it becomes opaque enough to prevent oscillation.

Solid-state saturable absorbers are crystals and glasses doped with materials that have saturable absorption at important laser wavelengths. Chromium in the +4 state is saturable in the 1- μm band of neodymium and erbium atoms. Vanadium in the +3 state can be used in the 1.3- μm band. Materials doped with cobalt and lead sulfide quantum dots are used in the 1.5- μm band. Semiconductors can also be engineered for use as saturable absorbers, with their wavelength range depending on their composition. These solid materials work much like dye-based passive Q switches, but are more convenient and less vulnerable to degradation over time.

These properties give passive Q switches an attractive simplicity, avoiding the need for active switching and reducing costs because they can be engineered to switch between transparent states all by themselves if the laser is being powered. However, they also come with important limitations. Their lack of active control makes synchronization difficult to impossible and makes them vulnerable to jitter. The dyes used in passive Q switches degrade over time. Performance and cost differ among materials and wavelengths, leading to tradeoffs between among different passive Q-switch materials, as well as between the low cost and simplicity of passive Q switches and the control and performance of active Q switches.

6.4.4.3 Q Switching in Lasers

Q switching depends on the choice of laser material as well as the Q switch. Optically pumped bulk solid-state lasers are generally

preferred because their large volumes and the long lifetimes of their upper laser levels combine to store large amounts of energy that can be Q-switched out. The smaller volumes of fiber lasers limit how much energy they can store for Q-switched pulses.

For solid-state lasers, upper-state lifetimes depend on the host material as well as the light emitter. For neodymium-doped crystals, they range from 1.2 ms for Nd–YAG to 50 μ s for Nd–YVO₄, with glass roughly halfway between. Atomic gas and semiconductor diode lasers work poorly with Q switches because their upper-state lifetimes range from a few nanoseconds to hundreds of nanoseconds. The 10.6- μ m CO₂ laser is an exception, because its upper state has a lifetime longer than a second, but its long wavelength requires special Q-switch materials.

In practice, the Q switch is integrated with the laser to optimize performance, and Q-switched lasers are sold as complete packages. Their big commercial advantage is their ability to generate high-energy nanosecond pulses. Microchip lasers based on tiny diode-pumped solid-state lasers can also be Q switched to generate pulses lasting 100 to 1000 ps, usually with passive semiconductor Q switches.

Later chapters will describe more about how Q switching is used on various laser types.

6.4.5 Cavity Dumping

Like Q switching, cavity dumping works by releasing energy stored in the laser cavity, but the details differ in important ways. Cavity-dumped lasers have totally reflecting mirrors on both ends of a high-Q resonant cavity. The cavity also contains a third optical element that normally is out of the optical path but can be switched into place to redirect the circulating light out of the reflective cavity to a separate output port.

Cavity dumping releases all the light energy circulating within the laser cavity in a single pulse lasting one round trip of the laser cavity:

$$\text{Pulse duration} = \frac{2Dn}{c} \quad (6-4)$$

where D is cavity length, n is the refractive index of the laser medium, and c is the speed of light in a vacuum. Thus, a 30-cm

long laser with a refractive index of 1 would produce pulses lasting 2 ns. This means that cavity dumping extracts light that is circulating within the cavity, in contrast to Q switching, which extracts the energy stored in a population inversion contained in a cavity that had been maintained at a low Q cavity where laser oscillation had been suppressed by high loss. This has important consequences both in the type of pulse generated and in the lasers which can be cavity dumped.

The easiest way to visualize cavity dumping is as a single mirror that pops up into the beam path and deflects the intracavity beam out of the cavity, as shown in Figure 6-5. In practice, however, cavity dumping normally is done by switching acousto- or electro-optical devices from a transparent state to one that reflects light out of the cavity.

Cavity dumping essentially yields a chunk of intracavity light from a continuous-wave laser operating at a stable power level




Figure 6-5. Cavity dumping can be visualized as a mirror popping up into the beam path and directing light out of the laser cavity. However, cavity dumping usually uses electro- or acousto-optical switches to redirect the light.

without coupling any light through its output mirror. Think of it as switching on a continuous-wave laser for exactly one cavity round trip, but the switching is not instantaneous so it is not a square wave that rises instantly from zero to full power and then drops back instantly to zero. Both total energy and peak power are lower with cavity dumping because Q switching extracts energy from all the atoms in the upper laser level, not the light they had already emitted. Cavity-dumped pulses are also shorter because it takes more than a single round trip of the laser cavity to stimulate emission from the whole excited-state population in Q switching.

Cavity dumping can serve several purposes. One is to generate pulses with different shapes or shorter lengths than available from Q switches. Another is to generate short pulses from lasers unsuitable for Q switching because of issues such as short upper-state lifetimes. Cavity dumping can also be combined with Q switching or modelocking to select pulses or to improve performance.

6.4.6 Modelocking and Ultrashort Pulses

Modelocking generates the shortest pulses that can be produced in a laser cavity, from picoseconds all the way down to femtoseconds. It gets its name from the idea of locking many longitudinal modes together so the laser oscillates simultaneously on all of them. But it may be simpler to understand by starting with the goal of the process, producing a very short clump of photons that bounces back and forth in a laser cavity, emitting a pulse of light every time the pulse hits the output mirror.

The idea of modelocking goes a big step beyond Q switching, but uses similar but much faster active and passive devices to modulate light inside the laser cavity. As in Q switching, a modelocked laser normally is pumped continuously to produce a population inversion, and the modulating device switches on briefly to allow light to circulate. However, for modelocking the modulator becomes transparent only briefly once every time light makes a round trip of the cavity.

For example, consider a laser cavity 30 cm long with refractive index of 1, so the cavity round-trip time is 2 ns. The modelocking modulator then might switch on for 1 ps (1/2000th of the round-trip time) every 2 ns. The rest of the time it would have high enough attenuation to keep the laser below threshold. Once things got started, this would produce a clump of photons bouncing back

and forth within the laser cavity every 2 ns, at a repetition rate of 500 MHz.

The process would be the same whether the modulator was an active acousto- or electro-optical device or a saturable absorber, which saturated only briefly during the peak of the pulse when a large number of photons bleached the material. But the process is not that simple because modelocking only works if the laser meets certain requirements.

The requirements come from a basic limitation of quantum mechanics called the *transform limit* described in Section 4.6.2. That says that the product of the duration of a light pulse Δt and the range of frequencies it contains $\Delta\nu$ must be at least a minimum value—0.441. So the shorter the pulse the laser produces, the broader the range of frequencies the laser must emit.

As you learned in Section 4.2.3, a laser cavity resonates on many longitudinal modes, depending on the gain bandwidth of the laser material and the length of the laser cavity. The broader the gain bandwidth, the more modes. The transform limit tells us that a laser needs a broad gain bandwidth to produce ultrashort pulses. That means that only laser materials with broad gain bandwidth are suitable for modelocking, and that such lasers must be operated in laser cavities with resonator mirrors that reflect light across a range of wavelengths.

Each longitudinal mode that oscillates in a laser has its own resonant wavelength, and those modes are equally spaced in frequency. That means that the modulation device is locking the modes together so they all stay in phase with each other—giving the process its name of modelocking. That also means that an evenly spaced time series of modelocked light pulses is equivalent to a series of continuous light waves oscillating at uniformly spaced light frequencies. (Evenly spaced in frequency does not mean evenly spaced in wavelength, although the difference is small if the spacing is close.)

You may wonder how a series of pulses in time can also be a series of continuous waves uniformly spaced in frequency. That arises from constructive and destructive interference among the frequencies. Add waves of those frequencies together in phase, and they add together to produce a series of uniformly spaced pulses. Mathematically speaking, the series of pulses are a Fourier transform of the spectrum of electromagnetic waves. So it all fits together, as shown in Figure 6-6.




Figure 6-6. Modelocking locks together many longitudinal modes in a laser cavity, so adding the waves together coherently produces a series of ultrashort pulses uniformly spaced by one cavity round trip. (Only a few modes are shown; modelocking actually locks together many more modes that are small compared with the cavity length.)

6.4.7 Self-Terminating Lasers

Certain lasers can only operate for less than a fraction of a second because of their internal physics. These lasers emit for brief intervals, then turn themselves off because of quirks in their physics. They are called *self-terminating lasers*, and they can be stopped by different processes.

One limitation is an inability to depopulate the lower laser level during laser operation, building up the population there and ending the population inversion. That can be an issue in three-level lasers, such as ruby, that have the ground state as their lower laser level. It can also be a problem in four-level lasers if the lower laser level has an inherently long lifetime for other reasons, as in the copper-vapor lasers described in Section 7.5.2. Three-level lasers can be operated if they are pumped hard enough to keep more than half the laser species in higher energy states. Other mechanisms

such as buffer gases can be added to get the laser species out of the lower laser level in four-level systems.

Another limitation is the difficulty of sustaining high-voltage electrical discharges in gases near atmospheric pressure. Those discharges can break down in nanoseconds, limiting the excimer lasers described in Section 7.7.

6.5 WAVELENGTH CONVERSION

So far we have described laser wavelength as a fixed feature of particular lasers. Some lasers have gain across a range of wavelengths, so their output wavelength can be tuned by adjusting their cavity optics or operating conditions, and others can emit on multiple transitions, but the gain bandwidth and transitions remain fixed. However, there are ways to change the wavelength of laser light. The two most important approaches are the optical pumping described in Section 6.3.1 and use of the nonlinear optical effects described in Section 5.6.

6.5.1 Importance of Wavelength Conversion

What makes wavelength conversion important is the rarity of good laser transitions. Laser action has been demonstrated on countless thousands of transitions, but only a few are used in commercial lasers. In part, that reflects the nature of physics; different materials have different optical properties, and some simply make better lasers. It also reflects the way technology works; investments of ingenuity, time, and money pay off in more knowledge and improved material properties. These two factors interact to shape laser performance.

The use of neodymium in solid-state lasers is a good example. The first neodymium lasers were demonstrated in 1961, a year after the first laser. They emitted nearly $1.06\ \mu\text{m}$ in the near infrared, could be easily excited by readily available lamps, and were reasonably efficient. Companies invested in developing crystals and types of glass that could serve as hosts for neodymium lasers, and that investment made neodymium lasers the solid-state lasers of choice for decades. Other solid-state lasers might be almost as good, but it was not worth investing much money in improving them unless they would make much better lasers.

When diode pumping began to replace lamp pumping, engineers found that neodymium could be pumped efficiently with semiconductor diode lasers emitting at 808 nm, and by good fortune that wavelength was easy to produce from readily available gallium arsenide diode lasers. Only later did ytterbium come into favor, when new diode lasers were developed that could pump ytterbium more efficiently than 808-nm diodes could pump neodymium. Ytterbium also benefited from interest in fiber lasers because it worked better in fibers than neodymium. In addition, ytterbium emitted at nearly the same wavelength as neodymium, so it could be used in many of the same applications.

That similarity also became a limitation of solid-state lasers. The best of them emitted in a narrow band between 1.03 and 1.08 μm in the near infrared. A few weaker solid-state lasers were available in the near infrared, but nothing very worked well at visible wavelengths except ruby, and that was limited to pulsed operation. For many years, gas lasers were the only source of continuous beams at visible wavelengths; red semiconductor diode lasers were not available until the mid-1980s. Only a few lasers could generate yellow, green, or blue light, and through the 1980s they remained expensive and often horribly inefficient, leaving few lasers in the visible spectrum.

Laser technology has progressed since then. Blue semiconductor diode lasers were a breakthrough in the 1990s. I remember the thrill in a packed auditorium when Shuji Nakamura showed the first blue diode laser ever shown in America. But new applications keep emerging that need specific new wavelengths not readily available from any standard commercial laser.

Now technological advances are filling these gaps in two different ways, often combined. One is to use readily available commercial lasers, typically semiconductor diodes, to optically pump a laser or laser-like source to emit the desired wavelength. The other is to use the nonlinear optical techniques described in Section 5.6 to convert the output of a laser to a different wavelength, such as generating the second harmonic to convert the 1064-nm near-infrared output of a neodymium laser to 532 nm, producing a bright green beam. Combining these techniques, optically pumped semiconductor lasers (OPSLs) are excited by light to produce near-infrared wavelengths not available from electrically pumped diode lasers or any other type, then the second harmonics of those beams are generated to shift the output to visible wavelengths unavailable




Figure 6-7. Wavelength conversion may go through multiple steps to generate a new wavelength. In this example, a semiconductor diode pump laser emitting at 900 to 1000 nm in the near infrared pumps an optically pumped semiconductor laser (OPSL), converting much of the pump power 1154 nm. Then a harmonic generator doubles the laser frequency, converting most of the power from 1154 to 577 nm, a yellow wavelength hard to make with conventional lasers.

from any other source, as shown in Figure 6-7 for production of 577 nm yellow light. Some energy is lost in the process, but it is more efficient than using a tunable dye laser to produce the desired wavelength.

6.5.2 Optical Pumping for Wavelength Conversion

The attraction of optical pumping for wavelength conversion is that it converts cheap photons from readily available pump lasers into more valuable photons at hard-to-get wavelengths. If the pump source is a diode laser, the optical-pumping process also generates a beam of much higher optical quality, important for many applications such as laser cutting and welding.

All laser applications are sensitive to some extent to beam quality, but their sensitivity varies widely. Some materials-working applications may work equally well with lasers at 950 or 1060 nm if the beam quality is the same. However, some sensing applications require a specific wavelength to measure the concentrations of materials properly. Tunable lasers can be used for some wavelength-sensitive applications in the laboratory, but they can be complex to operate, so it can be worth the trouble to go through multiple optical-pumping steps to get the desired wavelength. For example, holmium-doped fiber lasers are good sources of light around 2100 for some medical applications, but they require pump light at 1950 nm. The best source of that wavelength is a thulium-fiber laser that requires pumping at 1600 nm. The best source of 1600 nm is an erbium-doped fiber laser that is pumped by diode lasers at even shorter wavelengths. So the most practical way to produce the desired wavelength requires three optical-pumping steps.

Another aspect of diode pumping is that it can produce the high-quality beams needed for the nonlinear optics processes used to change wavelengths by harmonic generation and other processes. That is what happens in converting near-infrared lasers to visible light, as shown in Figure 6-7. Harmonic generation is far more efficient with the high beam quality of an OPSSL than with the poor beam quality of a standard semiconductor diode laser.

6.5.3 Nonlinear Optics for Wavelength Conversion

The nonlinear optical effects described in Section 5.6 offer another way to convert laser output to other wavelengths. The simplest of these effects is second-harmonic generation or frequency doubling, which in one step converts the fundamental wavelength normally emitted by a laser into light at half that wavelength. One example is the common green laser pointer, which includes a harmonic generation crystal that converts the 1064-nm neodymium line in the near infrared to 532 nm in the green, as shown in Figure 5-12. Although the harmonic generation crystal could be called an accessory, it is normally built into the laser, and users normally see the whole assembly as a green laser pointer, without realizing that the laser inside is emitting infrared light that is converted into green light.

Harmonic generation and other nonlinear processes work best with lasers that generate high power in a tightly focused beam. That comes from the fact that strength of the second-order nonlinearity depends on the square of the laser intensity. Double the laser intensity, and the strength of the nonlinear effect increases by a factor of four. Increase the laser intensity by a factor of 10, and the strength of the nonlinear effect rises by a factor of 100. Concentrating laser energy into short pulses works as well as focusing the laser beam more tightly or increasing the power in a continuous beam. Multiplying all those factors together works even better, so some of the strongest nonlinear responses come from short-pulse, high-power lasers. That means the best results typically come from pulsed solid-state lasers, which can combine short pulses with high energy and intensity. Semiconductor diode lasers fall short because their beams are not tightly focused, so they cannot reach the high intensities needed for efficient harmonic generation.

Near-infrared lasers have another advantage in that doubling the wavelengths of lasers emitting at 800 to 1400 nm generates

second harmonics in the 400- to 700-nm visible range, where laser sources are scarce. Often a second-harmonic generator is packaged with a near-infrared solid-state laser and sold as a visible diode-pumped solid-state laser delivering watts of continuous power at selected wavelengths. This has the advantage of making many wavelengths available at watt-level powers that otherwise would not have been available at reasonable cost. These wavelength-shifted solid-state lasers have largely replaced many of the visible-wavelength gas lasers described in Chapter 7.

Lasers packaged with wavelength-conversion optics are widely used because it makes them as simple to use as other lasers emitting at their normal emission lines. A green laser pointer with internal wavelength-conversion optics is just as easy to use as a red laser pointer built around a red-emitting semiconductor diode laser. Ease of use is important because nonlinear optics often require careful alignment to obtain the best performance.

Although third- and higher-order harmonics can be produced directly, it is more efficient to combine second-harmonic generators with sum-frequency generators in series, as shown in Figure 6-8. The most efficient way to generate the third harmonic is by adding the first harmonic to the second. The most efficient way to produce the fourth harmonic is putting two frequency doublers in series, with the first doubling the fundamental frequency, and the second doubling the second harmonic to produce the fourth. The

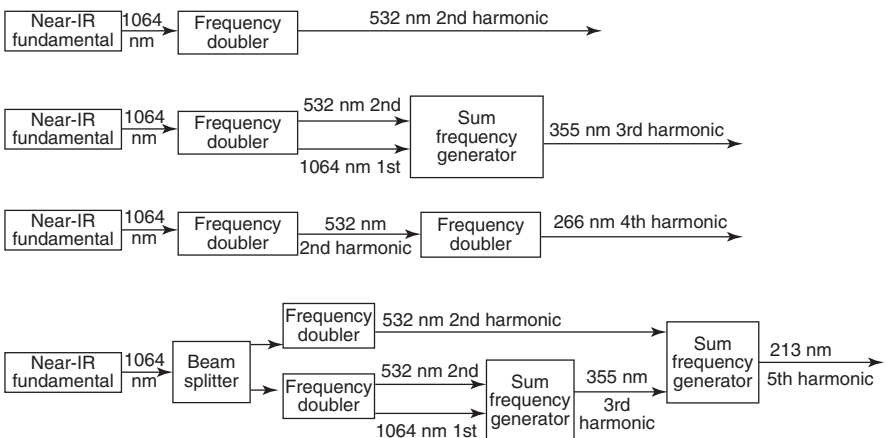


Figure 6-8. Generating second through fifth harmonics from a near-infrared laser by combining frequency doubling with sum-frequency generation.

fifth harmonic can be generated by adding the second and third harmonics.

The strength of nonlinear effects depends on the internal structure of the material. Nonlinear effects are very weak in most materials such as ordinary glass, but as described in Section 5.6.5, nonlinear effects are much higher in crystals with high levels of internal strain.

Harmonic generators and sum- and different-frequency generators are also available as discrete components that can be added to lasers. However, they can require careful adjustments such as phase-matching of the fundamental and second-harmonic wavelengths to best performance.

Parametric amplifiers and oscillators are available as standard commercial products, as described in Section 11.2. These are highly sophisticated systems that have become highly automated in recent years, greatly increasing their ease of use.

6.5.4 Raman Shifting of Wavelength

The Raman scattering process described in Section 5.6.4 can be used to shift laser wavelength. The Raman shift is measured in frequency units, and in most cases, it reduces the photon frequency and energy, shifting it to longer wavelengths. Like harmonic generation, it is a nonlinear process, which increases with the square of the power density.

The Raman shift measured in energy or frequency is an inherent characteristic of the material's energy levels and is the same for a particular material for any illumination wavelength. For silica, the Raman shift peaks at about 13 terahertz (13,000 GHz), so it would shift the 282-terahertz frequency of a 1064-nm neodymium laser to 269 THz, or 1144 nm, a wavelength otherwise difficult to produce which can be frequency-doubled to make yellow light at 572 nm. Other Raman-active materials have different Raman shifts.

Raman wavelength shifters can be integrated with near-infrared or visible lasers to generate new wavelengths or can be built as accessories to use with other lasers.

Another option is using Raman-active materials as the gain medium in Raman lasers or oscillators. These are not exactly lasers because they do not store energy in a population inversion, but they produce laser-like beams and otherwise act like lasers. Raman gain media also can be attached to lasers to produce new

wavelengths. Raman lasers and amplifiers are usually fiber-optic devices and are covered in Section 9.6.

6.6 LASER OSCILLATORS AND OPTICAL AMPLIFIERS

So far, we have described lasers that generate light on their own, with no external input other than pump power. Properly speaking, those are laser *oscillators*. Laser amplifiers are something subtly different, and that difference is important in the laser world.

It can also be confusing because the acronym “laser” specifies only light *amplification* by the stimulated emission of radiation. In practice, the usual definition of “laser” includes the requirement for internal feedback so it reaches laser threshold and oscillates, generating its own beam. In contrast, an *optical amplifier* amplifies light from an external source and lacks mirrors that form a resonant cavity.

This difference leads to a fundamental difference in behavior. The gain of the laser material and the reflective properties of the laser cavities combine to determine the wavelength of a laser oscillator, but the wavelength of an optical amplifier depends on the input signal that stimulates emission. Laser oscillation is relatively easy to measure because optical power is very low below threshold, but increases rapidly in power and decreases in line width above threshold. Amplification is more difficult to measure, particularly if the gain is low.

Early developers focused their attention on laser oscillators because they wanted devices that generated beams internally. Optical amplifiers (sometimes called laser amplifiers) can be used to boost the power available from pulsed laser oscillators, but today the main use of optical amplifiers is to boost the signal strength in fiber-optic communication systems.

6.6.1 Laser Oscillators

As you learned in Chapter 3, a conventional laser is an oscillator, with mirrors on two ends forming a resonant cavity. Laser oscillation starts when a spontaneously emitted photon stimulates emission from another excited atom. If the original photon and the stimulated emission are aimed along the axis of the cavity (or within the narrow range of angles reflected by the cavity mirrors),

they bounce back and forth between the mirrors, with some light transmitted through the output mirror when it reflects the rest of the beam back into the cavity. To sustain oscillation, the power must increase enough during a round trip of the laser cavity to offset the power lost within the cavity and the power exiting through the output mirror. A laser oscillator is the optical counterpart of a radio oscillator as a signal source.

A laser oscillator can produce pulses or a continuous beam. In either case, the laser output builds quickly as light oscillates back and forth between the mirrors, stimulating emission from excited atoms or molecules. From an electronic standpoint, the mirrors provide feedback to stabilize laser wavelength. In many pulsed lasers, light emission depletes the population inversion, terminating the pulse. In other pulsed lasers, internal components such as the Q switches described in Section 6.4 modulate gain, switching oscillation on and off. In a continuous laser, an external energy source continually excites the laser medium, and the power stabilizes at a level where gain in a round trip of the cavity offsets the power emitted or lost in the cavity.

Light circulates back and forth between the mirrors in a laser oscillator, sustaining the oscillation. Individual photons bounce back and forth until they are absorbed or emitted through the output mirror. The oscillation wavelength is selected by the optics of the resonant cavity, and by the nature of the laser medium.

6.6.2 Optical Amplifiers

An optical amplifier is a rather different device, as shown in Figure 6-9. It is a single-pass device that is transparent on both ends rather than reflective.

The input in the figure arrives as a single pulse, which is amplified as it stimulates emission inside the amplifier. The pulse grows larger as it passes through the amplifier, eventually reaching a maximum before exiting. The exit window should reflect no light back through the amplifying medium or optics should be added to prevent amplification.

This means that an optical amplifier and a laser oscillator serve fundamentally different functions. A laser oscillator generates a signal at a desired wavelength. An optical amplifier amplifies light from an external source, which makes only a single pass through the amplifier stage. The amplifier may be a tube of laser gas, a slab




Figure 6-9. Light oscillates back and forth in the laser oscillator at top. Light makes only a single pass through the optical amplifier at bottom.

or cylinder of a solid-state laser material, a semiconductor chip, or an optical fiber doped to amplify light passing through its core.

The two major uses of optical amplifiers are for producing pulses more powerful than available from an oscillator and to amplify signals in fiber-optic communication systems.

6.6.3 Master Oscillator Power Amplifier (MOPA) Configurations

Internal dynamics limit the peak power a laser oscillator can generate in a pulse, but external amplifier stages can amplify that pulse to higher powers. This arrangement is called a *MOPA* and essentially separates the generation of a pulse from its amplification.

The master oscillator produces the input pulse, which then makes a single pass through each amplifier in a chain. The master oscillator output should be at a wavelength where the amplifier stages have high gain.

Most MOPAs are solid-state lasers operated at relatively low repetition rates. To generate very high pulse powers, amplifiers are arranged in series, so the pulse passes through a sequence of amplifiers with successively larger gain volume, which can produce higher powers. The beam may be spread over larger areas in later amplifier stages to generate more energy. Thus, the pulse may pass




Figure 6-10. Master oscillator power amplifier puts a pulse through a series of larger amplifiers to increase pulse power.

through a small rod, then be expanded to pass through a larger rod, and then be expanded again to pass through a slab or disk, as shown in Figure 6-10.

Amplifier shapes are chosen to limit power density so it does not damage the output end of the amplifier, which is the most vulnerable area. Disks or slabs are used at high powers to aid in dissipating excess heat.

The MOPA design is widely used in lasers with very high pulse power, such as the world's biggest laser, the National Ignition Facility (NIF) at the Lawrence Livermore National Laboratory. In NIF and other extremely high-power systems, long cooling times are needed between pulses.

6.6.4 Chirped-Pulse Amplification

Another important use of optical amplifiers to generate high-power pulses is in *chirped-pulse amplification*. The idea grew from efforts to generate extremely high peak power in extremely short pulses without damaging optical materials and surfaces, and earned its inventors, Donna Strickland and Gerard Mourou, the 2018 Nobel Prize in Physics.

To avoid damage, chirped-pulse amplification stretches short pulses over longer periods of time and then recombines them and compresses the pulse in time. This is done by first passing the short pulse through optics which delay longer (or shorter) wavelengths to extend the pulse length, reducing the peak power to manageable levels that will not cause optical damage when amplified.

After the light passes through the amplifier, the longer amplified pulse is passed through another set of optics that compress it by delaying the wavelengths that led the dispersed pulse until

the other wavelengths catch up. This squeezes the pulse to a much shorter duration with a much higher power. In this way, pulse duration can be increased by a factor of a thousand or more for amplification and then decreased by the same factor (or more) to multiply peak output power by a corresponding amount. This technique is at the heart of lasers that have reached peak powers greater than a petawatt (10^{15} watts) in ultrashort pulses.

6.6.5 Optical Signal Amplifiers

Low-power optical amplifiers boost the strength of weak optical signals in fiber-optic communication systems. In this case, the input signal makes a single pass through the length of an optical amplifier.

Optical amplifiers are important for this application because they can amplify signals across a range of wavelengths in the gain band of the active material. Optical fibers can simultaneously transmit signals at many different wavelengths, which optical amplifiers simultaneously amplify. Thus, if the input contained separate signals at 1551, 1552, 1553, and 1554 nm, an optical amplifier would increase the strength of all four signals.

This allows optical amplifiers to serve as signal boosters on long-distance fiber-optic communication systems. The amplifier takes the weak signals delivered by an input fiber, increases their power, and then transmits them through the next span of fiber to the next amplifier. Optical amplifiers can also boost the strength of an optical signal before a receiver or raise the strength of the signal from a transmitter at the start of a long length of fiber.

Three types of optical amplifiers are used in communications. The most common is an optical fiber with its light-carrying core doped with erbium, which functions much like a fiber laser and is covered in Section 9.5.

A second is an optical fiber that relies on a process called stimulated Raman scattering to amplify a weak signal, as described in Section 9.6.

A third is a semiconductor optical amplifier, which will be covered in Section 10.6.4.

6.6.6 Amplified Spontaneous Emission: An Intermediate Stage

The division between oscillation and amplification is not as clear as it might seem. Ideally, you may think spontaneous emission of a

single photon could stimulate the entire cascade of stimulated emission in a laser beam, so all photons are coherent with each other. Something close to that can happen if the upper laser level is long-lived, so spontaneous emission is unlikely and stimulated emission can build to a high level before another photon is emitted spontaneously in a direction that would resonate in the laser cavity.

However, high-gain materials spontaneously emit many photons, which can also stimulate emission by other atoms in the upper laser level. Spontaneously emitted photons are at roughly the same wavelength, but are not in phase with each other, so amplification of spontaneous emission reduces the coherence of the beam. If spontaneous emission dominates, the result is *amplified spontaneous emission*, which is sometimes called *superradiance*. It can produce plenty of light, but the light is not very coherent. This can occur in very high-gain laser materials, particularly if the cavity is not highly selective.

Amplified spontaneous emission is desirable for applications in imaging systems where the speckle from coherent light causes undesirable noise. It is also attractive for optical coherence tomography because it has high spatial coherence, allowing tight focusing. However, amplified spontaneous emission can also create undesirable background noise in fiber-optic communication systems that carry signals at multiple wavelengths. Figure 6-11 shows the effect in an erbium-doped fiber amplifier. The four narrow lines represent signals at four wavelengths rising above low-level amplified spontaneous emission. Later amplifiers operated at high gain will amplify the spontaneous emission as well as the signal, producing background noise across the whole gain bandwidth of the amplifier, and a series of such amplifiers will build up noise that obscures the signal.




Figure 6-11. Amplified spontaneous emission in an erbium-doped fiber amplifier creates noise between the wavelengths carrying signals, shown as narrow vertical lines.

6.7 WAVELENGTH OPTIONS

We normally think of lasers as emitting light at only a narrow range of wavelengths. However, lasers can be designed to offer a variety of wavelength options.

6.7.1 Fixed-Wavelength Monochromatic Lasers

Laser oscillation in a resonant cavity normally concentrates emission in a narrow range of wavelengths where the laser transition has peak gain. The lower the gain and the narrower the cavity resonance and the gain bandwidth, the narrower is the range of wavelengths. Broader bandwidth is possible by using high-gain media in cavities with broad gain bandwidth. The choice of the transition and the design of the cavity select the wavelength.

As mentioned in Section 4.2.3, if gain bandwidth is broad enough and the laser cavity is relatively long, the laser cavity may allow oscillation on two or more longitudinal modes as shown in Figure 4-4. Many laser applications are insensitive to multimode operation. If applications require single-longitudinal mode operation, optics can be added to restrict the cavity to oscillating in a narrow range of wavelengths that prevent oscillation outside of a single mode.

6.7.2 Multitransition Lasers

Many laser materials have multiple laser transitions which can oscillate simultaneously in laser cavities with mirrors that reflect across a wide enough bandwidth. The wavelengths may be spaced broadly, as in argon-ion lasers, or closely, as in carbon-dioxide lasers. The ease of producing multiwavelength oscillation differs widely among laser types and is generally easier for closely spaced lines, such as CO₂ lasers. Multitransition lasers are acceptable for many applications that require high power, but few applications require them.

6.7.3 Wavelength-Tunable Lasers

Lasers with a broad gain bandwidth can be made tunable by operating them in an optical cavity that can be adjusted to select an oscillation wavelength within the material's gain band.




Figure 6-12. Prism in a tunable laser cavity disperses wavelengths; moving the rear cavity mirror selects one wavelength to oscillate.

One approach is to insert a prism or diffraction grating that disperses light to form a spectrum and then align the cavity mirrors or the dispersive element to select the desired wavelength. Figure 6-12 shows the idea using a prism to disperse light generated by optically pumping dye molecules dissolved in solution. Stimulated emission from the dye passes through the prism, which bends different colors in different directions. The rear cavity mirror is moved so it reflects one wavelength back through the prism and dye cell to the output mirror. In the figure, the cavity mirror is lined up to reflect green light, in the middle of the spectrum, so it resonates in the laser cavity.

Light waves at the red and violet ends of the spectrum are bent to the sides beyond the mirror so they are not reflected back through the dye cell to the output mirror. The dashed lines show where the rear cavity mirror would have to be to select those wavelengths. Only the green light is passed back and forth through the laser cavity, so stimulated emission amplifies the green light to high power, overwhelming the red and violet light, which do not oscillate in the laser cavity. (No real dye emits over such a broad range of wavelengths; a typical dye bandwidth is 10 to 70 nm, only a fraction of the visible spectrum. Actual tunable laser cavities are more complex.)

Laser wavelength can be tuned in other ways. The prism can be turned to aim the desired wavelength toward a fixed mirror. Alternatively, a reflective diffraction grating could replace both the rear cavity mirror and prism, and the grating could be turned to select

the wavelength reflected back through the dye cell and the output mirror.

Other optical elements can also tune wavelength or limit oscillation to a very narrow range of wavelengths for precision measurements. For example, a device called a *Fabry-Perot etalon* can be inserted into a tunable laser cavity to limit oscillation to a narrow range of wavelengths. The etalon is a pair of closely spaced, partly reflecting surfaces that bounce light back and forth between them. Interference effects allow the etalon to transmit only wavelengths that resonate between the two reflective surfaces. If the etalon is put inside the laser cavity, tilting affects how far the beam travels through it, changing the resonant wavelength transmitted. Combining an etalon with a movable mirror cavity can select a very narrow range of wavelengths.

High-performance tunable lasers have far more elaborate designs, such arrangements of several mirrors that form the cavity into a ring. Using such advanced techniques, the spectral bandwidth of a continuous-wave tunable laser can be limited to well below one part per billion.

6.7.4 Broadband Lasers to Emit Short Pulses

A major use of broadband lasers is to generate light pulses of very short duration. As you learned in Section 4.6.2, a broad range of wavelengths is needed to generate extremely short pulses because the product of bandwidth times pulse duration is a constant, as shown in equations 4-14 and 4-15.

Thus, concentrating laser energy into pulses lasting 100 femtoseconds (100×10^{-12} second, or a tenth of a trillionth of a second) requires laser light spanning a bandwidth of 4410 gigahertz. That corresponds to about 15 nm of laser bandwidth at a center wavelength of 1000 nm. Producing shorter pulses may require using devices that spread the range of wavelengths in the laser beam, typically done using nonlinear devices like those described in Section 5.6.

6.8 LASER-LIKE LIGHT SOURCES

A number of optical techniques using lasers or laser-related concepts can produce light with coherence, intensity, and/or

monochromaticity that may resemble laser light, but is not produced by stimulating emission from a population inversion. Section 6.5.3 described some examples, and Chapter 11 gives more details, but this section gives a quick overview.

6.8.1 Harmonic Generation

As you learned in Section 6.5.3, nonlinear processes can generate harmonics of laser frequencies. The simplest is *second-harmonic generation* or frequency doubling, as you saw schematically in Figure 5-12. Passing laser light through a nonlinear material generates coherent light at twice the input frequency or, equivalently, half the wavelength. This is a handy way to generate shorter wavelengths. Additional steps can generate the third, fourth, and fifth harmonics. Firing extremely intense laser pulses into gases can generate even higher harmonics, as described in Section 11.5.2.

6.8.2 Raman Shifters and Raman Lasers

Another nonlinear effect that can produce laser-like light is *stimulated Raman scattering* described in Section 5.6.4. Devices called *Raman shifters* can change the wavelength of a laser source and may be packaged separately or with a laser. Raman scattering in optical fibers can also be used in fiber amplifiers and lasers, described in Section 9.6.

6.8.3 Optical Parametric Sources

Optical parametric light sources use the sum- and difference-frequency generation described in Section 5.6.3 to convert light from a laser source at one wavelength to tunable light at other wavelengths. They are used both in single-pass amplifiers and in resonant oscillator configurations, which are described in Section 11.2.

6.8.4 Supercontinuum and Frequency Combs

Illuminating a strongly nonlinear material with very intense light source can spread the light energy across a broad range of wavelengths, called a *supercontinuum*. Generally, the light is first concentrated into very short pulses to produce extremely high peak power, which is focused to achieve the high intensity needed to produce a supercontinuum. A supercontinuum is a repetitively

pulsed light source that can span the whole visible spectrum or the even wider range of an octave of wavelengths, as described in Section 11.3, and a building block for the frequency combs described in Section 11.4.

6.8.5 Superluminescent LEDs

Amplified spontaneous emission sometimes can dominate output at the threshold of laser oscillation. Often, this is simply a transitional stage when a laser is starting up. But it is possible to operate semiconductor diodes as *superluminescent LEDs (SLEDs)*, where amplified spontaneous emission dominates because the diode is designed to minimize feedback that would produce stimulated emission and cause laser oscillation.

The attraction of a superluminescent LED is its combination of the high brightness of a laser with the low coherence and broader bandwidth of an LED, 5 to 100 nm. The low coherence avoids noise that can arise from coherent effects, particularly *speckle*, which produces bright spots in an image illuminated by a laser source in air. Look closely at a spot projected on wall by a laser pointer and you can see speckle. (Speckle is easiest to see in the area around the laser spot, particularly if it is magnified by a lens.) Speckle arises from air turbulence, such as the twinkling of stars, and degrades laser-projected images. Thus, “SLEDs” are used in fiber-optic gyroscopes and low-coherence optical coherence tomography for medical imaging.

6.9 WHAT HAVE WE LEARNED?

- Many types of materials can produce laser light. They are classed by their state and composition.
- Laser action in neutral or ionized atomic gases is a mature technology powered by electric discharges in the gas.
- Molecular gases typically emit on infrared transitions and are powered by steady electric discharges.
- Short-lived molecules called excimers produced by pulsed electrical discharges generate laser pulses in the ultraviolet.
- Solid-state lasers are based on nonconductive solids such as glasses and crystals, which typically contain 1% or less of the laser species. They are optically pumped with a light source matched to absorption band of the laser species.

- Semiconductor diode lasers are not classed as solid-state lasers because they are powered by electricity passing through the semiconductor.
- Fiber lasers are solid-state lasers in which the solid is an optical fiber with light-emitting atoms in its light-guiding core.
- Most semiconductor lasers are electrical diodes, in which recombination of electrons at holes in a thin junction layer produces a population inversion.
- Lasers may be powered by light or electric current passing through the laser medium.
- Optical pumping uses lamps or other lasers.
- Diode pumping uses semiconductor diode lasers.
- Optical pumping essentially converts pump light into a tightly focused laser beam at longer wavelength.
- Electrical pumping works only in materials that conduct electricity such as semiconductors and gases.
- Energy storage within the laser material and the laser cavity is an important factor in producing laser pulses.
- Diode lasers can be modulated by turning the drive current off and on.
- Quasi-continuous-wave lasers produce a series of pulses with peak powers higher than the steady output of continuous-wave lasers. They typically emit light 0.1% to 10% of the time.
- Q switching builds up energy stored in excited atoms in the laser medium and then releases the energy by allowing the light to escape from the laser cavity in a powerful burst. It works best in optically pumped solid-state lasers.
- Cavity dumping redirects light circulating in a laser cavity out of the cavity to produce a brief laser pulse.
- Modelockers are optical switches that allow light emission only for a brief fraction of the time light takes to make a round trip of the laser cavity. This produces a clump of photons that bounce back and forth within the laser cavity, with the output mirror transmitting part of the photons each time the clump reaches it.
- Both active and passive devices can produce Q switching and modelocking.
- Laser wavelengths can be changed by nonlinear effects including harmonic generation, sum- and difference-frequency generation, and Raman shifting.
- A laser oscillator has feedback from the resonant cavity and generates a beam on its own at a wavelength depending on the laser

medium and the cavity. It is the optical equivalent of a radio-frequency oscillator.

- An optical amplifier lacks a resonant cavity and feedback; it amplifies light from an external source.
- A MOPA is a multistage pulsed laser, with a laser oscillator followed by one or more optical amplifier stages.
- The MOPA design can produce higher energy pulses than a single-stage oscillator.
- Chirped-pulse amplification starts by stretching a short-duration input pulse into a longer pulse with lower peak power, then amplifying it, before compressing the amplified pulse to short duration. In this way, it can generate extremely high peak powers without damaging the amplifier.
- Optical signal amplifiers boost the strength of weak signals in fiber-optic communication systems. They can simultaneously amplify signals on multiple wavelengths within their gain bands.
- Superluminescent LEDs produce bright amplified spontaneous emission without lasing, avoiding speckle and other coherence effects.
- Laser oscillation normally concentrates emission in a narrow range of wavelengths, but some lasers can oscillate on multiple transitions with broadband optics.
- Tunable laser emission requires a laser medium with a wide line width and a cavity that can be adjusted to resonate across a range of wavelengths.
- Wide laser bandwidth is needed to generate very short laser pulses.

WHAT'S NEXT?

In Chapter 7, you will learn about gas lasers, the most diverse family of lasers.

QUIZ FOR CHAPTER 6

1. Which of the following is not considered a solid-state laser material?
 - a. A crystal of synthetic ruby doped with chromium.
 - b. A semiconductor diode containing gallium arsenide.

- c. A glass optical fiber with its light-guiding core doped with erbium.
 - d. A slab of glass containing neodymium.
 - e. A crystal of yttrium vanadate containing neodymium.
2. What is a dielectric material?
 - a. An opaque solid that conducts electricity.
 - b. A transparent semiconductor that conducts electricity poorly.
 - c. A transparent crystal that does not conduct electricity.
 - d. A metallic solid that conducts electricity readily.
 - e. A glass that is electrically conductive.
 3. What is not an advantage of semiconductor diode lasers for optical pumping?
 - a. Diode lasers efficiently convert electrical power into laser power.
 - b. Diode lasers are available with output wavelengths that match strong absorption lines of solid-state laser materials.
 - c. Diode laser output can easily be coupled into the light-guiding core of a fiber laser.
 - d. Diode lasers can produce much higher pulse energies than lamps.
 - e. None of the above.
 4. What is the most essential requirement for electrical pumping of a laser material?
 - a. It must be transparent.
 - b. It must be a gas.
 - c. It must be a semiconductor.
 - d. It must be opaque.
 - e. It must conduct electricity.
 5. What is required for a laser to operate as an oscillator?
 - a. Stimulated emission must occur within a resonant cavity.
 - b. Stimulated emission must occur without a resonant cavity.
 - c. The gain within the cavity must equal the loss within the cavity.
 - d. The gain within the cavity must equal or exceed the loss plus the power leaving through the output mirror.
 - e. Spontaneous emission must occur.
 6. What is needed for an optical amplifier?
 - a. Stimulated emission must occur within a resonant cavity.
 - b. Stimulated emission must occur without a resonant cavity.
 - c. The gain within a resonant cavity must equal the loss within the cavity.

- d. The gain within the cavity must equal or exceed the loss of light in the cavity and through the output mirror.
 - e. Spontaneous emission must occur.
7. How many signals can an optical amplifier amplify at separate wavelengths?
- a. One only
 - b. Two
 - c. One per amplifier stage
 - d. Four
 - e. As many as can fit within the amplifier's gain bandwidth
8. How does Q switching work?
- a. It builds up a large population of atoms in the upper laser level by preventing laser oscillation and then switches to allow laser oscillation, which extracts the energy by stimulated emission to produce an energetic pulse.
 - b. It pumps the interior of the laser with a burst of electrical energy and then turns the output mirror fully transparent.
 - c. It allows laser oscillation between two totally reflective mirrors, which builds up a large amount of light inside the laser cavity. Then it shifts from transparent to reflective to reflect all that circulating light out of the laser cavity.
 - d. It produces intense pulses of electricity which trigger powerful emission of light from every atoms inside the laser cavity.
 - e. It creates a strong electric field that increases the rate of spontaneous emission by the laser material by more than a factor of 1000.
9. Which can produce light pulses of the shortest duration?
- a. Quasi-continuous-wave operation.
 - b. Q switching.
 - c. Cavity dumping.
 - d. Modelocking.
 - e. Direct electrical modulation.
10. You are modelocking a laser that emits at 800 nm with a line width of 40 nm. What is the shortest pulse you can produce? (Hint: Check equations 4-14 and 4-15.)
- a. 2.4 femtoseconds
 - b. 24 femtoseconds
 - c. 39 femtoseconds
 - d. 240 femtoseconds
 - e. 390 femtoseconds

11. You have a narrow-line laser with bandwidth of only one megahertz that emits at 800 nm. What is the shortest pulse you could produce with that laser?
 - a. 441 femtoseconds
 - b. 441 picoseconds
 - c. 4.41 nanoseconds
 - d. 44.1 nanoseconds
 - e. 441 nanoseconds
12. What wavelength do you get when you frequency-double the 1064-nm output of a neodymium–YAG laser?
 - a. 532 nm
 - b. 800 nm
 - c. 1062 nm
 - d. 1596 nm
 - e. 2128 nm

GAS LASERS

ABOUT THIS CHAPTER

Most gas lasers share some key features. Gas is put into a tube, mirrors are put on the ends, and an electric discharge is passed through the gas to produce a population inversion. The details vary quite a bit. The discharge may pass across the tube or along its length. Some molecular gases emit in the mid to far infrared on vibrational or rotational transitions. Atomic gases tend to emit in or near the visible. Ionized gases may emit in the visible or ultraviolet. One odd family of short-lived molecules emits bright ultraviolet pulses. The result is a varied group of lasers that share important features and can illustrate some important aspects of laser physics. Two families of gas lasers, carbon dioxide (CO₂) and excimer lasers, are important commercially, but other gas lasers have become niche products. This chapter will explore the basics of gas lasers and teach you about the most important types.

7.1 THE GAS-LASER FAMILY

The family of gas lasers grew rapidly after Ali Javan, William R. Bennett Jr., and Donald R. Herriott demonstrated the first gas laser at Bell Laboratories in December 1960. Their helium–neon laser was the first laser to generate a continuous beam, and the red version first demonstrated in 1962 remains in use. Since then, laser action has been demonstrated at literally thousands of wavelengths in a

wide variety of gases, including other rare gases, metal vapors, and many different molecules.

Gas lasers vary widely in their characteristics. The weakest commercial lasers emit under a thousandth of a watt, but the most powerful you can buy emit thousands of watts in a continuous beam. In the past, experimental high-energy gas lasers have produced megawatt-class beams for brief intervals in military experiments. Some gas lasers can emit continuous beams for years, and others emit pulses lasting a few billionths of a second. Their outputs range from deep in the vacuum ultraviolet—at wavelengths so short that they are blocked completely by air—through the visible and infrared to the borderland of millimeter waves and microwaves.

What makes the gas-laser family so diverse? Gases are easy to study; therefore, extensive data were collected on their spectra and energy levels. Experiments are straightforward, and in the early days of gas-laser research, some modest university laboratories ran up impressive lists of discoveries by pumping gas into a glass tube, testing it, then replacing with other gases and testing them. For many years, gas lasers were standard for a broad range of laser applications, from simple classroom demonstrations to sophisticated instrumentation.

Sales of most types of gas lasers have faded as semiconductor and solid-state lasers have replaced them for many applications. However, two types of powerful gas lasers have remained important products, with annual sales in the hundreds of millions of dollars. Section 7.6 covers the CO₂ laser emitting at 10 μm in the infrared, a wavelength not readily available from solid-state or semiconductor lasers, which is important in medicine and industrial applications. Section 7.7 covers *excimer* or *rare-gas-halide lasers* emitting in the ultraviolet, which are also used in medicine and industry. For convenience, we will start with a generic gas laser and then cover those specialized types later.

7.2 GAS-LASER BASICS

Most gas lasers share many features with the generic gas laser in Figure 7-1. The laser gas is contained in a tube with cavity mirrors at each end, one totally reflecting and one transmitting some light to form the output beam. Most gas lasers are excited by passing an electric current through the gas. At low to moderate




Figure 7-1. Generic gas laser.

powers, the discharge is usually *longitudinal*, along the length of the tube, as shown in Figure 7-1. However, higher-power lasers and some types that operate at high gas pressure are often excited *transversely* across the width of the tube. Electrons in the discharge transfer energy to the laser gas, generally with another step needed to excite atoms or molecules to the upper laser level and generate a population inversion. Stimulated emission then resonates within the laser cavity and produces the laser beam.

7.2.1 Gas-Laser Media

Gas-laser media can be divided into three main types based on the light-emitting species.

Atomic gases include both neutral and ionized atoms, which emit light on transitions between electronic energy levels. Typically, atomic gas or ion lasers emit in the near ultraviolet, visible, or near infrared. Most of these lasers can emit a continuous beam. Examples are the helium–neon and argon-ion lasers. Metal vapors are the light-emitting species in a few atomic gas lasers, notably the helium–cadmium and copper-vapor lasers. These metals are solids at room temperature; therefore, the laser tube must be heated to vaporize some of the metal contained inside it before the laser can emit light.

Molecular gases are molecules that usually emit on transitions between vibrational and/or rotational energy levels, although

molecular nitrogen (N_2) emits on electronic transitions. Vibrational and rotational transitions can generate pulses or continuous output at near-, mid-, and far-infrared wavelengths. The CO_2 laser is the most powerful commercial gas laser.

Excimers are diatomic molecules, which are stable in an electronically excited state, but dissociate in the lower laser level. They emit pulses of ultraviolet light on electronic transitions as they fall apart and drop to a lower energy level. Most consist of one atom of a rare gas and one atom of a halogen. Examples are the krypton fluoride and xenon chloride lasers.

Only a single species emits light in most gas lasers; the only significant exception is an ion laser emitting on lines of both argon and krypton. However, many gas lasers contain additional species that serve other purposes, such as transferring energy to the laser species, helping to depopulate the lower laser level, or helping to remove waste heat. In helium–neon lasers, helium atoms capture energy from electrons passing through the gas and then transfer the energy to neon atoms, exciting them to the upper laser level. CO_2 gas lasers contain N_2 molecules that absorb energy from an electric discharge and transfer it to CO_2 molecules, as well as helium atoms that help CO_2 drop from the lower laser level (maintaining the population inversion), and assist in heat transfer. The optimum gas mixture for laser operation depends on the operating conditions and power level as well as on the light-emitting species.

Total gas pressure is an important variable because it affects energy transfer and how well the gas conducts electricity. The stable discharge required for continuous laser operation is easiest to sustain at pressures of a small fraction of one atmosphere. Pulsed lasers can operate at much higher pressures because they do not require sustaining a stable discharge for a long time.

7.2.2 Gas Replacement, Flow, and Cooling

Most low- to moderate-power gas lasers operate inside sealed tubes that isolate the pure gas from the atmosphere and allow operation at the optimum pressure for laser emission. Early gas lasers needed periodic gas replenishment or replacement because of gas leakage or contamination during operation. Great strides have been made in glass-sealing technology, and helium–neon lasers are now rated to operate for 20,000 hours or longer.

Some sealed gas lasers require periodic replacement of the laser gas because contaminants accumulate and gradually degrade laser action. This is a particular problem with excimer lasers, so their tubes are designed for periodic purging and refilling with fresh laser gas. Argon-ion lasers can be refurbished by adding new gas, cleaning the tube, and replacing some internal elements.

Gas flows through the tubes of many higher-power lasers, both to keep the gas at optimum operating temperature and to remove contaminants produced by electric discharges. Many flowing-gas lasers operate in a closed cycle, but others have open cycles that exhaust spent gas. Innocuous gases such as the CO_2 , N_2 , and helium in CO_2 lasers can be exhausted to the atmosphere, but some other laser gases are hazardous and must be collected in cartridges.

Gas lasers emitting a few milliwatts usually do not require active cooling, but higher-power gas lasers typically require active cooling with fans, closed-cycle refrigeration, external venting, or flowing water. Because CO_2 lasers are more efficient than most other gas lasers, they can be operated at higher powers without cooling.

7.2.3 Gas-Laser Excitation

The most common way to excite gas lasers is by passing a *longitudinal* electric discharge along the length of the laser tube, as you saw in Figure 7-1. As in a fluorescent tube, an initial high-voltage pulse ionizes the gas so it conducts electricity. Once the gas is ionized, the voltage is reduced to a lower level that can sustain the modest direct current needed to excite the laser gas. Operating at low gas pressure allows the stable discharge needed for a continuous laser.

Moderate-energy gas lasers can also be excited by a strong microwave field that excites atoms or molecules in the gas.

High-power and/or pulsed gas lasers are often excited *transversely*, by an electric discharge perpendicular to the length of the tube, as shown in Figure 7-2. This approach can pump more electrical energy into the laser medium faster than a transverse discharge and can be used at gas pressures too high for a stable longitudinal discharge.

As we will see in Section 7.9.3, a few gas lasers are pumped optically, usually by an electrically excited gas laser with shorter wavelength. Although their overall efficiency is limited, this approach can generate otherwise unobtainable wavelengths.

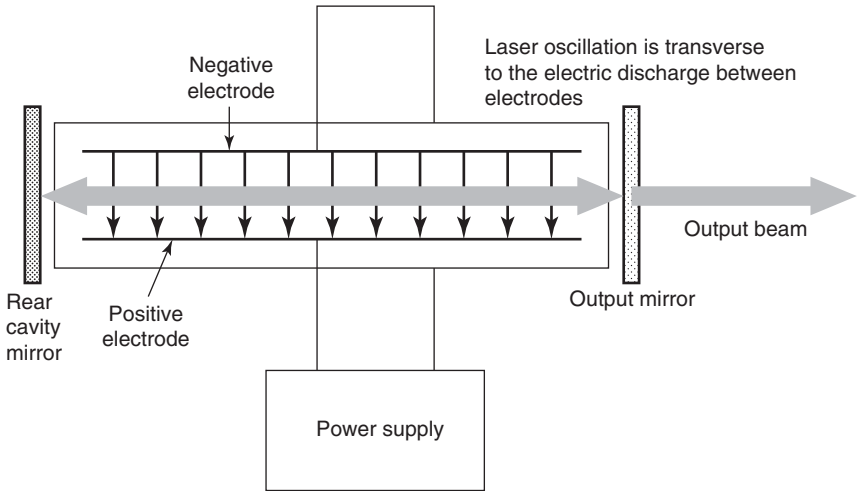


Figure 7-2. Transverse excitation of a gas laser.

7.2.4 Tube and Resonator Types

So far, our discussion of laser tubes and resonators has been vague on details such as placement of mirrors and the nature of the laser tubes. The reason is that things are a bit more complex than they might seem.

Gas-laser tubes need an output window on one end that lets light through but separates the laser gas from the air. They also need mirrors at both ends to form a reflective cavity. Applying a reflective coating to a window at the end of the laser tube can cut manufacturing costs, but it exposes the mirror coating to the discharge inside the laser tube, shortening its life. Using separate windows and mirrors improves performance, but adds to cost.

A common way to separate windows and mirrors is by mounting the windows at the ends of the tube at Brewster's angle, described in Section 4.7. The cavity mirrors are mounted outside the Brewster window, as shown in Figure 7-3. Brewster windows strongly attenuate one polarization, so strong amplification of the other polarization produces an output beam polarized in the direction that suffers less attenuation. Figure 7-3 also shows curved cavity mirrors defining a stable confocal resonator, which produces a good-quality, diffraction-limited beam with low divergence and is widely used in continuous-wave gas lasers at visible wavelengths.

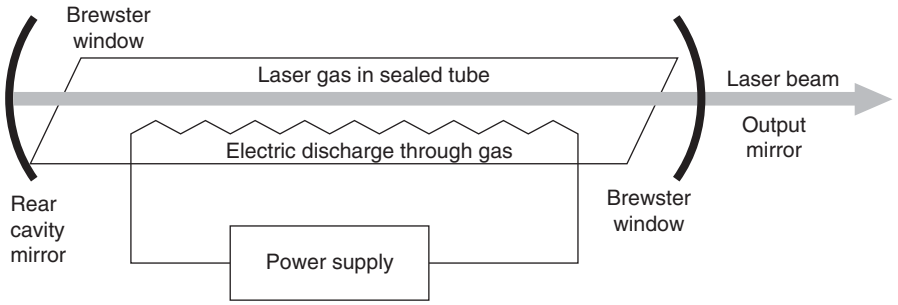


Figure 7-3. Laser tube with Brewster-angle window and confocal resonator.

Different types of cavities are used in high-gain, pulsed lasers, notably the rare-gas halide excimer lasers described in Section 7.7. It is relatively easy to make high-gain lasers, so they often are designed with only a single highly reflective mirror in the laser cavity with the output mirror having very low reflectivity. This design gives high-gain gas lasers comparatively large beam divergence and diameter.

Most tubes are glass or ceramic, and their prime role is to confine the laser gas and control excitation and gas flow. Some CO_2 lasers are operated in hollow waveguides that confine the laser light optically.

7.2.5 Wavelength and Bandwidth

Table 7-1 lists the nominal wavelengths and maximum powers of important gas-laser transitions. The electronic transitions have distinct nominal wavelengths; where ranges are given, they span a series of distinct lines, not a continuous range. Ranges are given for vibrational and rotational transitions because they include many closely spaced levels that can blend together to form a continuum at higher pressure.

The continual motion of atoms and molecules inevitably spreads gas-laser emission across a range of wavelengths. The velocity of the molecules along the length of the laser cavity causes a Doppler shift that increases or decreases wavelength, depending on the speed and direction of motion. The resulting *Doppler broadening* can spread laser output across a range of wavelengths proportional to the average speed of gas molecules, which depends on

Table 7-1. Wavelengths and power levels of major gas lasers, grouped by transition types involved. Lasers listed as being “megawatt class” have produced several hundred kilowatts or more in military weapon experiments, but that power is not used in any present system.

Type	Approximate wavelength (nm)	Approximate power range (W) [‡]	Normal operation
Electronic transitions			
Molecular fluorine (F ₂)	157	1–5 (avg.)	Pulsed
Argon fluoride excimer	193	0.5–50 (avg.)	Pulsed
Krypton fluoride excimer	249	1–100 (avg.)	Pulsed
Argon ion (UV lines)	275–305	0.001–1.6	Continuous
Xenon chloride excimer	308	1–100 (avg.)	Pulsed
Helium–cadmium (UV line)	325	0.002–0.1	Continuous
Nitrogen (N ₂)	337	0.001–0.01 (avg.)	Pulsed
Argon ion (UV lines)	333–364	0.001–7	Continuous
Krypton ion (UV lines)	335–360	0.001–2	Continuous
Xenon fluoride excimer	351	0.5–30 (avg.)	Pulsed
Helium–cadmium (UV line)	354	0.001–0.02	Continuous
Krypton ion	406–416	0.001–3	Continuous
Helium–cadmium	442	0.001–0.10	Continuous
Argon ion	488–514.5	0.002–25	Continuous
Copper vapor	510 and 578 nm	1–50 (avg.)	Pulsed
Helium–neon	543	0.0001–0.002	Continuous
Helium–neon	594	0.0001–0.002	Continuous
Helium–neon	612	0.0001–0.002	Continuous
Gold vapor	628	1–10	Pulsed
Helium–neon	632.8	0.0001–0.05	Continuous
Krypton ion	647*	0.001–7	Continuous
Helium–neon	1153	0.001–0.015	Continuous
Iodine and oxygen–iodine	1315	megawatt-class	Pulsed or CW
Helium–neon	1523	0.001	Continuous
Vibrational transitions			
Hydrogen fluoride (chemical)	2600–3000 [†]	0.01–150, megawatt class	Pulsed or CW
Deuterium fluoride (chemical)	3600–4000 [†]	0.01–100, megawatt class	Pulsed or CW
Carbon monoxide	5000–6500 [†]	0.1–40	Pulsed or CW
Carbon dioxide	9000–11000 [†]	0.1–20,000	Pulsed or CW
Vibrational or rotational transitions			
Far infrared	30,000–1,000,000 [†]	<0.001–0.1	Pulsed or CW

*Other wavelengths also available.

[†]Many lines in this wavelength range.

[‡]For typical commercial lasers.

atomic mass M and the gas temperature T . The average velocity in this case is measured as the *root mean square* (the square root of the sum of the squares) $\langle v \rangle$:

$$\langle v \rangle = \left(\frac{3kT}{M} \right)^{1/2} \quad (7-1)$$

where k is the Boltzmann constant. This velocity at normal operating temperatures is large enough to make Doppler broadening the main factor in determining the spectral width of most gas lasers. For example, the Doppler width (defined as full-width at half maximum) of a typical helium–neon gas laser is about 1.4 gigahertz or about 0.0019 nm. Although this is small compared with the 4.738×10^{14} hertz frequency of the laser’s 632.8-nm transition, it is much larger than the 1-megahertz bandwidth of one longitudinal mode of a typical helium–neon laser cavity, and large enough to include several longitudinal modes, separated by about 500 MHz.

Now that we have covered these basics, we will move on to important types of gas lasers.

7.3 HELIUM-NEON LASERS

Until the 1990s, the laser people were most likely to see was the low-power red helium–neon laser, which was used for supermarket checkout, construction alignment, and educational demonstrations. The least expensive of all gas lasers, the helium–neon laser, is often called a “He–Ne.” Its coherence and its visible beam long made the He–Ne the standard laser for educational demonstrations and recording holograms. Its coherence and spectral stability are important for some scientific and medical instruments, and it is still used for holography, but smaller and cheaper red semiconductor lasers have replaced the He–Ne in most other applications.

Sales of helium–neon lasers dropped from about 430,000 in 1991 to 39,000 in 2006, according to *Laser Focus World*. Some two-thirds of the He–Nes sold in 2006 were for use in scientific and biomedical instruments. More recent figures are not available. Sales appear to have continued declining, but some specialized markets remain and He–Nes remain in production.

7.3.1 Physical Principles

Helium–neon lasers can emit on several transitions among the energy levels shown in Figure 7-4. Electrons passing through a mixture of five parts helium and one part neon excite both species to high energy states, with the more abundant helium atoms collecting most of the energy. Excited helium atoms readily transfer that energy to neon atoms when the two collide, raising the neon atoms to the 5s and 4s energy levels. (The energy levels are spectroscopic labels convenient to use as identification, but you do not need to worry about the details until you go further into lasers.) The 5s and 4s energy levels of neon are metastable, so atoms stay in those states for a comparatively long time, producing a population inversion.

The transitions shown in Figure 7-4 are the strongest and most useful of the several possible transitions in helium–neon mixtures. Note that some transitions descend from the 5s level and others from the 4s level. The first helium–neon laser operated at 1153 nm




Figure 7-4. Energy levels and laser transitions in the helium–neon laser (not to scale).

in the infrared, but the demand for visible beams made the 632.8-nm red line the standard for helium–neon. He–Ne lasers emitting on the 543-nm green line are also available.

Once neon atoms drop to the lower laser level, they quickly drop through a series of lower energy levels to the ground state. Energy transfer from helium atoms can then raise them again to the upper laser level. The overall efficiency of the helium–neon is low—typically 0.01% to 0.1%—because the laser transitions are far above the ground state.

The overall gain of the helium–neon laser is very low, so care is needed to minimize laser cavity losses, as we will see later. However, the helium–neon laser is simple, practical, and inexpensive for gas lasers; mass-produced sealed-tube versions can operate continuously for tens of thousands of hours. Ionizing the gas takes about 10,000 volts, but a couple of thousand volts can maintain the current of a few milliamperes needed to sustain laser operation.

7.3.2 He–Ne Laser Construction

Figure 7-5 shows the internal structure of a typical mass-produced helium–neon laser. The discharge passing between electrodes at opposite ends of the tube is concentrated in a narrow bore, one to a few millimeters in diameter. This raises laser excitation efficiency and also helps maintain good beam quality. The bulk of the tube




Figure 7-5. A mass-produced helium–neon laser (courtesy of Melles Griot).

volume is a gas reservoir containing extra helium and neon. Gas pressure within the tube is typically a few tenths of a percent of atmospheric pressure.

Mirrors are bonded directly to mass-produced helium–neon tubes by a high-temperature process that produces what is called a “hard seal,” which slows the helium leakage that otherwise might limit laser lifetime. The mirrors must have low loss because of the laser’s low gain. The rear cavity mirror is totally reflective. The output mirror transmits only a few percent of the intracavity power to produce the laser beam. One or both mirrors have concave curvature to focus the beam within the laser cavity, which is important for good beam quality. An option is to seal the rear of the laser cavity with a Brewster window and mount the rear mirror separate from the laser tube to give polarized output.

The output power of helium–neon lasers depends on tube length, gas pressure, and diameter of the discharge bore. Researchers have found that output power is highest when the product of gas pressure (in torr; 760 torr equals one atmosphere) times bore diameter (in millimeters) is 3.6 to 4. Extending the tube can raise output power somewhat, but the improvements are limited, and 50 mW is about the maximum practical output.

The helium–neon laser can also be operated in a ring cavity to sense rotation. The ring laser is actually a square or triangular array of tubes, with mirrors at each corner reflecting the beam into the next arm. A ring laser “gyroscope” can detect rotation about an axis perpendicular to the ring plane by measuring phase differences in light traveling in different directions around the ring.

7.3.3 Practical Helium–Neon Lasers

Mass-produced helium–neon lasers can deliver 0.5 to about 10 milliwatts of red light. They range in size from little bigger than fat pens, 10 centimeters long and 1.6 cm in diameter (4 by $\frac{5}{8}$ inch), to 30 cm (a foot) or larger and 2 to 5 cm (1 to 2 inches) in diameter. A few special-purpose helium–neon lasers are considerably larger and can produce up to 50 mW in the red.

Although other helium–neon wavelengths listed in Table 7-1 have been available commercially since the mid-1980s, most helium–neon lasers emit the 632.8-nm red line. Output is much weaker, typically no more than a couple of milliwatts, at other visible wavelengths: 543 nm in the green, 594 nm in the

yellow-orange, and 612 nm in the red-orange. Milliwatt power is available on infrared lines at 1.153, 1.523, and 3.39 μm .

Complete new red helium–neon lasers ready for laboratory or other use start at a few hundred dollars; other wavelengths may be more expensive.

Helium–neon lasers typically emit TEM_{00} beams, with diameter about a millimeter and divergence about a milliradian. The typical Doppler broadened bandwidth of a red helium–neon laser is 1.4 GHz, which corresponds to a coherence length of 20 to 30 cm, adequate for holography of small objects. To get a longer coherence length, you must buy a laser limited to a single longitudinal mode with 1-MHz bandwidth, with 200- to 300-m coherence length.

7.4 ARGON- AND KRYPTON-ION LASERS

Like helium–neon lasers, argon- and krypton-ion lasers are powered by electric discharges passing through elements of the rare-gas (group VIII) column of the periodic table. All three emit continuous beams. However, there are also crucial differences. Argon and krypton lasers have their visible laser transitions in ions rather than the neutral atoms in the helium–neon laser. For that reason, they are often called simply *ion lasers*, although ions are the light-emitting species in some other gas lasers, notably helium–cadmium. Argon and krypton generate more powerful beams and emit at shorter wavelengths than helium–neon, with argon more powerful and more important commercially.

Historically, argon- and krypton-ion lasers were long the highest-power visible lasers emitting continuous beams into the 1990s. Improvements in solid-state, fiber and semiconductor lasers changed that, and sales of argon and krypton lasers have declined because of their high cost, complexity, and low efficiency. Solid-state and semiconductor lasers now are available at or close to the main ion-laser lines, but argon and krypton lasers still are used in some scientific and biomedical instruments.

7.4.1 Properties of Argon and Krypton Lasers

The active medium in rare-gas ion lasers is argon or krypton at a pressure of roughly 0.001 atmosphere. Tubes containing mixtures of the two gases can emit on lines of both elements across the visible

Table 7-2. Major wavelengths of ion lasers, with strongest lines noted

Argon (nm)	Krypton (nm)
275.4	337.4
300.3	350.7
302.4	356.4
305.5	406.7
334.0	413.1
351.1	415.4
363.8	468.0
454.6	476.2
457.9	482.5
465.8	520.8
472.7	530.9
476.5	568.2
488.0 (strong)	647.1 (strong)
496.5	676.4
501.7	752.5
514.5 (strong)	799.3
528.7	
1090.0	

spectrum, a capability used mostly in laser light shows. Table 7-2 lists important lines of the two elements in the near-ultraviolet, visible, and near-infrared parts of the spectrum.

The laser lines of argon and krypton are in ions with one or two electrons stripped from their outer shells. Ultraviolet wavelengths shorter than 400 nm come from atoms with two electrons removed (Ar^{+2} or Kr^{+2}). Visible output comes from singly ionized atoms (Ar^+ or Kr^+). Argon is the more efficient laser gas and is used for most applications, but krypton lines are spread across more of the visible spectrum.

As in helium–neon lasers, an initial high-voltage pulse ionizes the gas, and then lower voltage produces a sustained discharge. The current ionizes the atoms and excites them to high-energy states and the upper laser level, as shown for singly ionized argon in Figure 7-6. Laser transitions occur between many pairs of upper and lower levels; there is no room to show them all in the figure. If the optics permit, argon and krypton lasers can oscillate simultaneously on several different visible wavelengths, each produced by a transition between a different pair of levels.

All the lower laser levels have very short lifetimes. Argon ions quickly drop from the lower laser level (an excited state of the ion)




Figure 7-6. Energy levels that produce visible argon lines.

to the ion ground state by emitting an extreme-ultraviolet photon at 74 nm, which does not leave the laser tube. The ground-state ion can recapture an electron to become a neutral atom or again be excited to the upper laser levels. Similar things happen in krypton.

The need to ionize atoms makes argon and krypton lasers energy-hungry. Visible-wavelength lasers draw discharge currents of 10 to 70 amperes, more than a thousand times the level in helium–neon lasers, although the low resistance of the ionized gas permits voltages of 90 to 400 volts. Ultraviolet lasers require even more energy because they must remove two electrons from the light-emitting atoms. Ionization heats the gas, so the Doppler bandwidth of a single ion-laser line is about 5 GHz.

Typical visible output powers range from a few milliwatts to 25 W for argon; ultraviolet output is lower. Most argon and krypton lasers sold today emit less than 1 W and are used in scientific and biomedical instruments.

7.4.2 Practical Argon and Krypton Lasers

Argon and krypton lasers have low gain, so cavity losses must be minimized. Materials must withstand the intense extreme-ultraviolet generated when argon atoms drop from the lower laser

level, so the tubes are generally ceramic. Cavity optics may allow oscillation on multiple lines, notably the 488 and 514.5 nm lines of argon, or oscillation can be restricted to a single wavelength. Generally, the optics produce TEM_{00} beams and include Brewster windows to select linear polarization.

As in helium–neon lasers, confining the discharge to a narrow region in the center of the tube enhances excitation efficiency. In newer ion lasers, a series of metal disks with central holes confine the discharge to the central region. Laser operation depletes the gas, so the tube must include a large gas reservoir. A return path is also needed for positive ions. With wall-plug efficiency only 0.01% to 0.001%, argon and krypton lasers need active cooling with forced air or, at powers above a few watts, flowing water.

Operating conditions inside argon and krypton laser tubes are extreme, so lifetimes are shorter than those of helium–neon lasers, typically ranging from 1000 hours for high-power tubes to 10,000 hours for low-power tubes.

Because of their inefficiency, high costs, and other drawbacks, argon- and krypton-ion lasers have largely been replaced by solid-state lasers operating in the visible.

7.4.3 Mixed-Gas Argon–Krypton Lasers

Argon and krypton can be mixed in a single laser tube to make a laser that oscillates simultaneously on many lines throughout the visible spectrum. The main use of such mixed-gas lasers is in light shows. Their lifetimes are relatively short because the two gases are depleted at different rates.

7.5 METAL-VAPOR LASERS

Vaporized metals are the light-emitting species in some gas lasers. The vapor may be ionized, as in the helium–cadmium laser, or neutral, as in copper-vapor lasers. The two families operate quite differently.

7.5.1 Helium–Cadmium Lasers

Helium–cadmium (He–Cd) lasers emit continuous beams at powers from under a milliwatt to tens of milliwatts, slightly more powerful than helium–neon, but less powerful than argon lasers.




Figure 7-7. Energy levels in the helium–cadmium laser.

The strongest He–Cd wavelength is 441.6 nm in the blue part of the spectrum, but weaker 325 and 353.6 nm lines in the ultraviolet are also used. Their main applications have been in industry, measurement, and instrumentation.

Figure 7-7 shows the energy levels involved in the He–Cd laser. An electric current flowing through a thin capillary bore excites helium atoms to high-energy states, and collisions transfer the energy to cadmium vapor, ionizing the atoms and raising them to the upper laser levels of the three transitions. One metastable state is the upper level of the two ultraviolet transitions; the other is the upper level of the 441.6-nm blue line. Cavity optics select the oscillating wavelengths. Cadmium normally has only two electrons in an incompletely filled outer shell, so it is much easier to ionize than helium, argon, or krypton, which all have full outer electron shells.

Cadmium is a solid at room temperature, and metal in the laser tube must be heated to about 250°C to produce the several millitorr

of cadmium vapor needed for the laser to operate. Helium pressure is about a thousand times higher, several torr. Extra cadmium is put in the tube to replace metal vapor that condenses on cool parts of the tube during laser operation. Tubes normally also include a helium reservoir to replace gas that leaks out. Discharge voltages are around 1500 V.

Gallium–nitride semiconductor lasers emitting at 445 nm are readily available that can replace the blue 441.6 nm He–Cd line, which has greatly reduced He–Cd usage.

7.5.2 Copper-Vapor Lasers

The copper-vapor laser emits at 511 nm in the green and 578 nm in the yellow. The internal physics limits it to pulsed operation, but it can generate thousands of pulses per second and average powers of tens of watts.

To obtain the 0.1-torr vapor pressure needed for laser action, metallic copper in the laser tube must be heated to 1500°C, which takes about half an hour after the laser is switched on. Helium, neon, or argon may be added to improve discharge quality and energy transfer. Once the laser is operating, its waste heat keeps the metal vaporized.

In operation, a pulsed discharge is fired along the laser tube, and collisions raise the copper atoms to one of two excited states which emit in the green and yellow. Copper atoms accumulate in the metastable lower states of those transitions, terminating the population inversion and stopping the laser pulse in tens of nanoseconds. Depopulation of the lower levels leaves the laser ready to fire another pulse, allowing repetition rates of several thousand pulses per second.

Copper-vapor lasers are inherently powerful, with a gain of 0.1 to 0.3 per centimeter. They do not require a high-gain cavity; commercial versions have a totally reflective rear cavity mirror, and an output mirror that reflects only about 10% of light back into the laser cavity. Average powers can reach tens of watts, with efficiency several tenths of a percent, attractively high for visible gas lasers.

Today, copper-vapor lasers have limited uses, primarily for high-speed imaging and research. Gold- and lead-vapor lasers were developed that operate similarly at different wavelengths, but they were never widely used.

7.6 CARBON DIOXIDE LASERS

The CO₂ laser is exceptionally versatile and highly efficient; up to 20% of the input power can be converted into laser light. It operates under a wide variety of conditions, emitting a steady beam at low gas pressure or pulses at high pressures. Powers of continuous beams can range from milliwatts to 20 kW; output can be on a single narrow line or spread across a series of lines in the 10- μ m band. Annual sales of CO₂ lasers for industrial materials-working are now hundreds of millions of dollars a year, but sales of fiber lasers now surpass those of CO₂.

The CO₂ laser emits on vibrational transitions of the CO₂ molecule in a band between 9 and 11 μ m in the infrared. The 10- μ m output is transmitted reasonably well by air and absorbed strongly by a wide range of materials, allowing it to cut, drill, and weld both metals and nonmetals. The CO₂ wavelength is also strongly absorbed by the water in tissue, so it has long been used for some types of surgery.

7.6.1 CO₂ Laser Transitions

The CO₂ laser is excited by an electric discharge passing through the laser gas, which contains nitrogen and (usually) helium as well as CO₂. The N₂ and CO₂ molecules absorb energy from electrons in the discharge. The lowest vibrational level of N₂ transfers energy easily to CO₂. Helium helps maintain the population inversion by getting CO₂ molecules to drop from the lower laser level to a lower level or the ground state.

The CO₂ molecule has three vibrational modes shown in Figure 7-8: the symmetric stretching mode ν_1 , the bending mode ν_2 , and the asymmetric stretching mode ν_3 . The molecule can vibrate at different rates in each mode, so each has its own set of energy levels (0, 1, 2, 3, etc.). The laser transitions occur when CO₂ molecules drop from the higher-energy asymmetric stretching mode excited by the discharge to the lower energy symmetric stretching or bending modes.

A CO₂ molecule in the asymmetric stretching mode would emit a 10.5- μ m photon if it dropped directly to the symmetric stretching mode and a 9.6- μ m photon if it dropped directly to the second excited level of the bending mode. However, the molecule changes

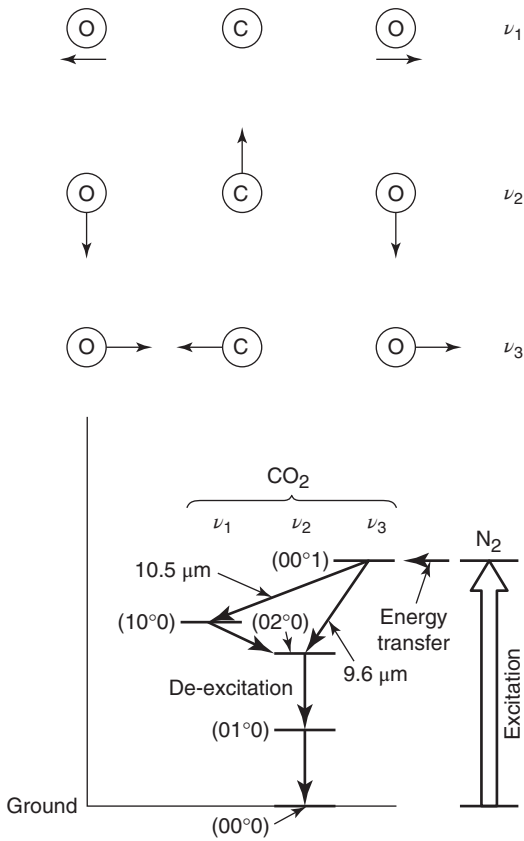


Figure 7-8. Vibrational modes and transitions of the CO₂ molecule.

its rotational state when it changes its vibrational state, so it does not emit those precise wavelengths. Speeding up its rotation takes some energy from the vibrational transition, so the emitted wavelength is longer than the nominal transition energy—for example, 10.6 μm rather than 10.5 μm . If the molecule slows down its rotation, the energy released adds to the energy from the vibrational transition, resulting in a shorter wavelength, such as 10.3 μm . This allows CO₂ lasers to emit on the range of closely spaced lines you saw earlier in Figure 4-5.

7.6.2 Sealed-Tube and Waveguide CO₂ Lasers

Low-power CO₂ lasers operate in simple sealed tubes like the generic gas laser in Figure 7-1. Continuous powers to a hundred

watts are possible by adding a flow-through gas reservoir to remove more heat. Typically, a high-voltage discharge passes between positive and negative electrodes at opposite ends of the tube, but the laser can also be excited by a radio-frequency-induced discharge. Water or a catalyst must be added to regenerate CO_2 molecules which the discharge splits into oxygen and carbon monoxide (CO).

The maximum power depends on length of the gain medium; a 1-m tube can generate about 50 W. Internal mirrors can fold the beam, but heat removal limits practical length.

An alternative approach to making sealed-tube CO_2 lasers is to shrink the tube's cross-section to a millimeter or two across, only about 100 times the 10- μm wavelength. A glass or ceramic tube of that size functions as an infrared waveguide, guiding the light waves through it (much like optical fibers guide visible light or hollow pipes guide microwaves). This structure avoids large diffraction losses that otherwise would occur with an output aperture that is so small compared with the wavelength. Gas flow is required in a waveguide laser, but it can be packaged as a sealed system with internal gas reservoir. Normally, waveguide CO_2 lasers emit continuous beams.

7.6.3 Longitudinal-Flow CO_2 Lasers

Higher-power CO_2 output is possible if fresh gas flows on the length of the laser cavity, in the same direction as the discharge current. This consumes gas, but pressures are low and the gases used in CO_2 lasers are not particularly hazardous or costly. Normally, some exhaust can be recycled by mixing it with fresh gas.

This straightforward design can produce output power per unit length somewhat higher than sealed-tube lasers, without the same limits on maximum total power. Continuous beams can deliver hundreds of watts. These lasers are more expensive to build and operate, so they are more costly than sealed CO_2 lasers.

7.6.4 Transverse-Flow CO_2 Lasers

Output power in a continuous CO_2 beam can reach about 10 kW per meter of tube length if the laser gas flow is *transverse* across the laser cavity rather than along its length. Transverse flow moves the gas through the laser cavity much faster than longitudinal flow along its length, removing waste heat and contaminants. The

electric discharge that drives the laser is also applied perpendicular to the tube axis, so it goes through a shorter length of gas, but pressure remains low to maintain a stable discharge.

Like longitudinal-flow lasers, the gas pressure is low and the output beam is continuous. Typically, the gas is recycled, with some fresh gas added. Transverse flow is normally used only in very high-power lasers, with total outputs in from kilowatt class to 20 kW.

7.6.5 Transversely Excited Atmospheric (TEA) CO₂ Lasers

Most sealed and flowing-gas CO₂ lasers operate at pressures below 0.1 atmosphere to stabilize the electric discharge needed for continuous operation, although the lasers can be pulsed. Operating at pressures near one atmosphere or higher requires pulsing electric discharges transverse to the laser cavity. Such *transversely excited atmospheric pressure* or *TEA* lasers are compact and can deliver powerful pulses lasting 40 ns to 1 μs, which are sought for some applications. The intense pulsed discharges break down CO₂, so gas replenishment or reconstitution is required.

7.6.6 Gas-Dynamic CO₂ Lasers

Rapid expansion of hot, high-pressure CO₂ mixed with other gases (often produced by combustion) through nozzles into a near-vacuum can produce a population inversion in CO₂. The rapid expansion cools the gas, but some molecules remain in high-energy levels after flowing through the nozzles. Arranging a pair of parallel mirrors on opposite sides of the gas flow can create a laser cavity that extracts stimulated emission from the gas to produce a laser beam. The result is called a gas-dynamic laser. Invented in the mid-1960s, it eventually reached hundreds of kilowatts in military experiments and inspired military efforts to develop megawatt-class flowing-gas lasers that burned other chemicals described in Section 7.9. However, those technologies never proved practical.

7.6.7 Optics for CO₂ Lasers

The 10-μm wavelength of the CO₂ laser sets it apart from other gas lasers described so far, which emit at much shorter wavelengths. Glasses and other materials used in visible and near-infrared optics

are opaque at 10 μm , so other materials must be used for transmissive optics. Reflective optics are often preferred where practical.

CO_2 lasers have good but not extremely high gain, so laser cavities normally have a totally reflective rear mirror and a partly reflective output mirror, which are usually separate from the windows at the end of the laser tube. Metal mirrors are often used in the laser cavity, with output coupling through a hole in the mirror rather than through a partly transmissive coating. Windows on the laser tube are transparent to the 10- μm laser beam, but generally not to visible light.

Research and measurement applications often require limiting CO_2 emission to a single laser line, and these lasers are made with tuning optics that select one specific wavelength. These wavelengths are discrete lines at the low pressures used in most CO_2 lasers, but broaden to become a continuum at the high pressures used in TEA lasers.

Most industrial and medical CO_2 laser applications are insensitive to the specific wavelength and instead require delivery of the highest possible power. These lasers are made with cavity optics reflective throughout the 9 to 11 μm range, to extract as much energy as possible from the CO_2 laser cavity in a good-quality beam.

Important transmissive materials used in the 10- μm band were shown in Figure 5-8. Most are opaque at visible wavelengths, but a few are transparent, including sodium chloride (which is transparent but must be protected from moisture) and zinc sulfide (which appears orange).

You cannot see 10- μm beams, and they can emerge from optics that look opaque to the eye, so special care is needed in using CO_2 lasers. Dust particles burn when exposed to very high-power CO_2 beams, an effect called “fireflies.” Infrared viewers that can reveal 10- μm beams are available, but take care in your selection. Many instruments called infrared viewers are sensitive to shorter bands near 1 μm or 3 to 5 μm , and do not show light in the 10- μm range. To see 10 μm , you need a “thermal” infrared viewer, so-called because the thermal (heat) radiation from room-temperature objects peaks near 10 μm . Safety regulations require that industrial CO_2 lasers be operated in enclosures that block the powerful and potentially dangerous 10- μm laser beam but transmit visible light so the operators can see inside, but research labs may not follow the same precautions.

Table 7-3. Major excimer lasers

Type	Wavelength (nm)
F ₂ *	157
ArF	193
KrCl	222
KrF	249
XeCl	308
XeF	350

*Not a rare-gas halide, but usually grouped with excimers.

7.7 EXCIMER LASERS

Excimer lasers are a family in which light is emitted by short-lived molecules made up of one rare-gas atom (e.g., argon, krypton, or xenon) and one halogen (e.g., fluorine, chlorine, or bromine). They are often called rare-gas halides. Table 7-3 lists the most important excimer lasers and the closely related molecular fluorine laser.

First demonstrated in the 1970s, excimer lasers have become important because of their unique ability to generate high-power ultraviolet pulses. Some are used in research, but many are used in two specialized but important markets. One familiar to the public is refractive eye surgery called LASIK, which reshapes the shape of the cornea to correct vision, so the patient will not need corrective lenses. The other is the multibillion dollar business of photolithographic fabrication of integrated semiconductor electronics. The short ultraviolet wavelengths of excimer lasers have been vital for packing huge numbers of transistors and other components onto chips, although further reductions will require new light sources.

7.7.1 Excimer Laser Physics

Excimers are peculiar diatomic molecules that are bound together only in an electronically excited state. The laser is excited by passing a short, intense electrical pulse through a blend of gases. Normally, 90% or more of the mixture is a buffer rare gas (typically helium or neon) not directly involved in the reaction. The rare gas that becomes part of the excimer (argon, krypton, or xenon) is a few percent of the mixture. The halogen atoms come from another gas, either halogen molecules such as F₂, Cl₂, or Br₂, or molecules that contain halogens such as nitrogen trifluoride (NF₃).

The discharge splits halogen-containing molecules, freeing excited halogen atoms to react with the rare gas to form electronically excited molecules such as xenon fluoride (written XeF^* , with the * meaning excited). The details are complex and depend on the gas mixture. The excited dimers remain in the electronically excited upper laser level for about 10 ns and then drop to the unbound lower laser level, releasing a photon as the molecule splits. This process rapidly empties the lower laser level, sustaining a population inversion. The duration of the electrical drive pulses and the molecular kinetics limits laser pulse length to several nanoseconds to about 100 ns.

Figure 7-9 shows the energy levels of a typical rare-gas halide as a function of spacing between the two atoms in the molecule, R (the rare gas) and H (the halide). The dip in the excited-state curve shows where the two atoms have a minimum energy, binding them together to form a metastable molecule. The horizontal lines are vibrational levels that exist in the potential well. No such minimum exists in the ground state curve, so the two atoms are unbound and drift apart.

Excimer laser repetition rates depend more on the power supply than on the gas. The principal limitation is speed of the high-voltage switches, which can operate at up to kilohertz rates in commercial lasers. Pulse energies of typical lasers range from




Figure 7-9. Internal energy of a rare-gas halide molecule in excited and ground states.

millijoules to joules; they drop as the repetition rate increases and vary somewhat among gases, with ArF, KrF, and XeCl generally the most energetic.

7.7.2 Excimer Laser Operation

Excimer lasers have such high gain that they almost do not need cavity mirrors. In practice, they have fully reflective rear mirrors and uncoated output windows that reflect a few percent of the beam back into the cavity and transmit the rest. The discharge is perpendicular to the length of the tube. Excimers are efficient for short-wavelength gas lasers, with wall-plug efficiency as high as 2%.

The corrosive halogens used in excimer lasers require special handling. The tubes must be made of halogen-resistant materials. After being filled with gas, the tubes are sealed and operated until the laser power declines to a point where the gas must be replaced. Tubes include a large volume of reserve gas to extend operating life, and often include a recycling system that helps regenerate the gas mixture. Eventually, the spent gas must be pumped out of the laser and replaced.

A gas fill can last many millions of shots for longer-lived mixtures such as xenon chloride, but at high repetition rates that does not add up to a long time. A 200-Hz laser in steady operation fires 720,000 pulses per hour. That means any excimer laser installation requires both gas supply and disposal.

7.7.3 Excimer Laser Configurations

Most excimer lasers are designed for particular applications. The 193-nm wavelength of ArF lasers is preferred for refractive surgery. Photolithography for integrated-circuit production largely uses 193-nm ArF lasers, but may use the longer 249-nm KrF wavelength for undemanding applications. These lasers are built into systems for surgery or photolithography. Excimer lasers are also used in a variety of other industrial applications including machining and annealing of surfaces. The type chosen depends on the application, and the laser is often sold as a system.

Excimer lasers may also be sold for laboratory use and other applications where flexibility is more important than optimizing production. Such lasers may be designed to operate with several different gas mixtures. The main gases used in such lasers are ArF,

KrF, XeCl, and XeF. Uses of the 157-nm F_2 laser have been limited by its low power level and the difficulty of using a wavelength at which the atmosphere is opaque.

7.8 NITROGEN LASERS

The 337-nm nitrogen laser resembles excimer lasers in having high gain and producing short ultraviolet pulses with high peak powers. It operates on a combined electronic-vibrational transition of molecular nitrogen (N_2), which is excited by a pulsed electric discharge at pressures from about 0.03 to one atmosphere. Excitation is efficient, and the gain is so high that it does not even require a rear cavity mirror (although adding one will double output power). However, the lower laser level has a 10- μ s lifetime, so nitrogen accumulates in that state and quickly self-terminates the population inversion and laser pulse. This limits laser efficiency to 0.1% or less, pulse length to a few nanoseconds, and pulse energy to about 10 mJ.

The low efficiency, pulse duration, and pulse energy limit average power of nitrogen lasers, which are smaller, easier to use, and much less expensive than excimer lasers. Operation is also simpler because they use harmless nitrogen rather than the corrosive and hazardous halogens required for excimer lasers. You can even use air, which contains 78% nitrogen, although laser operation is not very efficient.

Another unique attraction of the nitrogen laser is that it is the easiest laser for an amateur scientist or a student to build. Most gas lasers require very low pressure, hazardous gases, or precision mirror alignment. The nitrogen laser works best somewhat below atmospheric pressure, but it can operate at ambient pressure. The nitrogen laser requires a relatively simple power supply and does not require glass-blowing skills. You can find plans on the Internet by searching for “home-built nitrogen laser” or “nitrogen laser kit.” Be sure to be careful with the power supply; high voltage can kill.

7.9 CHEMICAL LASERS


Gas lasers powered by a chemical reaction are known as *chemical lasers*. As in a gas-dynamic CO_2 laser, a fuel reacts with an oxidizer while passing through a nozzle, and the exhaust expands

into a chamber, producing a population inversion. The process has the potential for tremendous power; military demonstrations have produced continuous-wave output of megawatts, although only for seconds at a time. That capability led to decades of military development, which eventually demonstrated megawatt-class chemical lasers for possible use as laser weapons.

However, chemical lasers proved to be bulky, and logistics officers eventually decided they were unsuitable for battlefield use because of their need for special fuels and the chemical hazards of the fuels and the reaction products. The Pentagon is instead developing weapon-class solid-state and fiber lasers that run on electricity, described in Section 13.10, with powers that can exceed 100 kilowatts. However, chemical lasers are interesting enough to deserve brief descriptions of the two best-developed types, the hydrogen fluoride/deuterium fluoride (HF/DF) laser and the chemical oxygen/iodine laser (COIL).

7.9.1 Hydrogen Fluoride (HF/DF) Lasers

Chemical reactions between a number of gases containing hydrogen and other gases containing fluorine can produce vibrationally excited hydrogen fluoride (HF) molecules that can be made to emit laser light. The processes are chain reactions that for the simple case of reacting molecular hydrogen (H_2) and molecular fluorine (F_2) produces a chain reaction:



Each step produces a vibrationally excited HF molecule (HF^*) that emits an infrared photon and a free atom of either hydrogen or fluorine, which initiates the next step in the chain reaction. The burning process resembles a rocket engine and can produce a steady laser beam as long as it has enough fuel. The exhaust gases flow through the nozzles and pass through a chamber with mirrors forming a laser cavity, with a laser beam emerging at the side, as shown in Figure 7-10. Expansion of the gas produces the population inversion needed for laser emission.




Figure 7-10. Laser beam is perpendicular to gas flow in an HF–DF chemical laser.

Excited HF^* molecules emit laser light on lines between about 2.6 and 3.0 μm , but the atmosphere absorbs those wavelengths. Replacing the common hydrogen-1 isotope with the heavier isotope deuterium (hydrogen-2) shifts the wavelength to a band between 3.6 and 4.0 μm , which the air transmits much better. HF lasers were planned for use in space, where no air would obstruct the beam; DF lasers were tested on the ground. However, as weapons they suffered from the bulk and complexity of the laser, the logistic problems of supplying the dangerous special fuels needed, and the need for large optics to focus the long wavelength onto a small area.

7.9.2 Chemical Oxygen–Iodine Lasers

The chemical oxygen–iodine laser (COIL) offered a crucial advantage—a much shorter wavelength of 1.315 μm . This would reduce the size of the output optics needed to focus the beam to a lethal intensity and propagate better through air.

The chemistry of a COIL is relatively complex. The COIL installed in a Boeing 747 for the U.S. Air Force’s Airborne Laser mixed chlorine gas (Cl_2) with a solution of hydrogen peroxide and potassium hydrogen dissolved in water. That released hot excited

oxygen molecules that were mixed with iodine (I_2) molecules to produce iodine atoms that absorbed energy from excited oxygen and emitted laser light at $1.315\ \mu\text{m}$. As in the HF/DF laser, the hot excited laser gas flowed rapidly through a laser cavity which extracted a powerful beam. However, Pentagon brass eventually decided that the need for special chemical fuels and the vulnerability of a weapon system in a Boeing 747 made the giant COIL laser unsuitable for its proposed mission in national missile defense.

7.10 OTHER GAS LASERS

Many other gas lasers have been demonstrated in the laboratory, and some have been manufactured commercially. There is not enough room to catalog them comprehensively, but a few types deserve brief mention.

7.10.1 Carbon Monoxide Lasers

The CO laser is a less-successful cousin of the CO_2 laser. Emitting on vibrational–rotational transitions of CO, which are mostly between 5 and $6\ \mu\text{m}$, it can be even more efficient than the CO_2 laser. Like the CO_2 laser, the CO laser can emit powerful continuous beams. However, CO lasers have been plagued by serious practical problems, including strong absorption of some lines by the atmosphere, and the need to cool the gas below room temperature for efficient operation. It remains the most powerful laser source in its wavelength range.

7.10.2 Far-Infrared Lasers

Many molecular gases have laser lines that can generate relatively low powers at wavelengths between about 30 and $1000\ \mu\text{m}$. Instead of being excited by an electric discharge, they are excited by light from another laser, usually CO_2 , which produces a population inversion leading to laser action.

This family is often called *far-infrared lasers*, and they work very similarly. The pump laser is tuned to a narrow band to match an absorption line that excites the gas to an excited vibrational level. The laser transition occurs between two rotational levels in the excited vibrational state. Typically, the pump beam enters at one

end of the tube, and the far-infrared beam exits at the other end. The pump and output beams pass through holes in the cavity mirrors. These lasers are generally used in laboratory applications.

7.10.3 Diode-Pumped Alkali Lasers

One of the newest types of gas lasers is a new variation on one of the very oldest proposals for a laser—optical pumping of an alkali metal vapor. The original idea was to use light from an alkali metal lamp to excite metal vapor in a laser cavity. A number of alkali metals were proposed, and eventually a 7.2- μm cesium was demonstrated in 1962, but it was inefficient and impractical, an obscure footnote in laser history.

What revived interest in alkali vapor lasers was the invention of high-efficiency semiconductor lasers and their use in optical pumping of other lasers, as described in Section 6.3.1. William Krupke looked at the spectra of alkali vapors and realized that the single valence electron in their outer shells had pairs of energy levels spaced so closely that they allowed extremely efficient optical pumping, with the emitted photon carrying all but a tiny fraction of the energy from the pump photon. For example, rubidium has a pump line at 780 nm and emits laser light at 795 nm—so only 1.9% of the pump photon energy is lost. Military labs have demonstrated high power, and there is some hope of the technology reaching the megawatt-class power sought for national missile defense. The new family of *diode-pumped alkali lasers* are a long way from becoming laser weapons, but their very high efficiency may lead to new applications for gas lasers.

7.11 WHAT HAVE WE LEARNED?

- Gas lasers vary widely in performance, depending on the laser species.
- Sales of low-power gas lasers have declined as semiconductor and solid-state lasers replaced them.
- High-power ultraviolet excimer lasers and infrared CO₂ lasers remain important products.
- Low-power gas lasers typically are glass tubes filled with gas at low pressure powered by a high-voltage discharge along their length.

- High-power gas lasers are typically excited by a discharge across laser cavity.
- Atomic gases emit on electronic transitions of neutral or ionized atoms.
- Molecular gases usually emit on vibrational or rotational transitions, but some emit on electronic transitions.
- Many gas lasers contain other gases in addition to the one that emits light.
- Some gas lasers require gas to flow through their tubes to remove contaminants and waste energy, and replenish the laser gas.
- In many gas lasers, windows are mounted at Brewster's angle to minimize losses, and mirrors are separate from the tube. Brewster windows polarize the beam.
- Thermal motion of gas atoms causes Doppler broadening, which broadens laser bandwidth.
- The red He–Ne laser remains the most common gas laser, but red diode lasers have replaced it for many applications.
- In a He–Ne laser, the electric discharge excites helium atoms, which transfer energy to neon atoms, which emit the low-power laser beam. It has low gain.
- Coherence length is an important advantage of He–Ne lasers.
- Argon- and krypton-ion lasers emit shorter wavelengths and more powerful continuous beams than He–Nes. Both emit on many lines.
- He–Cd lasers emit milliwatts of blue and ultraviolet light. The tube must be heated to vaporize cadmium, which is solid at room temperature.
- Neutral metal atoms emit short repetitive pulses in copper-vapor lasers, which have very high gain.
- The CO₂ laser is the most versatile and powerful gas laser, emitting on vibrational transitions at 9 to 11 μm .
- Sealed-tube CO₂ lasers can emit up to 100 W. Longitudinal-flow CO₂ lasers emit higher powers, and transverse-flow lasers produce up to 20 kW.
- TEA CO₂ lasers emit high-power pulses when a transverse discharge excites gas at pressure near one atmosphere.
- CO₂ lasers are used in industry and medicine.
- Excimers are short-lived rare-gas halide molecules that dissociate when they emit ultraviolet photons and drop to a lower energy state.

- Excimer lasers are the most powerful pulsed ultraviolet lasers; they are used in medicine and semiconductor fabrication.
- Chemical lasers get their energy from chemical reactions.
- Nitrogen lasers have high gain and emit 337-nm pulses with high peak power, but have low average power and efficiency.
- Diode-pumped alkali lasers have a potential for higher efficiency than other gas lasers, and perhaps reaching kilowatt-class power, but they face formidable challenges.

WHAT'S NEXT?

In Chapter 8, we will learn about solid-state crystalline and glass lasers, which have different properties than gas lasers.

QUIZ FOR CHAPTER 7

1. Which of the following general statements is not true about gas lasers?
 - a. Most gas lasers are excited electrically.
 - b. Only materials that are gaseous at room temperature can be used in gas lasers.
 - c. The laser tube cannot contain atoms or molecules other than the one species that emits light.
 - d. Gas lasers can emit continuous beams.
 - e. Gas lasers can emit ultraviolet, visible, or infrared light.
2. Which of the following lasers can emit the shortest wavelength in a continuous-wave beam?
 - a. Helium–neon
 - b. Helium–cadmium
 - c. Nitrogen
 - d. Argon fluoride
 - e. Carbon dioxide
3. Which of the following lasers emits the shortest wavelength?
 - a. Helium–cadmium
 - b. Krypton fluoride
 - c. Nitrogen
 - d. Argon fluoride
 - e. Argon ion

4. Gas lasers have electronic transitions across approximately what wavelength range?
 - a. 150–1500 nm
 - b. 250–1500 nm
 - c. 400–700 nm
 - d. 400–5000 nm
 - e. 150–15,000 nm
5. Why would an engineer select a red He–Ne laser over a red semiconductor laser?
 - a. The He–Ne is smaller.
 - b. The He–Ne uses power more efficiently.
 - c. The He–Ne has better beam quality.
 - d. The He–Ne costs less.
 - e. Only if designing a retro optics lab.
6. The strongest emission lines of the argon-ion laser are at
 - a. 275.4, 300.3, and 302.4 nm
 - b. 488.0 and 514.5 nm
 - c. 446 and 532 nm
 - d. 632.8 and 1152 nm
 - e. 647.1 and 799.3 nm
7. Application of a high voltage initially ionizes the gas in an argon laser. Then what happens?
 - a. The laser emits intense light in the vacuum ultraviolet until the argon ions drop to the metastable upper laser level.
 - b. Laser operation begins in the near ultraviolet and then voltage is reduced for laser operation at visible wavelengths.
 - c. The voltage is increased further until the argon gas crosses the threshold for laser emission.
 - d. The voltage is decreased, and the resulting current produces laser emission.
 - e. The power supply burns out because argon-ion lasers do not require high voltage.
8. What is fundamentally different between a helium–cadmium laser and a He–Ne laser?
 - a. He–Cd is pulsed; He–Ne emits continuous wave.
 - b. He–Cd emits a red beam; He–Ne emits a blue beam.
 - c. He–Cd oscillates on only one line; He–Ne can oscillate on several lines.
 - d. He–Cd lasers are solid state; He–Ne lasers are gas.
 - e. He–Cd lasers must be heated to generate light; He–Ne lasers do not require heating.

9. CO₂ lasers are excited by
 - a. Longitudinal electric discharge
 - b. Transverse electric discharge at low pressure
 - c. Transverse electric discharge at high pressure
 - d. Transverse electric discharge in moving gas
 - e. All of the above
10. Which gas laser generates the highest power for industrial applications?
 - a. Carbon dioxide
 - b. Copper vapor
 - c. Argon–fluoride
 - d. Argon ion
 - e. Carbon monoxide
11. What happens when excimer molecules drop to the lower laser level?
 - a. They emit light and dissociate because they are unstable in the ground state.
 - b. They emit light and accumulate in the ground state, ending the population inversion so laser action stops.
 - c. They absorb stimulated emission, reducing gain and terminating the laser pulse.
 - d. They dissociate, then each of the two atoms that were in the molecule remains in an excited state to produce a population inversion.
 - e. They react with each other and immediately return to the excited state.
12. Which of the following molecules is not used in an excimer laser?
 - a. ArF
 - b. HF
 - c. KrF
 - d. XeCl
 - e. XeF

CHAPTER 8

SOLID-STATE LASERS

ABOUT THIS CHAPTER

In this chapter, you will learn about solid-state lasers, in which light is emitted by atoms embedded in a crystal, glass, or other transparent solid. This chapter first defines solid-state lasers more precisely and explains their general operation. Then it describes the most important solid-state lasers, including the classic ruby laser, neodymium, ytterbium, and a family of related materials in which laser emission can be tuned across a range of wavelengths. Although fiber lasers and amplifiers technically are solid-state lasers, they differ in important details, so they are treated separately in Chapter 9. This chapter covers solid-state lasers with nonfiber form factors, including rods, slabs, and thin disks.

8.1 WHAT IS A SOLID-STATE LASER?

The first step in understanding solid-state lasers is to realize that the laser world defines “solid-state” differently than electronic engineers or physicists. Solid-state physics occurs in solids. Solid-state circuits are semiconductor devices that conduct electricity and perform electronic operations. In electronics, semiconductors are considered solid-state devices in contrast to vacuum tubes, where electrons move through a vacuum rather than through a solid.

Laser terminology separates semiconductor lasers from lasers made from other solids because the two classes of lasers are

powered differently. You learned in Chapter 6 that most lasers get their energy from electricity or light from outside sources. Most gas lasers described in Chapter 7 are pumped by electric discharges created when a high voltage ionizes the gas so that it carries a current. This chapter and Chapter 9 cover solid-state lasers powered by light; Chapter 10 covers semiconductor lasers powered by electricity. Those differences lead to important differences in how lasers work and can be used.

We will begin this chapter by looking carefully at these differences.

8.1.1 Differences among Solids

The electrical behavior of solids falls into three broad categories: conductors, semiconductors, and insulators. *Conductors* are metals like copper and silver with free electrons in their outer shell that can easily move through the solid to conduct electricity. *Semiconductors* are nonmetals such as silicon and germanium or compounds such as gallium arsenide, with electrons in their outer shells bound so loosely to the atom that applying a few volts across the material can start a current flowing. *Insulators*, also called *dielectrics*, are materials such as glasses and crystals that hold their outer-shell electrons so tightly that only a very high voltage could pull the electrons free, therefore they normally cannot conduct a current.

Electron bonding also affects how solids interact with light. The tight bonding of electrons that prevents current flow also means that the lowest energy levels in insulating materials are so high that visible light does not have enough energy to excite electrons. Thus visible light can pass right through a dielectric solid, making it transparent, as shown in Figure 8-1. Examples include glasses and many types of crystals.

Solid-state lasers need to be transparent so that pump light can pass through the solid to transfer energy to the laser species. However, the laser species must absorb the pump light to collect that energy. To meet those seemingly contradictory requirements, solid-state laser materials are made by adding a small amount of the light-absorbing laser species to the transparent *host material*. In the case of Maiman's ruby laser, the host material is aluminum oxide or sapphire and the laser species is chromium, which makes ruby when added to the crystal. Knowing how much chromium to add to the crystal is part of the science of solid-state lasers.




Figure 8-1. Electronic energy levels in insulating solids and semiconductors.

In contrast, as shown in Figure 8-1, semiconductors usually can conduct electricity as well as transmit laser light; therefore, an electric current can produce a population inversion and laser light can escape through the semiconductor. (Semiconductors can also absorb other visible light, as in solar cells.) Chapter 10 explains more about semiconductor lasers.

Conductors are metal-like materials in which the outer electrons are not bound strongly to atoms and are conducted freely on the surface or within the solid. However, they do not transmit light, preventing laser action or transparency in metals.

These differences are not as sharp as the definitions may make them sound. Materials we consider semiconductors may be transparent to infrared light, because those photons lack enough energy to release electrons from their bonds to atoms in the crystal so that they can conduct current. Look at Figure 5-8, and you can see that silicon, germanium, and gallium arsenide, materials we use as semiconductors, are transparent in the infrared.

8.1.2 Ruby and Other Solid-State Lasers

To show how a solid-state laser works, let's look at the very first laser. Maiman's ruby laser was pumped by a burst of white light from a photographic flashlamp, as shown in Figure 8-2. Light from the spring-shaped flashlamp illuminated the little ruby rod that it surrounded. Chromium atoms in the ruby absorbed green and violet light from the lamp and emitted red laser light. Thin metal films coated on the ends of the rod formed a reflective laser cavity, with the beam emerging through a small hole in the middle of the film on one end. The laser pulsed only during the flash of the flashlamp.

Today's diverse family of solid-state lasers has been refined in many ways, but the same principles underlie their operation. Photons from an external source excite atoms dispersed in a solid host, producing a population inversion. Spontaneous emission triggers a cascade of stimulated emission, which oscillates between the mirrors in a laser cavity, and produces a beam.

Solid-state lasers can take various forms. Often the solid is rod-shaped, but in the fiber lasers described in Chapter 9 it is a long thin fiber that works somewhat differently. Solid-state lasers may also be shaped as disks or slabs, as described later in this section.

Solid-state lasers can be made from many different materials, but today most solid-state lasers are made from a handful of




Figure 8-2. Structure of the first ruby laser.

materials chosen because they work well. The lasers can be packaged with nonlinear devices to generate other wavelengths from the fundamental laser frequency, as described in Section 6.5.3. Solid-state laser materials can be used in laser amplifiers as well as in oscillators, particularly for optical amplification in fiber-optic communication systems, or to produce short high-energy pulses, as described in Section 6.4.

8.1.3 Optical Pumping and Light Conversion

Optical pumping has been the key to success for most solid-state lasers. Typically, its importance comes from its ability to convert light from one form into another form that is more useful or valuable. The ruby laser turns bright pulses of white light from a lamp spreading in all directions into a tightly focused beam of coherent red light. Therefore, it changes the wavelength, the organization, and the direction of the light.

Optical pumping with a flashlamp is rather inefficient. Flashlamps convert only a fraction of the input electrical energy into light, and laser materials absorb only a fraction of the lamp's broad emission spectrum. The lamp light spreads in all directions, and only a fraction of it can be focused onto the laser rod or slab. Add it all together, and ruby lasers convert no more than one percent of the electrical energy into laser light. Other solid-state materials can convert up to a few percent of input lamp light into laser energy, but that is not very efficient, and it long limited the use of solid-state lasers.

Pumping with semiconductor diode lasers offers three big advantages over lamps. One is that a semiconductor diode converts much more of the input electric power into light than a lamp; in some cases, over half the input energy emerges as light. (That is the reason why LED lighting is far more efficient than incandescent lamps.) A second is that the laser light is concentrated in a narrow range of wavelengths, which can be matched to narrow bands where solid-state materials absorb most strongly. A third is the high directionality of the laser beam, which eases focusing into rods and fibers. The combined high efficiency has made diode pumping preferred for most solid-state laser uses.

Why not use semiconductor lasers directly rather than using them to pump solid-state lasers? One reason is that current semiconductor lasers do not generate the high-quality, tightly focused beams that solid-state lasers can produce, and that for many

applications, from cutting metals or plastics to pointing out items on a screen. Another is that optical pumping can be an efficient way to shift the light to a more useful wavelength, particularly when combined with nonlinear optics. Meanwhile, engineers are working to improve the beam quality of semiconductor lasers, as you will learn in Chapter 10.

8.1.4 The Solid-State Laser Industry

The solid-state laser industry (including fiber lasers) has grown greatly in recent years, and that growth has affected its structure. A handful of types are widely used in industry with huge sales, notably ytterbium-fiber lasers and neodymium solid-state lasers. Titanium-sapphire lasers are widely used in research. Erbium-doped fiber amplifiers are used in telecommunication networks.

Most types have found more limited or specialized markets, and many are still in development that may lead to expanded applications in the future. These are included in the list in Table 8-1.

8.1.5 The Inevitable Exceptions

Inevitably, ingenious optical scientists will come up with exceptions to the usual rules, and one of them is using optical pumping to energize semiconductor lasers. Optical pumping has long been used in semiconductor laser research because it is simpler and easier for making a few test devices. More recently, developers have found that optical pumping of semiconductors can produce wavelengths not possible with electrical pumping. Section 8.5 covers these optically pumped semiconductor lasers (OPSLs).

8.2 SOLID-STATE LASER MATERIALS

A solid-state laser material generally has two essential components: light-emitting atoms called the *active species* and a *host* solid in which those atoms are embedded called the *host*. In the ruby laser, the chromium atoms that give the ruby its reddish color are the active species, and aluminum oxide (Al_2O_3), also known as sapphire, is the host. The active species is identified first in most solid-state laser materials, often by its chemical symbol, followed by the host material. For example, a titanium-sapphire laser in which




Figure 8-3. Light-emitting atoms in a solid-state laser are embedded in a crystalline or glass host.

titanium is the active species and sapphire is the host, is usually written Ti–sapphire (or Ti:sapphire). Similarly, an Nd:YAG (or Nd–YAG) laser has neodymium as the active species doped into a synthetic crystal called yttrium–aluminum garnet (YAG or $\text{Y}_3\text{Al}_5\text{O}_{12}$).

Atoms of the active species are dispersed within the host, as shown in Figure 8-3. If the host is a crystal, the active species typically replaces an atom with similar valence in the host crystal. Thus, chromium atoms occupy aluminum sites in ruby because both have +3 valences. The active species is called a *dopant* because it is added in small quantities and replaces only a small fraction of atoms in the host.

Both the active species and the host affect the properties of a solid-state laser. The wavelength of the laser species depends mainly on the active species, but may be shifted by interactions with the host. Thermal and mechanical properties of the laser depend largely on the host, which usually makes up 99% or more of the laser volume.

Three types of solids are used as hosts: glasses, synthetic single crystals, and ceramics. A *glass* is a disordered or amorphous solid formed when certain liquids solidify under proper conditions. The best-known glasses are blends of oxide compounds dominated by silica (SiO_2), but blends of non-oxide compounds such as fluorides can also form glasses. Single crystals have better thermal properties than glasses, but they must be grown, a time-consuming process which takes much longer than making glass, and are limited in the size that can be produced. Some laser materials such as Nd–YAG can be produced in polycrystalline or *ceramic* forms that are clear

and transparent, like glasses or crystals. Ceramics have thermal properties similar to those of single crystals but can be produced in much larger sizes, allowing them to produce higher sustained power than either glass or single-crystal materials.

It takes considerable time and money to optimize the combination of an active species and a host for solid-state laser performance, so developers often stay with established combinations that are well characterized.

8.2.1 Active Species and Transitions in Solid-State Lasers

The active species in a solid-state laser absorbs pump light that excites it to a higher energy level, from which it generally drops to the metastable upper laser level, before stimulated emission drops it to the lower laser level, as you saw earlier in Figure 3-2. The most important active species are metals: chromium, neodymium, erbium, ytterbium, titanium, holmium, thulium, and iron. All but chromium, titanium, and iron are from a group of chemically similar elements called the rare earths. Table 8-1 lists commercial solid-state lasers using these materials. Note that only a handful of types are in wide use; most others have a few specialized uses. Many other similar solid-state lasers have been demonstrated in the laboratory, but have yet to gain commercial acceptance.

The light-emitting species often are called “ions” because they have a nominal chemical valence in host crystals. The most common valence state is +3, including Cr^{+3} , Ti^{+3} , Nd^{+3} , Er^{+3} , Yb^{+3} , Ho^{+3} , and the other rare earths. These elements form covalent bonds in solid-state laser materials, so this ionization state is more nominal than real. The “ions” are fixed in the crystal, and the “missing” electrons are nearby, typically bonded to oxygen atoms.

Solid-state laser transitions are between electronic levels of the active species or ion. Valence bonds with the host material and interactions with adjacent atoms affect energy levels, changing laser wavelength. For example, neodymium emits at 1054 nm when it is doped into phosphate-based glass, and at 1064 nm in YAG.

Many solid-state laser transitions are between pairs of well-defined energy levels and have gain over a narrow range of wavelengths, such as the 1064-nm Nd–YAG transition. However, interactions between vibrational energy levels in the host and electronic energy levels in the active species can produce *vibronic* states, spread over a range of energies, described in Section 8.6.

Table 8-1. Commercial solid-state lasers and their uses (includes lines used in fiber lasers)

Output (nm)	Name	Light emitter	Host	Usage
440–600	OPSL or VECSEL	OPSL or VECSEL	Semiconductor	Specialist
675–1100	Ti–sapphire	Titanium	Sapphire	Tunable and ultrashort
694	Ruby	Chromium	Sapphire (ruby)	Very rare
701–826	Alexandrite	Chromium	Alexandrite	Specialist
1030–1100	Yb–fiber or Yb–YAG	Ytterbium	Glass or YAG	Widespread, industrial
1047, 1054	Nd–yttrium lithium fluoride (YLF)	Neodymium	YLF	Specialist
1060	Nd–glass	Neodymium	Glass	Specialist
1064	Nd–YAG	Neodymium	YAG	Widespread, industrial
1064	Nd–vanadate	Neodymium	YVO ₄	Widespread, industrial
1530–1570	Er–fiber or Er–glass	Erbium	Glass, typically drawn into fiber	Widespread, communications
1800–3400	Cr–ZnSe/S	Chromium	Zinc selenide/sulfide	Specialist
1900–2050	Tm–fiber	Thulium	Glass, drawn into fiber	Specialist
2050–2150	Ho–fiber	Holmium	Glass fiber	Specialist
3400–5200	Fe–ZnSe/S	Iron	Zinc selenide/sulfide	Specialist

8.2.2 Host Materials

Host materials must meet stringent optical, thermal, and mechanical requirements. They must be reasonably transparent to the pump light and absorb little light at the laser wavelength.

The main optical requirement is that the host should transmit both input pump light and the output laser beam. Crystalline hosts are chosen for their transparency at the wavelengths of interest. Conventional oxide glasses are widely used at wavelengths to slightly beyond 2 μm , but absorb at longer wavelengths. Fluoride glasses and compounds such as zinc sulfide and zinc selenide are used as hosts when transparency is needed at longer wavelengths.

Thermal properties of the host are important because most pump energy in many solid-state lasers winds up as heat, and the dielectric solids used as hosts are thermal insulators, so they cannot dissipate heat readily. High temperatures can affect energy-level populations, usually decreasing laser gain, efficiency, and output power. Internal accumulation of heat can also change refractive index and degrade beam quality. In extreme cases, heat differentials can warp or crack the laser medium.

Crystalline laser hosts generally dissipate heat more efficiently than glass, so crystal laser rods can operate at higher powers and higher pulse repetition rates than glass rods of equal size. However, glass is easy to draw into long, thin optical fibers, which can dissipate heat efficiently because of their large surface area and relatively small volume.

Glass has the advantage that it is easier and much less expensive to shape into large chunks. Crystalline materials must be grown slowly under carefully controlled conditions to produce defect-free bulk crystals called *boules*, which must be machined into rods or other shapes suitable for laser use.

An alternative to single crystals is producing polycrystalline ceramic versions of laser crystals by heating and compressing powders or other granular material and then letting them solidify to form transparent laser material. The optical quality can be quite good, and ceramics can be produced in much larger sizes than single-crystal materials.

8.2.3 Pump Bands and Light Absorption

Efficient optical pumping requires a light source that matches the absorption bands of the laser material. Figure 8-4 shows absorption in one of the most important solid-state lasers, Nd-YAG, compared with the narrow spectrum of a diode pump laser and the broad spectrum of a flashlamp pump. The sharp spikes are absorption on narrow transitions between pairs of sharply defined electronic energy levels. The broader bands are weaker absorption on more dispersed bands.

Historically, the first optical pumping sources available for solid-state lasers were lamps. *Flashlamps* produce brief, intense flashes when an electrical pulse passes through a gas and are widely used in photography. *Arc lamps* emit intense light continuously as a steady electric discharge flows through gas; applications include




Figure 8-4. Broadband absorption of YAG doped with 0.62% Nd, compared with output of typical pump sources. (Courtesy of Northrop Grumman Corp.)

movie projectors. These lamps emit bright light at wavelengths absorbed by the broad absorption bands of neodymium and similar broad absorption bands in ruby.

Once bright semiconductor lasers were developed, they offered an alternative. Their narrow-line emission could pump the strong 808-nm peak of Nd–YAG shown in Figure 8-4. Concentrating pump light on a strong absorption line makes laser pumping much more efficient than lamp pumping. Another advantage of diode lasers is that they convert a large fraction of the input electrical energy into light, further enhancing diode-pumping efficiency. Some solid-state laser media such as ytterbium-doped materials covered in Section 8.4.3 lack the broad absorption bands needed for lamp pumping, so they did not become widely used until diode pumping became available.

8.2.3.1 Flashlamp Pumping

Flashlamp pumping has lost market share to diode-laser pumping, but it has strengths including its ability to generate high pulse energies.

The most common type of flashlamp now used for optical pumping is a long linear tube filled with a rare gas such as krypton or xenon. A brief high-voltage electrical pulse applied across the ends of the tube causes electrical breakdown of the gas, which

conducts a current pulse that generates a brief flash lasting about a millisecond (0.001 second). The laser pulse produced inevitably is shorter than the lamp pulse because it takes time for the lamp power to exceed laser threshold. Pulse-shortening accessories such as the Q switches described in Section 5.5.2 can make the pulses much shorter. Pulse repetition rate is limited by the switching electronics and the lamp.

8.2.3.2 *Arc Lamp Pumping*

Continuous-wave solid-state lasers can be pumped with electric arc lamps, in which a steady electric current flowing through a gas-filled tube produces intense light. Intense light is needed to sustain a continuous population inversion in solid-state laser materials such as Nd-YAG, which are capable of continuous laser operation. Solid-state lasers pumped with arc lamps can be pulsed with Q switches, mode lockers, or cavity dumpers at higher repetition rates than possible with flashlamps.

Arc lamps for laser pumping are available commercially, but the lamps are short-lived and inefficient, so they have largely been replaced by diode pumping, described in the next section.

8.2.3.3 *Semiconductor Diode-Laser Pumping*

Diode pumping with semiconductor lasers has gained wide acceptance for pumping solid-state and fiber lasers because of their high efficiency and small size. Their prime attraction is a much higher efficiency than lamps that comes from four advantages. First, diode lasers can convert much more of the input electricity into light than flashlamps, in some cases well over 60%. Second, diodes can be made to concentrate their emission on the wavelengths where solid-state media have their peak absorption, as shown in Figure 8-4. Third, diode lasers are highly directional and can be coupled into the light-guiding cores of optical fibers. Finally, diodes can reach high output powers and be assembled into arrays that can combine to generate high pump powers. These advantages have been particularly important for the fiber lasers described in Chapter 9.

Combining the high efficiencies of diode emission and optical pumping means that the wall-plug efficiency of diode-pumped lasers can reach or exceed 20% for bulk solid-state lasers and 40% for fiber lasers. That compares a few percent at best for flashlamp pumping.

8.2.3.4 Other Laser Pumping

Other types of lasers can be used to pump solid-state lasers when no suitable semiconductor laser matches their pump bands. Fiber lasers are among the sources used, particularly for solid-state or fiber lasers emitting at wavelengths longer than 2 μm .

8.3 SOLID-STATE LASER CONFIGURATIONS

Traditionally the laser medium in solid-state lasers was a long, thin rod, often thinner and a bit shorter than a round wooden pencil, with two parallel polished ends in a resonant cavity. The long length allowed amplification of the stimulated emission over a reasonable distance, and the thin diameter helped dissipate any accumulated heat. The design was developed for linear flashlamps, which could be placed alongside the rod in a reflective cavity. Later the design was adapted for diode pumping. That tried and true design remains widely used.

One established alternative is shaping the laser medium into a rectangular, elliptical, or cylindrical slab about as thick as a traditional laser rod, with pump light illuminating the wide flat sides. Limiting the slab thickness avoided heat dissipation problems that had occurred with fat rods, allowing them to produce higher power. You will read more about them in Section 8.3.2.


A much newer technique is reshaping the laser medium into a very thin disk that can be pumped by high-brightness diode lasers while resting on a heat sink that removes waste energy. This has proved successful in generating kilowatt-class continuous powers, a major advance over traditional rods. You will learn about this approach in Section 8.3.3 and in Section 8.5.

A very successful new approach is to stretch the rod into a longer and thinner optical fiber. The fiber laser design is flexible, so that it can be coiled, and also includes an internal light-guiding structure which improves performance and efficiency by how it confines both pump light and laser output, as you will learn in Chapter 9.


The performance of these configurations depends on the source used for optical pumping, which is the only practical way to produce a population inversion in a solid that does not conduct electricity. Let us look at these in more detail.

8.3.1 Pumping Laser Rods

Side pumping matches the geometry of readily available linear flashlamps to the long thin shape of a laser rod. The lamp is placed parallel to a laser rod inside a reflective cavity in one of the variations shown in Figure 8-5. Some pump light travels directly from the lamp to the laser rod, with light reflected from the inside of the reflective cavity toward the laser rod. Pumping is most efficient if the cavity is elliptical in cross-section and the laser rod and the lamp are placed at the two foci, so a light ray from the lamp at one



(A) Closely spaced linear lamp and rod (side view).



(B) Closely spaced linear lamp and rod (end view).

Figure 8-5. Flashlamp pumping of laser rods.




Figure 8-5. (Continued)

focus is reflected to the laser rod at the other focus. Alternatively, two lamps can be arranged on opposite sides of the laser rod in a dual elliptical cavity. This approach can be used for either flash-lamps or arc lamps, but requires cooling at high power or with arc lamps.

The same side-pumping geometry can be used for pump diodes, with the diodes or arrays of diodes arranged along a line to replicate the shape of a linear lamp, like LED lighting can be arranged in a long thin cylinder to reproduce the shape of a linear fluorescent lamp, as shown in Figure 8-6 for an array of chips fabricated on a single laser wafer.

An alternative approach to diode-pumping of small laser rods is to aim output of a single diode laser through the end of the rod along the length of the laser, as shown in Figure 8-6. That couples the pump beam into the rod, but requires coatings on the rods to control the light—one to reflect pump light back into the laser cavity at the other end while coupling laser light through the output mirror and other to reflect the oscillating beam back toward the output mirror while transmitting the pump beam. This approach is used in small lasers, notably in green laser pointers, which contain a 1064 nm neodymium laser that is frequency-doubled to the green.




Figure 8-6. Diode-pumping of solid-state laser rods.

Although diode pumping can greatly improve the performance and efficiency of solid-state rod lasers, the rod geometry has fundamental limits on heat dissipation for large sizes. Diode pumping reduces the heat load, but even crystals and ceramics conduct heat poorly. The refractive index of solids changes with temperature, so temperature differences within a rod can cause index differentials that distort or disperse the beam in the solid. These thermal properties limit the output power available from solid-state laser rods and lead to development of other laser geometries.

8.3.2 Pumping Slab Lasers

The rod geometry for solid-state lasers suffers from fundamental limits on its size arising from the need to dissipate the heat that accumulates within an optically pumped solid. The larger the block of solid, the longer it takes for the heat to work its way out to the edges. Early efforts to develop solid-state lasers that could produce very high energy pulses were thwarted by heat that accumulated in

laser logs, causing optical effects that limited power and in extreme cases shattering the glass.

Molding the laser material into a slab only about as thick as a thin laser rod eased the problem. Slabs could be rectangular, elliptical, or circular as long as they were thin. Pump light could illuminate the wide sides of the slab, and heat could dissipate from the sides or be removed by chillers or fans. Slab lasers were first used in optical amplifiers with banks of flashlamps illuminating the sides, with cooling between pulses.

The development of diode pumping opened new opportunities for the slab geometry. Diode pumping is much more efficient than lamp pumping, and deposits much less waste heat, easing cooling requirements. Diode pumping also allows other geometries. Diode-pumped slabs can be used as oscillators as well as amplifiers, and the diodes can pump from the edges as well as the wide sides. In some cases, the slabs may bounce the laser beam on a zigzag path in the plane of the slab to extract light energy from more of the slab volume. When cooled, such slabs can generate up to kilowatts of continuous power.

8.3.3 Thin Disks and Folded Surface-Emitting Cavities

Diode pumping also opened a new option called the *thin-disk laser*, shown in Figure 8-7. The thin disk at the bottom of the figure is a centimeter-scale disk of laser medium a fraction of a millimeter thick. Pump diodes illuminate its transparent top face, while the reflective bottom face functions as the rear cavity mirror in the laser oscillator. The thin disk rests on a heat sink which removes waste heat from the disk.

Two types of laser medium are used in thin-disk lasers. One is a thin layer of nonconductive solid-state material, often ceramic YAG doped with neodymium or ytterbium, which can reach kilowatt powers. The other is an *optically pumped semiconductor laser* or *OPSL* which is composed of the same type of semiconductors used in diode lasers, but lacks the internal structures needed to conduct current. You will learn more about OPSLs in Section 8.5. The rest of this section focuses mainly on conventional dielectric solid-state thin disks.

Figure 8-7 shows a single thin-disk laser being pumped by two pump diodes, which illuminate it from the sides. The bottom surface of the thin disk is reflective, so light resonates within a vertical




Figure 8-7. Diode-pumping of a single thin-disk laser mounted on a heat sink.

cavity with the output mirror at top as shown. Thin disk lasers can produce even higher power if two or more are placed in series in a W- or zigzag-shaped cavity with the reflective-bottomed disks at the bottom and with pump diodes and cavity mirrors on top, so light passes through each of the thin slabs twice while resonating in the bent cavity. The powers of individual slabs add together up to a total of four slabs.

A major attraction of dielectric thin-disk lasers is their ability to convert more than 50% of the light from pump diode arrays into a single high-quality output beam for applications such as laser machining. In practice, the cavity optics must be more complex than the simple arrangement shown in order to generate high power efficiently. The disk absorbs only part of the pump light on a single round-trip through the disk, so the laser cavity redirects the reflected light back onto the disk so more can be absorbed. With pump light spread evenly across the face of the disk, the thermal gradient is linear from top to bottom, avoiding thermal lens effects

that could distort the beam. Active cooling of the heat sink removes waste heat.

The cavity mirrors transmit only a few percent of intracavity light to the output beam, leaving much higher power circulating within the laser cavity. This high intracavity power can be used for harmonic generation.

Thin-disk lasers can generate pulsed or continuous output. Dielectric types using ytterbium-doped ceramics can generate continuous output up to 6 kW from a single disk and up to 16 kW from a four-disk version. Laboratory versions have demonstrated more than 10 kW from single disks and tens of kilowatts from multidisk lasers and have attracted military interest for possible weapon use.

8.3.4 Fiber Lasers

Diode pumping also enabled the development of high-power fiber lasers. Technically, they are solid-state lasers, but they work somewhat differently and have become so important that they are worth a chapter of their own and are described in Chapter 9.

8.4 MAJOR SOLID-STATE LASER MATERIALS

So far, this chapter has focused on solid-state lasers in general, without going into much detail on the specific materials. However, solid-state laser materials determine the details of how solid-state lasers work, therefore, we need to look at them carefully. This section covers solid-state materials that usually generate a fixed wavelength. Sections 8.5 through 8.7 describe different material families, including OPSLs and a family of solid-state lasers with particularly broad bandwidth and the ability to produce ultrashort pulses.

8.4.1 Ruby Lasers

Ruby was the first laser material, but it is far from ideal. It is a three-level laser system, with the ground state also being the lower laser level, so its efficiency normally is lower than that in four-level lasers such as neodymium. Ruby is limited to pulsed operation at low repetition rates, so it is used only for a few applications that require its high-energy red pulses, but it is of historic importance.

Natural ruby is a gemstone, but ruby lasers use a synthetic ruby made by growing the crystal from aluminum oxide with 0.01 to




Figure 8-8. Energy levels of chromium atoms in ruby laser.

0.5 percent chromium added. Figure 8-8 shows ruby energy levels. Ground-state chromium absorbs light in bands centered near 400 and 550 nm, well matched to flashlamps. After being excited to the upper excited state, chromium drops to a pair of closely spaced metastable levels, releasing energy to the crystalline lattice. Then stimulated emission drops them to the ground state, mainly on the lower-energy laser line at 694.3 nm. Once the pump lamp stops, chromium atoms accumulate in the ground state and absorb stimulated emission, ending the laser pulse.

The long lifetime of the upper laser level allows a ruby rod to store energy effectively, and ruby lasers can be Q-switched to produce short pulses with typical energies of a few joules. Q switching limits pulse duration to 10 to 35 nanoseconds and pulse energies to a few joules, but yields peak power in the 100-megawatt range. Ruby can also be used in oscillator amplifiers.

Ruby is limited to a wall-plug efficiency below 1%, but the crystal conducts heat well and resists damage from excess optical energy as long as its surface is kept clean of dirt that could absorb laser energy. The laser properties of a three-level system degrade




Figure 8-9. Major energy levels in 1.06-micrometer neodymium laser showing two pump lines. Pump diodes are more efficient at 808 nm, so that line is used more than the shorter 750 nm line.

rapidly as temperature increases, so ruby must be operated at no more than a few pulses per second to prevent heat build-up.

8.4.2 Neodymium Lasers

The most common bulk solid-state lasers use neodymium atoms in various solid hosts. Figure 8-9 shows its four-level energy-level structure along with the two strongest pump lines at 750 and 808 nm. The most widely quoted value for the laser transition is 1064 nm, which is the value in Nd-YAG. The figure rounds the wavelength to 1.06 μm as a reminder that interactions between neodymium atoms and hosts can change the wavelength by about 1%. The figure does not show a number of secondary laser transitions that also occur in neodymium but are much weaker than the main line.

The two primary pump bands shown are in the 700–850 nm range. Each raises neodymium atoms to one of the two excited levels. The detailed absorption spectrum is more complex, as is shown in Figure 8-4, with a strong peak at 808 nm. Diode pumping uses the 808 nm line where neodymium absorption is strongest and diode emission is efficient. Lamp pumping is spread across the visible spectrum.

The optically excited neodymium atoms quickly drop to the metastable upper laser level, releasing energy to the crystalline

lattice. When they drop from the metastable state to the lower laser level, they emit on the strong 1.06-micrometer laser transition and then quickly drop to the ground state.

8.4.2.1 Neodymium Hosts

The most common neodymium host is yttrium–aluminum garnet or YAG, a hard, brittle crystal with the chemical formula $Y_3Al_5O_{12}$. Dozens of other crystals have been tested as neodymium hosts; two others often used are yttrium lithium fluoride (YLF) and yttrium vanadate (YVO_4). Silicate and phosphate glasses are also important neodymium hosts.

Neodymium atoms replace some yttrium atoms in YAG. Typical YAG rods are cylinders 6 to 9 mm (0.24–0.35 inches) in diameter and up to 10 cm (4 inches) long, cut from single-crystal boules. The crystal has good thermal, optical, and mechanical properties, but its growth is slow and difficult, limiting sizes. Larger pieces of Nd–YAG can be made by fusing powdered Nd–YAG under pressure to make a transparent polycrystalline ceramic with thermal and mechanical properties like the single crystal. Generally, manufacturers do not specify if lasers are made from single-crystal or ceramic material.

Nd–YAG emits most strongly at 1064 nm. It emits about 20% as much light at 1319 nm and significant power at 946 nm, but those lines are rarely used. Packaging an Nd–YAG laser emitting on the three strongest lines with harmonic generators produces the shorter wavelengths listed in Table 8-2. Continuous beams can generate the second harmonic, but producing the third, fourth, and fifth harmonics generally requires pulses with high peak power. The 532-nm green line and the 473-nm blue line are the most important visible lines.

The optical and thermal properties of Nd–YAG allow both continuous and pulsed laser operation. The small size of Nd–YAG rods and the optical properties of neodymium atoms limit energy

Table 8-2. Important Nd–YAG harmonics

Fundamental	1064 nm	946 nm	1319 nm
Second	532 nm	473 nm	660 nm
Third	355 nm		
Fourth	266 nm		
Fifth	213 nm		

storage in a typical rod to about half a joule, far less than in a ruby rod. However, most of that energy can be removed from the rod in a Q-switched pulse, and the energy can be replenished in well under the millisecond duration of a flashlamp pulse, so a repetitively pulsed Nd–YAG rod can generate high average powers, as well as high peak power in Q-switched pulses lasting 10 to 20 nanoseconds.

Nd–YAG lasers can be quite powerful. The average power of repetitively pulsed lasers can exceed a kilowatt, although most operate at much lower powers. Power generally is lower in continuous lasers and is lower at short-wavelength harmonics than at the fundamental wavelength.

Yttrium lithium fluoride (YLiF₄) does not conduct heat as well as YAG and is softer, but its refractive index changes less with temperature, and it suffers fewer heat-related problems. As in YAG, the neodymium atoms replace some yttrium atoms in the YLF crystal. Nd–YLF can store more energy than Nd–YAG, so it can generate higher-energy Q-switched pulses. The crystal is birefringent, so unless it is restricted to a single polarization it generates two wavelengths, 1047 and 1053 nanometers, each with its own polarization orientation. Commercial versions emit pulses or continuous beams.

Neodymium can also be doped into YVO₄, in which it replaces some yttrium atoms. Two other rare earths, gadolinium and lutetium, may be used in place of yttrium. All three are similar in size and have similar properties. Nd-doped vanadates have the same primary line as Nd–YAG, 1064 nm, but the secondary lines are at 914 and 1342 nm, with the latter line much stronger than that from Nd–YAG. Nd–YVO₄ has higher gain efficiency than Nd–YAG, making it attractive for very small lasers, such as green laser pointers. Vanadate properties also differ from those of YAG lasers in other subtle ways, making them a better choice for certain applications.

8.4.2.2 Neodymium–Glass Lasers

Glass can also be doped with neodymium for solid-state lasers. The output wavelength depends on the glass composition; it is 1062 nm for silicate glass, 1054 nm for phosphate glass, and 1080 nm for fused silica. Its main attraction is the ease of making glass with good optical quality in sizes much larger than available for crystalline laser hosts. Nd–glass lasers have a larger output bandwidth than crystalline lasers, allowing mode-locking and generation of shorter optical pulses. Because it has lower gain than Nd–YAG, an equal

volume of Nd-glass can store more energy than YAG, so the glass laser can generate higher-energy pulses.

The principal trade-off is that glass has poorer thermal characteristics than YAG, so glass lasers need more time to cool between pulses. Thus, glass laser oscillators cannot be pumped continuously and normally operate at much lower repetition rates than Nd-YAG lasers.

However, the attractions of larger volumes and higher energy are compelling for use in giant research lasers, such as the National Ignition Facility at the Lawrence Livermore National Laboratory in California. The system is a giant oscillator amplifier with 192 beam lines focused onto a single target to produce nuclear fusion. It contains more than 3000 slabs of neodymium glass, each weighing 42 kilograms and measuring 81 by 46 by 3.4 centimeters. Such big chunks of glass need a long time to cool down, so they can fire only a few shots a day, but they produce tremendous peak powers during their brief pulses.

Neodymium-doped glass can also be used in fiber lasers, but most fiber lasers are doped with ytterbium or other rare earths, which have properties more suitable for use in fiber lasers than neodymium.

8.4.2.3 *Neodymium Laser Configurations*

Neodymium laser rods are pumped with lamps or semiconductor diode lasers, as shown in Figures 8-5 through 8-7. Typically, lamp-pumped lasers have wall-plug efficiency in the 0.1–1% range, but the efficiency of diode-pumped lasers often is well over 10%.

The high gain in neodymium lasers allows the use of stable or unstable resonators. A stable resonator can produce the standard Gaussian TEM₀₀ beam with a bright central spot, but it does not extract laser energy from as much of the laser volume as an unstable resonator. The beams from unstable resonators differ from those of stable resonators in the near field, but in the far field both have bright central spots.

External amplifiers can boost the output power and energy produced by neodymium oscillators, as described in Section 6.5. Different stages do not need to have the same host material, but the amplification band must overlap the oscillator output. Thus, a Nd-YAG oscillator may be used with Nd-glass amplifier stages.

Neodymium lasers are versatile and can take a variety of forms. Some are made specifically for medical applications, materials

working, laser marking, or used in instrumentation or measurement. Others are made for general-purpose research or laboratory use. As a result, neodymium lasers can look strikingly different, ranging from a compact green laser pointer you can fit in your pocket to a massive drilling laser installed on a factory laser, yet Nd–YAG rods lie at the core of both.

The wide range of applications leads to different design choices. General-purpose laboratory lasers give the user many options, such as leaving room for the addition of accessories such as harmonic generators and Q switches. Lasers designed for specific applications typically allow few modifications.

The smallest models can operate without active cooling, but larger types require either forced-air cooling (i.e., a fan) or flowing-water cooling to remove substantial amounts of waste heat. A 100-watt laser that is 1% efficient generates 9.9 kilowatts of waste heat!

8.4.2.4 *Harmonic Generation and Wavelengths*

The near-infrared wavelength of neodymium lasers is fine for many purposes, but visible or ultraviolet light is better for many others. Fortunately, neodymium lasers generate high enough powers so that nonlinear harmonic generation can readily produce shorter wavelengths, as listed in Table 8-2 for commercial Nd–YAG lasers. YLF and YVO₄ lasers can also produce harmonics at similar wavelengths.

Frequency doubling is the simplest form of harmonic generation and can be done with continuous beams. Figure 8-10 compares the arrangement used to generate the 532-nm second-harmonic to the more complex configurations used to produce the third and fourth harmonics at 355 and 266 nm in the ultraviolet. Because the conversion is inevitably incomplete, some light remains at the fundamental 1064-nm wavelength and must be blocked or separated from the other wavelengths as shown.

Some neodymium lasers are packaged with internal harmonic generators, so the packaged laser emits the harmonic wavelength rather than the fundamental, but descriptions may not clearly indicate the nature of the light source. Green laser pointers are a particular concern; they contain a neodymium laser and a harmonic generator that converts its 1064-nm output to 532 nm. Some inexpensive green laser pointers do not block either the 808-nm pump beam or the 1064-nm neodymium fundamental, posing potential eye hazards.

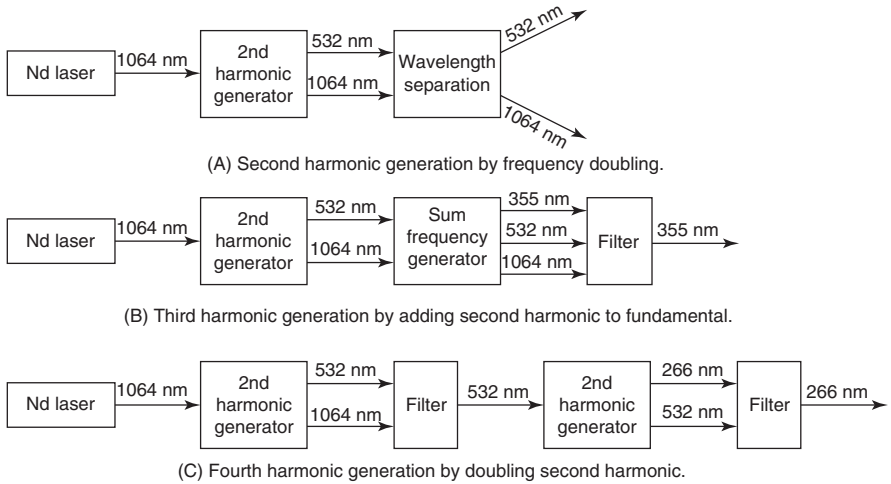


Figure 8-10. Generation of second, third, and fourth harmonics of a Nd-YAG laser.

Sum-frequency generation can also convert near-infrared light from neodymium lasers into the visible and near ultraviolet. For example, combining the 1064 and 1342 nm Nd-YAG wavelengths in a suitable nonlinear material generates sum-frequency output at 593 nm.

8.4.3 Ytterbium-Doped Lasers

Like many other rare-earth elements, ytterbium has strong laser lines in the near infrared. However, ytterbium lacks broad pump bands, which makes it impractical for lamp pumping. The development of diode pumping opened the door to practical ytterbium lasers emitting at 1030 to 1100 nm and pumped by either 915, 940 or 975 nm diode lasers.

Ytterbium can be doped into YAG and other materials to make both crystalline and ceramic bulk lasers. However, its primary use is in glass fiber lasers, described in Chapter 9. The rest of this section describes ytterbium as a dopant for use in bulk lasers.

8.4.3.1 Ytterbium Energy Levels and Efficiency

Ytterbium has a simple and favorable energy-level system in YAG, shown in Figure 8-11. Like many other rare-earth elements,




Figure 8-11. Energy levels and pump and emission lines for ytterbium-doped YAG. Note that the longest pump wavelength is within 10% of the shortest laser wavelength. That is a sign of very high efficiency.

ytterbium has strong laser lines in the near infrared. It also has an absorption band from 915 to 945 nm and a sharp peak at 975 nm, where diode lasers are particularly efficient. Look at the ytterbium lines showing pump energy and energy emitted on the laser transition, and you will note that they are noticeably closer than shown for neodymium in Figure 8-9.

The drawings are not exactly to scale, but the numbers tell the story. The *quantum defect*—the fraction of the pump photon energy that does not emerge in the laser photon emitted on the laser transition—is an impressively low 9% for ytterbium for a 940-nm pump and 1030-nm laser emission. In contrast, the 808-nm pumping the 1064-nm Nd–YAG transition loses 24% of its energy. That is only a theoretical lower limit on efficiency, but it is still important because diode-pumping of neodymium is considered efficient by traditional laser standards.

Inevitably, there are tradeoffs as the quantum defect shrinks toward zero. A 1030-nm Yb–YAG laser pumped at 975 nm has a 5.6% quantum defect. That puts the lower laser level close enough to the ground state to give it an appreciable population at thermal equilibrium and room temperature. Such lasers are called

quasi-three-level systems because they are intermediate between the three- and four-level lasers described earlier. That can lead to some unusual and unexpected behavior.

8.4.3.2 Ytterbium Wavelength Range

Another important feature of ytterbium lasers is their wavelength range. For simplicity, Figure 8-11 shows only a few key energy levels of ytterbium. Actually, many more closely spaced lines exist near the upper and lower laser levels that cannot be shown properly on that scale. That is because energy levels in rare earths can be split into a series of closely spaced levels by electric fields, called the Stark effect. This differs from the splitting of energy levels in other solid-state lasers into *vibronic states* caused by lattice vibrations called phonons that you will learn more about in Section 8.6.

Instead of having the narrow linewidth that we expect in lasers, transitions involving energy levels with many closely spaced states have a broad gain band that can emit light on a broad range of wavelengths, depending on operating conditions and the optics of the laser cavity. Figure 8-11 shows two examples, a 1030-nm transition dropping to a slightly lower state than a lower-energy 1050-nm stage. Actually, both the upper and lower laser levels are essentially continuous bands that allow ytterbium to emit at wavelengths from less than 1030 nm to around 1100 nm, depending where in those bands the light-emitting atoms start and end in those levels.

Ytterbium lasers typically are used more for fixed-wavelength applications in industry than for applications that require broad bandwidth. The main fixed wavelengths used are 1030 nm, where output is high, with other options including the slightly longer 1050 nm and matching the 1064 nm of neodymium laser.

Like neodymium, ytterbium lasers can be used to generate second harmonics, which typically are at 515 nm because of the higher fundamental power available at 1030 nm.

8.4.3.3 Ytterbium Laser Hosts

Ytterbium is like neodymium in its ability to operate well in many crystalline, ceramic, and glass hosts. YAG is on the top of the list for bulk ytterbium laser hosts because of its potential for high gain in thin disks as well as rods, with fixed wavelength typically 1030 nm, but sometimes 1050 nm or 1064 nm. The peak absorption in Yb-YAG is around 940 nm. Yttrium vanadate and a number

of other crystalline hosts also have been tested, but are not in widespread use.

Glasses are good hosts for ytterbium-fiber lasers, described in Section 9.4.1.

8.4.4 “Eye-Safe” Solid-State Lasers

Neodymium and ytterbium lasers have earned their shares of the solid-state laser market because of their excellent performance, but their wavelengths in the 900 to 1350 nm range have a serious drawback for some applications. These near-infrared wavelengths are not visible to the human eye, but the eye does transmit them to the light-sensitive retina, where a comparatively modest beam can cause permanent damage. That is a serious problem for applications such as laser radar and atmospheric measurements, which involve sending near-infrared beams into the open air where people could easily encounter them and suffer eye damage.

Fortunately, water strongly absorbs wavelengths longer than 1400 nm, and the ocular fluid is largely water, so it blocks those wavelengths. That means near-infrared wavelengths beyond 1400 nm pose minimal danger to the retina, although high powers could burn the cornea, the surface layer of the eye. That has led both military and civilian users to seek solid-state lasers in the retinal-safe range at wavelengths from 1400 to 2100 nm so that they do not endanger the eyes of distant bystanders, users or troops in training and can use sensors and instrumentation similar or identical to that used with neodymium or ytterbium lasers. Several types are available at fundamental wavelengths or at wavelengths produced by nonlinear techniques.

8.4.4.1 Erbium YAG and Glass Lasers

The rare-earth element erbium is well known for its emission band centered on 1.55 μm , which has become widely used in fiber lasers and amplifiers for telecommunications. That band can also be used for bulk lasers in glass and crystals at a broader range of wavelengths stretching to 1650 nm.

Figure 8-12 shows the key energy bands in erbium lasers, which are spread in a number of closely spaced states across a band of energies both in the ground state at bottom and at the two laser levels above it. At the left are the energy levels in the 1550-nm Er-glass laser transition, which occurs between the upper laser level in




Figure 8-12. Energy levels and pump and emission lines for two types of diode-pumped erbium lasers. At left are the energy levels for the 1550-nm Er band in glass, which can be pumped in two diode bands, at 980 and 1450 nm. At right are the transitions and energy levels involved in the 2940-nm Er–YAG laser in the mid-infrared, used for medical treatment.

the middle and the ground state at bottom. Two absorption lines are used for diode pumping, one at 980 nm and the other at 1450 nm. High-power diodes are readily available at 980 nm, which is close to the peak of diode-laser efficiency, and that wavelength is preferred for many purposes. The 1450-nm pump wavelength is close to the laser output wavelength, so the photon defect is much smaller than for pumping at 980 nm, but that benefit is offset by the lower wall-plug efficiency and output power of 1450-nm diode lasers. Erbium-doped glass fiber lasers and amplifiers are covered in Chapter 9, but Er–glass can also be used in bulk lasers. Erbium-doped YAG has similar transitions at 1619 and 1645 nm that can be diode pumped for use in bulk lasers.

A second erbium transition in YAG is shown at the right. This is also pumped by 980-nm diodes, which excite erbium atoms to a second excited state as shown in the figure. However, the laser transition in this system is between the second and first excited states, emitting at 2940 nm in the mid-infrared. Water in bones and tissue absorbs very strongly at that wavelength, so Er–YAG bulk lasers are used for dentistry and some types of surgery.

8.4.4.2 Thulium–YAG and Glass Lasers

Thulium is a rare-earth element with a 2010-nm laser line in YAG that can be pumped with 980-nm diode lasers. When doped into glass fibers, it can emit up to a kilowatt of laser light at 1900 to 2070 nm, but uses other pump sources, as covered in Section 9.4.3.

8.4.4.3 *Holmium YAG and YLF Lasers*

Holmium is a rare-earth element that emits near 2100 nm when doped into crystals of YAG or YLF. Pumping usually is a multi-stage process. For example, a diode-pumped erbium laser emitting at 1550 nm could pump a thulium fiber laser emitting at 1950 nm to pump holmium fiber lasers. Holmium output in glass is limited by rapid increases in glass absorption at wavelengths beyond 2.05 μm .

8.4.4.4 *Hybrid Laser-Nonlinear Sources*

Neodymium lasers or other solid-state lasers emitting near 1 μm can be combined with nonlinear sources to generate eye-safe wavelengths. For example, passing 1064-nm Nd-YAG pulses through a high-pressure cell containing pure methane can shift the wavelength to 1540 nm. Another option is neodymium pumping of optical parametric oscillators (described in Section 11.2.1) tuned to emit 1570 nm.

8.4.5 *Mid-Infrared Solid-State Lasers*

Solid-state lasers have long been scarce in the mid-infrared band, defined for our purposes here as a range of wavelengths from about 2.5 to 5.5 μm where the atmosphere is transparent enough for remote sensing and measurement. Standard oxide-based glasses do not transmit well in this region, but some crystalline hosts are clear enough for use in lasers. Semiconductor-based quantum cascade lasers have developed rapidly in recent years and are covered in Section 10.10.

8.4.5.1 *Erbium Lasers in 3- μm Band*

The rare-earth erbium has lines in the 3- μm band, with wavelength depending on the host material. In Er-YAG, the laser line is at 2940 nm; for the newer crystal Er-YSGG (yttrium scandium gallium garnet) emission is at 2790 nm. Both can be pumped by 980-nm diodes or by flashlamps. Oxide glasses absorb too strongly in the 3- μm band to be used in lasers, but fluorozirconate glass (ZBLAN, $\text{ZrF}_4\text{-BaF}_2\text{-LaF}_3\text{-AlF}_3\text{-NaF}$) has much lower loss, allowing erbium-doped ZBLAN fiber lasers pumped by 980-nm diodes that emit at 2840 nm.

A major attraction of lasers near 3 μm is the high absorption of water in that range, peaking at 2.94 μm . The absorption is so strong that such lasers can cut bone and drill teeth, as well as be used in other dermatology applications that benefit from high absorption.

8.4.5.2 Chromium-Doped Chalcogenide Lasers

Other hosts for mid-infrared solid-state lasers are based on compounds that are transparent at those wavelengths such as zinc sulfide and selenide. The two are part of a family of compounds called *chalcogenides* that form both glasses and crystals and contain sulfur, selenium, and/or tellurium. Their bandgaps are 3.54 electron volts for ZnS and 2.7 eV for ZnSe. The two are semiconducting, so electrons with more than the bandgap energy can travel through the solid. They are transparent to photons with less energy, because the photons lack enough energy to raise a valence-band electron to the conduction band. Mid-infrared photons have even less energy—about 0.25 electron volt for a 5- μm photon—so ZnS or ZnSe are transparent in that band, like glass is in the visible band. That makes ZnS and ZnSe viable hosts for mid-infrared solid-state lasers.

Chromium ions in ZnSe and ZnS have laser lines that can be pumped efficiently at 1.5 μm , typically with erbium-doped lasers, and have gain across a wide band from about 1.8 to 3.1 μm . Cr-ZnSe lasers are tunable across their gain band, and millimeter scale chips can generate up to 10 watts at room temperature. Their energy levels are similar to those of the tunable vibronic lasers described in Section 8.6, making Cr-ZnS/ZnSe lasers one of the few available solid-state mid-IR lasers.

8.4.5.3 Iron-Doped Sulfide/Selenide Lasers

Iron doped into ZnSe or ZnS has laser lines at longer wavelengths, from 3.4 to 5.1 μm when pumped with laser light near 3 μm . Like the chromium-doped ZnSe/ZnS lasers, the iron laser lines can be tuned in wavelength.

8.5 OPTICALLY PUMPED SEMICONDUCTOR LASERS

In the early days of lasers, it was possible to clearly classify them into distinct classes of devices. Decades of advances have blurred those lines to the point where semiconductors are being optically pumped as well as dielectric solids. Section 8.3.3 mentioned that semiconductor thin disks could be optically pumped. Now it is time to look at a class of unusual lasers which have become important solid-state sources that are variously called *OPSLs* or *vertical external-cavity surface-emitting lasers (VECSELs)*.

The terminology is muddled, apparently because competing groups developed very similar products that do essentially the same

thing: generate near infrared light by optically pumping thin semiconductor disks that can be frequency-doubled into the visible. For practical purposes, they act much like other solid-state lasers but produce visible wavelengths not available from other solid-state lasers. But what is inside the box is special type of semiconductor.

The devices going by different names are functionally similar. The optically pumped disk is a semiconductor with the optical structure of a semiconductor diode laser with a laser cavity that is perpendicular to the semiconductor chip and emits light from its surface. Its internal structure guides light, but does not guide current, simplifying the design in ways that make it possible for them to be diode pumped to generate light at near-infrared wavelengths that are hard to produce otherwise and that can be frequency-doubled to visible wavelengths.

Strictly speaking, these devices are OPSLs. They also have vertical external cavities and emit laser beams from their surfaces, as shown in Figure 8-13. Therefore, both terms are reasonably accurate, although there also are electrically powered devices called VECSELs that you will read about in Section 10.7.2, so not all VECSELs are optically pumped.

Whichever label you prefer, OPSLs or VECSELs were developed to meet a long-unmet need for compact solid-state lasers able to emit watt-class beams of good quality at visible wavelengths. Outside of frequency-doubled neodymium lasers in the green, the only good-quality visible beams came from gas lasers, and they were available at only a few wavelengths, were mostly limited to a few milliwatts of power, and except for milliwatt-level helium–neon lasers, typically unwieldy and expensive. Frequency doubling of other near-infrared lasers did not offer the desired variety of visible wavelengths. Nor was it easy to frequency double electrically powered semiconductor lasers because of poor beam quality for diode lasers emitting from their sides and low power for surface-emitting diode lasers. (See Chapter 10 for more detail.) Optically pumping the surface of a thin disk in an external optical cavity did the trick.

The basic idea of an OPSL is similar to the thin-disk laser that you saw in Figure 8-7. However, the details differ significantly, as shown in Figure 8-13. The OPSL itself is not just a thin slab of laser medium; it contains two multilayer semiconductor structures that serve different functions. The layer at the top of the chip, illuminated by the diode pump, is a stack of thin semiconductor layers containing structures called *quantum wells* (explained in




Figure 8-13. Inner workings of an OPSL. A pump diode at upper right is focused onto the OPSL chip, exciting atoms in the quantum wells at the surface. A distributed Bragg reflector at the bottom of the chip serves as the total reflector in the optical cavity. A dielectric coating on the inner surface of the output coupler reflects the fundamental wavelength and focuses it into a frequency-doubling crystal that produces visible (green) light, which the output coupler transmits as the output beam. (Courtesy of Coherent Inc.)

Section 10.5.2.2) which trap the population inversion created by the pump light. The bottom layer is a distributed Bragg reflector, a stack of layers that reflect nearly all of the stimulated reflection and serves as the total reflector at the back of the laser cavity. A heat sink sits beneath it, as in a thin-disk laser, but power levels are at the watt level, not the much higher levels of thin disks.

The fundamental output of the OPSL resonates in a folded optical cavity vertical to the surface of the chip. A dichroic coating on the output coupler reflects the fundamental frequency down and to the left, focusing it on a frequency-doubling crystal and a totally

reflecting mirror behind the doubler. Any light remaining at the fundamental wavelength is reflected by the coating into the other side of the folded cavity in which it oscillates.

The second harmonic from the doubling crystal is transmitted by the output coupler to form the visible beam. It does not resonate in the linear cavity at left because that cavity contains no laser gain medium, just a harmonic generator. The resulting visible beam looks and acts just like it came from a laser, although it is actually the second harmonic of a near-infrared laser.

The standard semiconductor used in OPSLs is indium gallium arsenide (InGaAs), a particularly efficient diode laser. Its fundamental output wavelength can be set between 920 and 1154 nm to give doubled output at wavelengths between 460 nm in the blue and 577 nm in the yellow. The OPSL can also be designed for frequency tripling, which generates wavelengths as short as 355 nm in the ultraviolet. Standard OPSL wavelengths include the 488-nm line of the argon laser, which OPSLs are replacing in life science applications. Other wavelengths are tailored to wavelengths optimum for applications, including 460 nm for writing on film, and 561 and 568 nm for life sciences. Other semiconductors with fundamental output at different wavelengths can also be used in OPSLs, usually at longer wavelengths.

InGaAs can be powered electrically, so you may wonder why OPSLs are pumped optically. The reasons are that achieving the required high powers requires a large emitting area, which must be uniformly excited to produce the high beam quality needed for OPSLs. Electrical current must come from the sides, so it does not spread as evenly across the surface as optical pumping, which can come from above.

The vertical external-cavity structure is another important feature of the OPSL. It allows optical pumping and also provides room for the intracavity frequency-doubling optics.

In recent years, OPSLs have been very successful in displacing gas lasers from existing applications such as flow cytometry and fluorescence excitation because they are much smaller and more efficient than visible gas lasers and do not require the high voltages and active cooling required for watt-class gas lasers. The ability to customize OPSLs to emit at a specific absorption line has opened the door to new applications.

Some manufacturers do not clearly identify the type of laser inside the watt-class solid-state visible lasers they sell, and their

sales staff may not know what is inside the box. Often you will find they are OPSLs.

8.6 BROADBAND AND TUNABLE SOLID-STATE LASERS

Lasers with broad gain bandwidth are important because they allow both tuning of the laser wavelength and generation of extremely short pulses, both of which are important for laser applications. There are two main classes of solid-state lasers with broad gain bandwidth. One is called *vibronic* lasers because their laser transitions involve simultaneous changes in both vibrational states of the atomic lattice and electronic states of the light-emitting species, spreading them over closely spaced energy levels. The other groups are rare-earth elements such as ytterbium where electric fields produce a similar effect.

8.6.1 Vibronic Lasers

In vibronic lasers, the lower laser level is a band spanning a range of energy levels arising from atomic vibrations in the solid-state laser host, as shown in Figure 8-14. This makes it possible to tune




Figure 8-14. General energy-level structure of a vibronic laser.

their wavelength across a range determined by the width of the vibrational energy band. In many cases, the upper laser level is itself the bottom of a vibronic band created by vibrational sublevels of a higher-energy electronic state. Note the contrast between the transitions shown in Figure 8-14 and the simpler transitions shown for ruby and neodymium in Figures 8-8 and 8-9.

The excited atom can drop to anywhere within the lower vibronic band after it is stimulated to emit energy, but gain varies across the band. Without wavelength-selective optics, laser output is at the wavelengths with highest gain. With suitable wavelength-tuning optics, vibronic laser emission can be tuned across a comparatively broad range of wavelengths, which can vary up to $\pm 20\%$ from a central wavelength. Table 8-3 lists emission ranges and compositions of some vibronic lasers. Not all are available commercially. The table includes three two nonvibronic laser types covered in earlier sections, ytterbium in Section 8.4.3 and thulium in Section 8.4.4.2, and one mid-infrared vibronic laser and chromium-doped ZnSe in Section 8.4.5.2.

The broad gain bandwidth of vibronic lasers makes them suitable for both broadband and tunable output. As you learned in Section 4.6.2, broadband output is needed for short pulses, which have become a major application of titanium-sapphire lasers.

Table 8-3. Representative broadband vibronic lasers, plus two other tunable lasers

Name	Composition	Wavelength range (nm)
Alexandrite	Chromium-doped BeAl_2O_4	701–858*
Co– MgF_2	Cobalt-doped MgF_2	1600–2400 ⁺
Cr–Emerald	Chromium-doped $\text{Be}_3\text{Al}_2(\text{SiO}_3)_6$	729–842
Cr–Forsterite	Chromium-doped Mg_2SiO_4	1167–1345
Cr–GSGG	Chromium-doped $\text{Gd}_3\text{Sc}_2\text{Ga}_3\text{O}_{12}$	740–850
Cr–YAG	Cr^{4+} -doped YAG	1350–1650
Cr–ZnSe	Cr^{2+} -doped ZnSe	2000–3100
Fe–ZnSe	Iron-doped ZnSe	3400–5200
LiCAF	Cr-doped LiCaAlF_6	720–840
LiSAF	Cr-doped LiSrAlF_6	780–1010
Thulium–YAG	Thulium-doped YAG (nonvibronic)	1870–2160
Ti–sapphire	Titanium-doped Al_2O_3	675–1100
Ytterbium	Yb-doped glass or YAG (nonvibronic)	1030–1100

*To 858 nm at elevated temperatures; only to 826 nm at room temperature.

⁺Operates at cryogenic temperatures.

Their broad gain bandwidth also gives vibronic lasers the ability to generate wavelengths not available from other solid-state lasers and tune the output wavelength for research and other applications.

Aside from their broad bandwidth and tunability, vibronic lasers work in much the same way as other solid-state lasers and sometimes can operate in the same laser cavities.

8.6.2 Titanium–Sapphire Lasers

Titanium-doped sapphire has two major attractions for solid-state lasers: broad emission bandwidth and good material characteristics. The broad bandwidth makes Ti–sapphire suitable for two distinct types of applications described in Section 6.4: generating optical pulses shorter than one picosecond (10^{-12} s) and producing wavelength-tunable output. The characteristics and the convenience of solid-state lasers have made Ti–sapphire the laser of choice for these applications, mostly for research.

The laser crystal contains about 0.1% titanium, which replaces aluminum in the Al_2O_3 crystal lattice. In that sense, it is similar to the ruby laser, in which chromium atoms replace aluminum in the sapphire lattice. The Ti^{+3} ion interacts strongly with the host crystal, and this combines with the structure of the titanium energy levels to make the bandwidth of the laser transition the broadest of any solid-state laser, from 660 nm in the red to 1100 nm in the near infrared. The peak output is at 700 to 900 nm.

No single Ti–sapphire laser emits over that entire range. Tunable emission usually is limited to a range of 100 to 300 nm with any particular set of cavity optics; changing cavity optics gives access to other parts of the wavelength range, as shown in Figure 8-15. Normally, tunable models are offered with multiple sets of optics with ranges that overlap at their edges or with broadband optics centered near the 800-nm gain peak. Harmonic generation is possible, with pulsed Ti–sapphire lasers offering the second harmonic tunable at 350–470 nm, the third harmonic tunable at 235–300 nm, and the fourth harmonic tunable near 210 nm.

Picosecond and femtosecond Ti–sapphire lasers are designed to oscillate broadband and are mode-locked to concentrate their power into a series of short pulses spaced at regular intervals. As you learned in Section 6.4.6, mode-locking produces short pulses by locking together all cavity modes, with the pulse width decreasing as the spectral bandwidth increases. The bandwidth of




Figure 8-15. Output of a tunable titanium-sapphire laser.

a Ti-sapphire laser is broad enough to allow generation of pulses in the 10-fs range. External optical devices can further extend the bandwidth to allow generation of pulses in the 5-fs range, but the arrangements are complex.

Ti-sapphire pulses can be amplified in external amplifiers, and the chirped-pulse amplification technique described in Section 6.5.4 permits peak powers to reach very high levels, albeit for a very short time.

Crystal growth is relatively easy, but Ti-sapphire has such high gain that most lasers use small crystals. The crystal's excellent optical and mechanical properties allow Ti-sapphire lasers to produce repetitive pulses or continuous beams.

One early practical limitation of Ti-sapphire lasers was its absorption peak at 500 nm, which required pumping with costly and inefficient argon-ion lasers because the upper level had too short a lifetime for flashlamp pumping. Doubled neodymium lasers emitting at 532 nm have now become standard, and pumping has been demonstrated with blue diode lasers emitting at 445 to 455 nm, but the shorter wavelength suffers from low absorption.

8.6.3 Alexandrite Lasers

The first vibronic laser developed commercially was alexandrite. It is now used in medicine and cosmetic treatment as well as in research.

The light-emitting species is chromium at concentrations of 0.01% to 0.4% in BeAl_2O_4 , a mineral known as alexandrite. Its energy levels are similar to those of the ruby laser, as shown in Figure 8-16. The 380 to 630 nm pump bands in alexandrite are similar to those in ruby, but alexandrite has a vibronic band in its ground state that is absent in ruby.

Alexandrite can lase on a weak 680.4-nm fixed-wavelength transition to the ground state, but it is not as efficient as the equivalent 694.3-nm ruby transition. Its strongest room-temperature emission is on vibronic transitions at 700 to 830 nm. The pump bands of alexandrite extend far enough into the red for pumping by red diode lasers as well as lamps. However, red diode lasers do not generate light as efficiently as the near-infrared 808-nm diodes used to pump neodymium lasers.

Alexandrite has peculiar kinetics because two electronically excited states together function as the upper laser level. One is the




Figure 8-16. Comparison of ruby and alexandrite energy levels.

bottom of a vibronic band of energy levels, the other is a fixed state with longer lifetime and only slightly less energy. This combination makes alexandrite gain increase with temperature, unlike most other lasers in which gain drops as temperature rises.

Temperature also affects wavelength. After stimulated emission, chromium atoms drop to one of the vibrationally excited substates of the ground state, and then release the remaining energy by exciting a vibration in the crystal lattice. Rising temperature increases the steady-state population of the lowest vibrational sublevels, making it harder to invert the population between these levels and the upper laser level. Because these low-lying levels correspond to the higher-energy end of alexandrite's tuning range, rising temperatures shift the gain toward longer wavelengths, so heated alexandrite lasers can operate at wavelengths up to 858 nm.

Alexandrite has lower gain than neodymium lasers, making cavity design more difficult, but allowing alexandrite to store more energy. Like Nd–YAG, alexandrite can operate in pulsed or continuous wave. In pulsed operation, average powers can reach 100 watts, lower than the most powerful Nd–YAG lasers, but still among the more powerful available lasers. Harmonic generation can produce tunable light at 360 to 400 nm, 240 to 270 nm, and 190 to 200 nm.

8.6.4 Other Vibronic Lasers

Many other vibronic lasers have been demonstrated in the laboratory, but so far have found little practical use. One promising candidate is chromium-doped lithium–aluminum–strontium fluoride ($\text{Cr}^{3+}:\text{LiSrAlF}_6$), usually abbreviated Cr–LiSAF. It is tunable between 780 and 1010 nm, a broad enough range to produce pulses lasting tens of femtoseconds. The bandwidth does not quite match that of Ti–sapphire, but Cr–LiSAF has a broad absorption band in the red, so it can be pumped by 670-nm diode lasers instead of the green lasers required for Ti–sapphire. When pulsed, it can generate harmonics in the blue and ultraviolet.

8.6.5 Broadband Nonvibronic Lasers

Ytterbium is technically not a vibronic laser because its broad tuning range comes from different effects. However, its gain band is about 10% of its wavelength; broad enough it can be mode-locked for short-pulse generation, as described in Section 9.4.1.1. The fiber

gain medium is particularly attractive because it is rugged, simple, and allows stable operation in difficult situations—as one laser trade-show exhibitor demonstrated by mounting a mode-locked Yb-fiber laser on a shaking table and using it to generate a femtosecond frequency comb.

8.7 PULSED SOLID-STATE LASERS

As you learned in Section 6.4, lasers can be made to produce pulses of various duration and power in a variety of ways. Such pulsing can have a tremendous impact on laser performance. From a user's standpoint, a pulsed laser can be very different than one emitting a continuous beam.

Compare, for example, two neodymium lasers, one emitting a continuous beam of one watt and another emitting 1000 Q-switched pulses per second. Suppose the Q switch was able to collect and store half of the one watt of light energy the laser produced during each one-millisecond interval, and that it then released the energy in a 10-nanosecond pulse. During that Q-switched burst, the laser would emit half a millijoule of energy or an average power of 50 kilowatts during the pulse. In other words, the Q switch would concentrate the pulse so that its peak power was 50,000 times higher than the continuous beam, and the two would have quite different effects on whatever they illuminated.

Most solid-state lasers have upper-state lifetimes lasting microseconds to milliseconds, adequate for Q switching. Nd-YAG is particularly good because its upper-state lifetime is over a millisecond. Q switching is possible in Ti-sapphire, but does not work as well because its upper-level lifetime is just three microseconds. Semiconductor diode lasers have upper-state lifetimes of only a couple of nanoseconds, which together with their small cavity sizes make Q switching impractical.

Ultrafast lasers emitting picosecond or femtosecond pulses have an even more concentrated effect. A mode-locked pulse lasting 100 femtoseconds is so short that if focused to a small spot it can vaporize the surface layer without damaging the layer just underneath. That is the secret of how lasers can cut glass without cracking it.

Such examples show how controlling laser operation can transform what lasers do. But it is not as simple as adding a Q switch here

and a mode-locker there to get what you want. Some pulsing techniques cannot be used with some types of lasers. Bulk neodymium lasers can be Q switched because their upper laser level has a long lifetime, allowing the laser to accumulate a large population in the upper laser level, ready for the Q switch to trigger release of all that energy. But neodymium's bandwidth in crystal hosts like YAG is too small to allow mode-locking to generate 100-femtosecond pulses. Therefore, no one laser can do everything.

The bottom line is that solid-state laser technology can be complex, but it also offers a tremendous versatility. You may want to look back at Section 6.4 to learn more about pulsing techniques now that you know more about solid-state lasers.

8.8 WHAT HAVE WE LEARNED?

- Solid-state lasers are transparent solids in which laser action is powered by light. They include glasses, crystals, and ceramics.
- Semiconductor diode lasers are not considered solid-state lasers because they are powered by electricity.
- Generally, the laser species in a solid-state laser is a dopant present in small amounts in an electrically insulating host.
- In a ruby laser, chromium is the laser species and aluminum oxide is the host material.
- YAG is a crystal called yttrium–aluminum garnet that is used as a host material for neodymium and other laser species.
- The active species is identified first and then the host. Nd–YAG means neodymium is the active species and YAG is the host.
- Important laser species in solid-state lasers include neodymium, ytterbium, chromium, erbium, and titanium.
- Solid-state lasers may be shaped as rods, slabs, thin disks, or optical fibers.
- Optical pumping uses light from a lamp or semiconductor diode laser to power laser action.
- The pump source must be matched to absorption lines of the laser species.
- Most pump lamps are pulsed flashlamps.
- Diode pumping efficiently converts poor-quality diode-laser beams into high-quality beams at different wavelengths.
- Semiconductor diode lasers are preferred for laser pumping because they generate light more efficiently than other lasers.

- Neodymium lasers are the most common bulk solid-state lasers. They emit at $1.06\ \mu\text{m}$ and can be pumped by lamps or by diodes emitting at $808\ \text{nm}$.
- Neodymium and other rare-earth elements emit laser light on electronic transitions near $1\ \mu\text{m}$.
- Crystals have better thermal properties than glass, so they can operate at higher powers and repetition rates, but are hard to grow in large sizes.
- Ceramics are polycrystalline materials that can be produced in larger sizes than single crystals.
- Glasses can be produced in much larger sizes than crystals, but dissipate heat poorly in bulk.
- Harmonic generation can convert near-infrared output from solid-state lasers into visible or ultraviolet light.
- Ytterbium requires diode pumping. It can be used in bulk glass or ceramic hosts, but is largely used in fiber lasers (See Chapter 9).
- Lasers at wavelengths longer than $1.4\ \mu\text{m}$ pose much less eye hazard than those at shorter wavelengths.
- Erbium-doped lasers pose less damage than ytterbium lasers.
- A few solid-state lasers emit in the mid-infrared at 2.5 to $5.5\ \mu\text{m}$.
- Thin disks of Yb–YAG pumped by arrays of diode lasers can emit kilowatts of laser power.
- OPSLs are thin disks of semiconductor optically pumped by diodes that produce watts of visible laser power. They are also called VECSELs for vertical external-cavity surface-emitting lasers.
- Vibronic lasers have broad gain bandwidth, so they can produce short pulses and their output wavelength can be tuned.
- Titanium–sapphire is not a vibronic laser, but has the broadest bandwidth among solid-state lasers because other factors broaden its gain bandwidth.
- Q switching can generate high-energy nanosecond pulses from solid-state lasers such as neodymium, but does not work well with all solid-state lasers.
- Mode-locking can generate ultrashort pulses lasting 100 femtoseconds from vibronic lasers.

WHAT'S NEXT?

In Chapter 9, we will learn about fiber lasers, a fast-moving branch of solid-state lasers that have become so important that they need a chapter of their own.

QUIZ FOR CHAPTER 8

- Which of the following is not considered a solid-state laser?
 - Ytterbium-doped fiber laser.
 - Neodymium-doped YAG rod laser.
 - Gallium–arsenide semiconductor diode laser.
 - Titanium–sapphire crystal laser.
 - Optically pumped semiconductor laser.
- Which of the following is the best pump source for a solid-state laser?
 - A flashlight
 - A flashlamp
 - A helium–neon laser
 - An electrical discharge
 - A fluorescent tube
- What type of laser is the most efficient pump for neodymium lasers?
 - 808-nm semiconductor diode.
 - Argon-ion laser at 488 nm.
 - Helium–neon at 632.8 nm
 - Ytterbium-doped fiber laser at 1030 nm
 - Titanium–sapphire laser.
- Which of the following elements is not used as an active species in solid-state lasers?
 - Neodymium
 - Titanium
 - Chlorine
 - Chromium
 - Ytterbium
- What is an important advantage of glass as a host for solid-state lasers?
 - Glass has better thermal characteristics than crystals.
 - Glass is the only material transparent to all pump light.
 - Glass is used in flashlamps.
 - Glass melts at higher temperatures.
 - Glass rods or slabs can be produced in large sizes
- Which of the following lasers is a quasi-three-level system?
 - Ruby
 - Neodymium
 - Titanium–sapphire
 - Ytterbium
 - There is no such thing.

7. How is a thin-disk laser pumped?
 - a. From above the surface by a diode laser
 - b. From the edge by a diode laser
 - c. From above by a flashlamp
 - d. By an electric discharge through a vacuum
 - e. By a low voltage applied across the disk
8. What is the second harmonic of the ruby laser?
 - a. 266 nm
 - b. 347 nm
 - c. 355 nm
 - d. 532 nm
 - e. 694 nm
9. Which of the following wavelengths cannot be readily generated from a Nd–YAG laser?
 - a. 266 nm
 - b. 355 nm
 - c. 477 nm
 - d. 532 nm
 - e. 1064 nm
10. Which solid-state laser has the broadest gain bandwidth?
 - a. Titanium
 - b. Ytterbium
 - c. Neodymium
 - d. Ruby
 - e. Thulium
11. What differentiates vibronic lasers from other solid-state lasers?
 - a. Can be pumped efficiently with semiconductor lasers
 - b. Higher gain
 - c. Shorter wavelengths
 - d. Broader gain bandwidth and tunable output
 - e. All of the above
12. Which of the following solid-state lasers poses the least threat of eye damage?
 - a. Neodymium–YAG
 - b. Ruby
 - c. Ytterbium
 - d. Titanium–sapphire
 - e. Erbium
13. You visit a trade show and see a company offering several small lasers, each emitting a separate color, including blue, green, and

yellow. All the salesperson will tell you is that they are “solid state.” What are they?

- a. Nd–YAG lasers emitting different harmonics
 - b. Titanium–sapphire lasers made to emit specific wavelengths
 - c. Ytterbium lasers
 - d. OPSLs or VECSELs
 - e. Diode lasers
14. Which solid-state laser can generate femtosecond pulses?
- a. Nd–YAG
 - b. Titanium–sapphire
 - c. OPSL
 - d. Holmium
 - e. Ruby

CHAPTER 9

FIBER LASERS AND AMPLIFIERS

ABOUT THIS CHAPTER

In this chapter, you will learn about fiber lasers, which technically are a type of solid-state laser, but differ in important ways from the other solid-state lasers covered in Chapter 8. The crucial difference is that the light-emitting atoms are embedded in the light-guiding core of an optical fiber. The combination of diode pumping with light guiding in a fiber core makes fiber lasers particularly efficient and allows them to generate continuous powers to tens of kilowatts. Doped optical fibers are also used to amplify light in the long-distance fiber-optic cables that are the backbone of the Internet.

9.1 WHAT ARE FIBER LASERS?

Fiber lasers are solid-state lasers in which the laser medium is in the form of an optical fiber. As you learned in Section 2.5, optical fibers transmit light through a thin central core surrounded by an outer layer with lower refractive index. Total internal reflection at the boundary between core and cladding traps light in higher-index core so it travels through the length of the fiber.

In 1961, Elias Snitzer decided to try making a laser by adding light-emitting elements to glass. He was working on fiber optics

for the American Optical Company and decided to put the light-emitting elements in a long thin rod of high-index glass surrounded by a lower-index glass. He put mirrors on the ends to make the first glass laser, and although it was a straight rod, it had the core-cladding structure of an optical fiber, so it was also the first fiber laser. He demonstrated laser oscillation in neodymium-doped glass rods 30 to 300 μm in diameter and several inches long. Later Snitzer also made an optical fiber amplifier in the same way. A few other people experimented with fiber lasers in the following years, but key technology needed to make fiber lasers was not yet available. To understand why, let's look at the basics of fiber lasers.

9.1.1 Fiber Laser Basics

You learned the basics of fiber optics in Section 2.5. The basic idea of a fiber laser is to put the light-emitting element in the core so total internal reflection can guide the emitted laser light along the length of the fiber, as shown in simplified form in Figure 9-1.

The lightly shaded area is the fiber core, containing the active laser species. Spontaneous emission goes in all directions. Light emitted toward the side escapes from the fiber, as shown at upper left. Other spontaneous emission directed along the length of the fiber is guided by total internal reflection along the fiber and between the cavity mirrors until it exits through the output mirror.




Figure 9-1. Basic concept of fiber laser. The shaded core of the fiber is doped with an optically pumped light-emitting atom such as ytterbium. Light emitted by the ytterbium atoms is confined along the length of the fiber core by total internal reflection at the edge of the core. Cavity mirrors reflect the light back and forth so it oscillates within the cavity, with some light exiting through the output mirror. This simple view does not show how the laser is pumped and does not show amplification. A real fiber laser is longer and thinner.

That initial spontaneous emission stimulates more emission, building up intracavity power concentrated in the fiber core. Total internal reflection can guide light around bends, and most glass fibers are thin enough to be flexible, so fiber lasers can operate even when they are coiled or bent.

The entire resonant cavity of a fiber laser is contained within the fiber, generally including the cavity mirrors, which usually are a series of layers formed within special segments of fiber. Pump light typically is delivered through optical fibers, and often the light generated within the fiber is coupled into a beam-delivery fiber connected directly to the fiber laser output. Integration of the whole optical path within a fiber laser greatly simplifies operation and maintenance, and avoids intracavity contamination, big advantages for industrial applications that have contributed to the success of fiber lasers.

To show the basic idea of fiber lasers, Figure 9-1 omits the crucial matter of transferring energy into the laser medium. Like other solid-state lasers, fiber lasers were optically pumped, and in the 1960s the only practical sources of the intense light needed were lamps. However, optical pumping from the side with a lamp, as shown for laser rods in Figure 8-5, became much harder when trying to focus the light onto a much thinner fiber and limited both laser efficiency and the potential output power.

Instead, interest in glass lasers turned to making large rods to store a large amount of energy that would emerge in a single pulse. However, glass conducts heat more poorly than crystals, so big glass rods needed seconds, minutes, and eventually hours to cool down, limiting their use.

9.1.2 Pumping Fiber Lasers

Another possibility was coupling pump light into the end of an optical fiber, as is done for sending signals in fiber-optic communications. The best approach to end-pumping would be with the small beam produced by a tiny semiconductor diode laser. The diode laser was invented in 1962, and diode pumping was tried in the 1960s. But early diode lasers emitted too little light and were too short-lived to be useful.

That changed with the development of longer-lived and higher-powered diode lasers during the 1970s and 1980s. Diode lasers emitted light from small areas on semiconductor chips, and the

emitting area could be mounted directly onto the end of a fiber, so the narrow emitting stripe on the diode put light right into the small light-guiding core of the fiber.

By good fortune, the best-developed type of semiconductor diode lasers, based on gallium-arsenide compounds, emit at near-infrared wavelengths from 800 to 1000 nm, where rare earths such as neodymium have strong absorption lines. So diode lasers could produce the right wavelengths needed for optically pumping solid-state fiber lasers.

More developments followed, including better designs for diode-pumping fibers and development of monolithic arrays of diode lasers emitting watts of output. But the real catalyst that triggered the emergence of fiber lasers was expansion of fiber-optic telecommunication networks that needed a way to amplify optical signals without turning them into electronic form.

9.1.3 The Emergence of Fiber Amplifiers

Fiber-optic telecommunication cables spread across continents in the 1980s and began crossing under the seas in 1988. Although optical fibers were remarkably transparent, they could not carry signals over unlimited distances. The light faded slowly, and after passing through 50 to 100 km of glass, the signal had to be amplified or it would become too faint to detect reliably. In early fiber-optic systems, the only way to amplify a faint incoming optical signal was to use a detector to convert the signal into electronic form, amplify the resulting electronic signal, and then convert the amplified electrical signal into optical form. It was cumbersome and expensive, so developers set out to develop all-optical amplifiers.

In theory an optical signal amplifier is just a length of laser medium through which a signal passes to be amplified, with no mirrors on the ends. Diode lasers were used to generate optical signals for communications, so developers tried to make diode-laser amplifiers, but they were too noisy. However, diode-pumped fiber lasers proved to be outstanding optical amplifiers. David Payne at the University of Southampton found that fibers doped with the rare-earth erbium could amplify light in the 1550-nm band where glass optical fibers have the lowest loss, a perfect match for telecommunication systems. Better yet, fiber amplifiers have tremendous bandwidth, so they could carry signals on many separate wavelengths, vastly increasing the capacity of

fiber-optic networks. The technology worked so well that it fueled an economic boom that became a bubble at the turn of the century.

Fiber lasers were a spinoff of the development of fiber amplifiers. In the 2000s, power levels increased dramatically and continuous-output industrial fiber lasers reached multi-kilowatt powers. Since then, fiber lasers have become a billion-dollar business, with uses ranging from heavy-duty laser machining to precision laboratory instruments.

9.1.4 Attractions of Fiber Lasers

A major point of this brief history is to point out how very successful fiber lasers and amplifiers have been. Without fiber amplifiers, the Internet would not be able to transport the huge amount of digital traffic it carries today. Fiber lasers are revolutionizing industrial machining, enabling new technologies such as additive machining, which builds up metal structures layer to layer using lasers, as you will learn in Chapter 13.

The success of fiber lasers comes from their advantages over older types of lasers:

- *Much higher wall-plug efficiency:* Early gas and solid-state lasers converted only a small fraction of input power into light, making energy costs high and cooling a necessity. The most efficient commercial fiber lasers can convert nearly half of the input electricity into light, slashing operating costs and making higher laser powers practical.
- *High surface-to-volume ratio of fiber eases heat removal:* Accumulation of heat inside bulk optics such as laser rods can degrade laser beam quality and performance. Fibers are so thin that heat is always close to the surface, limiting the need for active cooling.
- *Fiber-coupled beam path and beam delivery:* No optical surfaces are exposed to dirt, dust and the air from the pump diode to the end of the fiber delivering the laser beam, avoiding surface damage and greatly reducing maintenance needs. Fiber beam delivery is compatible with robotic positioning systems.
- *High beam quality at high power in a compact laser:* Fibers deliver a tightly focused, high-quality beam, and fiber lasers are compact because of their size and efficiency. Figure 9-2 shows a multi-kilowatt fiber laser.



Figure 9-2. A commercial multimode ytterbium-doped fiber laser that generates a steady 20-kilowatt beam for welding metals weighs in at about 600 kilograms and is comparable in size to a family refrigerator. The fiber architecture makes it remarkably compact for such a powerful laser. Photo by Alexei Markevitch, Courtesy of IPG Photonics.

- *Versatile and rugged technology:* Models range from multi-kilowatt continuous-output laser machine tools for cutting and welding to femtosecond and narrow-line output for laboratory use.

To understand how this is done, let's look inside fiber lasers.

9.2 OPTICAL FIBER STRUCTURES

Simple optical fibers have a circular core of glass with a high refractive index surrounded by a circular cladding with lower refractive

index. Total internal reflection at the core-cladding index confines light traveling through the fiber in a high-index core surrounded by a lower-index cladding, which is needed to keep light from leaking out of the fiber at points where the fiber touches something.

You learned about the major types of optical fibers used in optical applications, including passive light transmission and imaging. The fibers used in fiber lasers and amplifiers are called *active* because they generate light, and they differ from passive fibers in important ways. This section also describes some advanced fibers that may be used in both active and passive applications.

9.2.1 Single-Clad Doped Fibers

You saw in Figure 9-1 the simplest type of doped fiber, with a single central core doped with a light-emitting element such as erbium or ytterbium. The refractive index is slightly higher than that of the single cladding—typically 1% and often less—but that small difference is enough to confine light directed within about seven degrees of the fiber axis. Generally the core carries light in only a single mode for the best possible beam quality, but multimode cores may be used at high power.

This type of core is used in low-power fiber lasers and in optical amplifiers, which amplify an input signal making a single pass from one end of the amplifier to the other. Pump light typically enters through one end of the fiber.

9.2.2 Dual-Clad/Dual-Core Doped Fibers

A more efficient pumping scheme is the dual-clad structure shown in Figure 9-3. The small inner core has the highest refractive index and is doped with a light-emitting element such as ytterbium. It usually is single mode at the laser wavelength. All the ytterbium laser light is generated in the inner core and confined there by total internal reflection at the boundary with the lower-refractive-index inner cladding. This inner cladding is sometimes called the outer core because it is the middle of three layers and functions as a cladding to the inner layer, but is a core for guiding the pump light.

The inner cladding (outer core) has a refractive index higher than that of the outer cladding. This lets the outer core guide the pump light along the fiber. Total internal reflection at the boundary between the outer core and the cladding bounces the pump light




Figure 9-3. Internal structure of a dual-core fiber. The core sometimes is called the inner core. The middle layer is called either outer core or inner cladding. The outer layer shown may be called outer cladding or sometimes just cladding.

back and forth through the outer core as it travels along the fiber. Pump light following that path through the inner core also passes repeatedly through the inner core, where it excites the laser species. The goal is to transfer as much pump energy as possible to the laser species, which emits at a longer wavelength.

The relative size of the inner and outer cores depends on the type of laser. The outer core should be large to collect and guide pump power; in some cases, it approaches one millimeter in diameter. The inner core should be small to concentrate emission and produce a high-quality beam, but not so small that the light intensity becomes high enough to cause undesired nonlinear effects or optical damage. Keeping the inner core smaller than $10\ \mu\text{m}$ can restrict it to supporting only a single transverse mode, giving better beam quality. However, optical tricks such as coiling the fiber can confine light in a single mode even if the inner core is as large as $40\ \mu\text{m}$.

For simplicity, Figure 9-3 shows the outer core as cylindrical, but such strong symmetry is not desirable because it could direct pump light on paths that never enter the inner core to excite the laser species. Instead, the outer core normally has an asymmetric shape that reflects the pump light at varying angles so it passes repeatedly through the inner core as it travels along the fiber. Outer

cores instead have profiles that are D-shaped, rectangular, hexagonal, and/or off-center to assure asymmetry.

9.2.3 Single- and Multimode Fiber Cores

Fibers with both single- and multimode transverse modes may be used in fiber lasers. Modes transmitted depend on core size, wavelength, and refractive-index difference between core and cladding (or outer core). As for other lasers, beam quality is higher when a fiber laser operates in a single transverse mode.

With careful design, continuous wave single-mode fiber lasers can deliver up to 10 kilowatts in a high-quality beam. Some increase may be possible, but care is needed to avoid increasing light intensity to a level that could cause nonlinear effects that degrade beam quality.

Multimode operation is an option for many applications, particularly in laser materials working. Industrial fiber lasers have reached powers to 100 kilowatts by combining beams from several fiber sources into large-core multimode fibers. Another option is combining beams from many separate lasers to produce a single high-power beam of good quality, described in Section 9.3.7.

9.2.4 Microstructured Fibers

Conventional optical fibers are made of solid glasses with different refractive indexes for the core and cladding to control how they guide light. An alternative is using *photonic crystal fibers* with internal microstructures that control how they guide light. Such fibers can be used in fiber lasers as well as in fibers for light transmission.

The internal microstructures are produced by stacking glass rods and tubes together in the desired arrangement, then heating them until the glass softens so it can be stretched into a long thin fiber. Such microstructured fibers can be made with light-guiding capabilities impossible with solid-glass fibers, such as single-mode transmission in an exceptionally large or small single mode.

Fiber lasers can be made by doping the central regions of internal microstructures with ytterbium or another light-emitting element. In this case, a solid central region may be surrounded by a cladding-like region with regularly spaced holes; the air in the holes

effectively reduces the index of the glass to confine light in the core by total internal reflection. The solid central region retains a higher refractive index because it lacks holes.

Microstructured fibers can also be made of arrays of internal holes with only thin walls between them, often with larger holes in the center of the fiber. Such fibers guide light in a different way that can be attractive for beam transmission or for use as nonlinear elements. Hollow fibers can also be filled with gas for use as a nonlinear component or as a compact laser. The details are complex and beyond the scope of this book, but can be found elsewhere by searching for *photonic crystal fiber* and *photonic bandgap fibers*.

9.3 FIBER LASER DESIGN AND EFFICIENCY

The basic concept of a fiber laser is simple. A transmission fiber collects light from a pump laser (usually a diode) and couples it into one end of a rare-earth doped fiber through a rear cavity mirror that transmits the pump wavelength but reflects the wavelength generated by the fiber, as shown in Figure 9-4. The pump light excites rare-earth atoms in the doped fiber core to the upper laser level, and spontaneous emission triggers a cascade of stimulated emission on the rare-earth transition. The laser light travels back and forth along the doped fiber, where the rear cavity mirror reflects all the light,




Figure 9-4. Pumping a simple fiber laser. Pump diode at the right excites the light emitter in the doped fiber, which produces stimulated emission and oscillates in the laser cavity.

and the output mirror reflects some light and transmits the rest of the stimulated emission as a laser beam. The output beam usually is coupled directly into a beam-delivery fiber.

Details vary considerably among types of fiber lasers. This section focuses on the key concepts, without going into the specifics of the materials involved, and winds up with a summary of the attractions of fiber lasers. Section 9.4 covers the light-emitting materials used in fiber lasers.

9.3.1 Fiber Choice and Configurations

The selection of doped fibers is crucial in fiber laser design. In most cases, the doped central core has a small diameter for single-mode operation. At lower powers, fiber lasers may have a single central core doped with a light-emitting element that also carries pump light, as shown in Figure 9-1. However, fiber laser operation is more efficient with the light emitter in a single-mode central core surrounded by an inner cladding (sometimes called an outer core) that carries the pump light. Total internal reflection at the boundary of the inner and outer claddings guides the pump light, as shown in Figure 9-3. Section 9.2.2 described the advantages of the latter design.

Either type of fiber may be coiled as shown in Figure 9-4, but other arrangements are possible. The length of the gain fiber may range from centimeters to the 100-meter range, depending on performance requirements and design details. Two key variables are the concentration of the laser species in the fiber core and the length of the gain fiber, which together control the gain in the resonant cavity. Good design requires balancing fiber concentration and length to get the desired pumping level; too much of either can lead to problems.

One problem is gain saturation, shown in Figure 9-5. Power increases rapidly above laser threshold, then levels off to a uniform rate along the dashed line. At higher powers the rate of increase decreases, as the curve rolls over to the right. You can think of that as the laser running out of pump light or not being able to increase stimulated emission as input power increases, reducing laser efficiency. This can occur if the gain fiber is too long or contains too much of the light-emitting element.

The nonlinear effects described in Section 5.6 also can impact output by draining away some laser energy by processes that




Figure 9-5. Gain saturation occurs as pump power increases. Power rises quickly just above laser threshold, then levels out in the linear region along the dashed line. Gain saturates when the output does not increase in proportion to the pump power, and the curve starts declining at right.

convert the light to other wavelengths when the power density reaches high levels.

9.3.2 Cavity Optics in Fiber Lasers

An important feature of most fiber lasers is that the beam path is entirely through optical fibers from the exit face of the pump diode through the gain fiber and a beam-delivery fiber to the beam's destination, as shown in Figure 9-6. This is important because it avoids light losses and contamination issues that can occur at glass-to-air and air-to-glass interfaces.

Any time light passes between two media with different refractive indexes, such as glass and air, some light is reflected and lost. The effect, called *Fresnel reflection* is most obvious when you look from a lighted room through a window into the dark. The glass window reflects light from the room like a mirror. Silica glass with refractive index around 1.5 reflects only about 4% of the incident light, but that is enough to look bright at night. Optics can be coated




Figure 9-6. Fiber lasers typically have the beam path completely in fiber from the pump diode at left through the laser and beam-delivery fiber to the workpiece at right. The cavity and output mirrors, shown here as vertical lines, are optical elements called fiber Bragg gratings embedded in fibers.

to reduce that reflection, but fiber lasers can avoid those losses altogether if the light beam travels entirely within glass fibers with their ends pressed tightly together with no intervening air.

An all-fiber beam path also avoids dirt accumulation on optical surfaces. At low power, dust and dirt causes light scattering and losses. At high power, it can damage optical surfaces. Dust and dirt absorb light, heating it to a point where it may burn and darken the surface further so it absorbs more light, ultimately causing damage. All-fiber beam paths avoid exposed surfaces other than the one at the end of a beam-delivery fiber, which can be periodically replaced. The result is reduced maintenance costs and better performance, particularly for lasers in factory environments.

If you look carefully at the all-fiber beam path in Figure 9-6, you may think we forgot something very important—the mirrors at the ends of the gain medium, marked by labels and little arrows. You do not see the mirrors because they are ingenious optical structures called *fiber Bragg gratings* that are embedded in short lengths of special fibers included in the fiber laser. These special fibers contain a series of zones with high and low refractive index formed by exposing the fiber to ultraviolet light. A series of these zones of proper spacing reflects certain wavelengths and transmits others, such as the multilayer interference coatings described in Section 5.3.3.

This makes it possible to make an all-fiber optical cavity for a fiber laser. A fiber Bragg grating on one end transmits pump light but reflects the laser light generated inside the gain fiber. A fiber

Bragg grating on the other end reflects all the pump light back into the laser cavity, but serves as an output mirror for the laser light generated in the fiber.

9.3.3 Diode-Pumping Fiber Lasers

So far, we have shown you only one approach to pumping fiber lasers—illuminating one end of the gain medium with a single pump diode through a single fiber. In reality, there are several ways to pump fiber lasers, and the pump technique is crucial to performance, particularly at high powers.

End-pumping with a single pump diode as in Figure 9-4 is the simplest approach for low-power fiber lasers. For higher powers, separate fibers can collect light from many pump lasers and couple them together into the end of a fiber laser. Light from the linear arrays and stacks of diode lasers described in Section 10.5.3.3 also can be used for end-pumping.

An alternative approach that can deliver even more pump diode power into a fiber laser is called *side pumping*, but it is quite different from the perpendicular side pumping of laser rods shown in Figure 8-6. Instead, this side pumping is delivered by a fiber spliced to the fiber laser at a steep angle, as shown in Figure 9-7. The fiber coming in from above left carries pump light guided in a large cladding. At center, the pump light is coupled into the inner cladding carrying pump light and surrounding the doped core (gray) of the gain fiber. The coupling process requires making




Figure 9-7. Side coupling of pump light (above left) into a dual-cladding gain fiber (horizontal). Dashed light rays from the pump laser are said to enter the cladding (white) of the gain fiber from the side, but the input actually is at a steep angle. The pump light then travels through the inner cladding of the fiber laser, pumping the light-emitting species in the core of the dual-cladding fiber laser.

careful physical connections so the pump light can easily enter the gain fiber, but it can provide high power and efficient coupling for high-power lasers.

This type of side pumping can be highly efficient in transferring pump light into the doped fiber, giving high optical-to-optical efficiency that is essential for high-power fiber lasers. Many pump fibers can be coupled in this way.

9.3.4 Pumping Fiber Lasers with Other Fiber Lasers

Fiber lasers can also pump other fiber lasers. That might seem an odd procedure, but it has distinct advantages in generating wavelengths that are hard to produce otherwise or squeezing extra power out of fiber lasers. As with diode pumping, the pump light generally enters the end of the doped fiber.

Although diode pumping has brought important advances to development of fiber and solid-state lasers, particularly at high power, it does have its limits, particularly in the limited range wavelengths where high-power pump diodes are available. As you will learn, diode powers and efficiencies peak between 800 and 1050 nm. Diode lasers can pump in some other bands in the visible and near infrared, but the power and efficiency are lower.

The best way to pump fiber lasers with other absorption bands can be using a two-stage process—first diode pumping a fiber laser emitting the desired absorption band, then using that fiber laser to pump the second laser. For example, a diode-pumped erbium-fiber lasers emitting near 1600 nm can pump a thulium-fiber laser emitting at 1950 nm, which, in turn, can pump a holmium-fiber laser at 2100 nm, as described in Section 9.4.4. This has been done in other systems with pump bands at 1500 nm or longer.

Another reason for fiber-laser pumping is to generate high-quality beams at high power. Although fiber lasers dissipate waste heat well, like all other lasers they have their limits that depend on the amount of heat that the pump laser deposits in the fiber. That heat depends on the *quantum defect*, the difference between the energy of the pump photon E_{pump} (a quantum of energy) and the energy of the output photon emitted by the laser E_{laser} . Mathematically, the quantum defect D_q is as follows:

$$D_q = 1 - \frac{E_{\text{laser}}}{E_{\text{pump}}} \quad (9-1)$$

It is the amount of energy left behind as heat by each pump photon that excites an atom that emits a photon of laser light. The larger the quantum defect, the more heat left over for the fiber laser to dissipate. Fiber lasers are very efficient by laser standards, but when you are trying to produce 10 kilowatts of laser light, every little bit of energy saved helps avoid overheating.

Equation 9-1 is useful for showing that the quantum defect measures energy, but we usually measure the wavelength of light, which is inversely proportional to the wavelength, so we can simplify calculations by writing the equation in terms of the wavelength λ to give

$$D_q = 1 - \frac{\lambda_{\text{pump}}}{\lambda_{\text{laser}}} \quad (9-2)$$

Equation 9-2 is easier to use because the math is easier and you are usually given wavelength rather than energy.

9.3.5 Fiber Laser Efficiency

Commercial fiber lasers offer wall-plug efficiency that can reach 50% at high powers. That impressive efficiency combined with excellent beam quality of fiber lasers have been central to their success, particularly in the industrial laser market.

Efficiency is an important laser parameter, and as you learned in Section 4.5 laser efficiency depends on many factors and tends to be low. Semiconductor diode lasers can reach wall-plug efficiency close to 70%, but their beam quality is poor. The high *optical-to-optical conversion efficiency* of fiber lasers lets them convert that poor quality diode beam into an excellent quality fiber beam, with only a modest loss in efficiency. The high conversion efficiency benefits from the tight confinement of pump light in the fiber cladding and laser light in the core.

Table 9-1 shows the impact of efficiency on laser operation by calculating how much power is needed to generate a 10-kilowatt beam at various efficiencies, and how much power becomes waste heat. At low efficiency, staggering amounts of input power wind up as waste heat, which must be dissipated to avoid damaging the laser. Increasing efficiency from 1% to 5% pays a huge dividend, reducing waste heat by a factor of five. Waste heat generated also can compound input power requirements because refrigeration

Table 9-1. Impact of increasing energy efficiency on operating requirements for a 10-kilowatt laser, in terms of input power needed and waste heat generated.

Laser output (kW)	Efficiency (%)	Input power (kW)	Waste heat (kW)
10	1	1000	990
10	5	200	190
10	10	100	90
10	20	50	40
10	30	33	23
10	40	25	15
10	50	20	10
10	60	16.7	6.7
10	70	14.3	4.3
10	80	12.5	2.5
10	90	11.1	1.1

equipment needs power to operate it. Some lasers can be cooled with flowing water, but that has its own economic and environmental costs. Cooling equipment and high-power electric supplies also require valuable space in factories or laboratories.

The higher the power requirement for any laser, the more is the allure of high efficiency. The high overall efficiency of fiber lasers has made them particularly successful in industrial applications such as cutting and welding metals. If you are using kilowatts of power, increasing efficiency is more than going green; it is a way to boost profits.

9.3.6 Limits on Fiber Laser Power

The concentration of power that contributes to efficient pumping of fiber lasers also can cause troublesome nonlinear effects that can limit laser power and performance.

Glass has inherently low nonlinearity, but nonlinear effects are cumulative, so they build up with distance. They also are proportional to the light intensity inside the glass, so the smaller the fiber core, the stronger the nonlinear effects. The nonlinear effect with the lowest threshold in fiber lasers and amplifiers is called *stimulated Brillouin scattering*. It arises when photons scatter off atoms in the glass, reducing their frequency by 10 to 20 gigahertz by transferring a small fraction of their energy to vibrations in the glass (phonons). The photons also bounce backward, draining energy from the laser beam. The effect can limit fiber laser output at multi-kilowatt levels in single-mode fibers.

High powers also can damage fibers. Ends are the parts of fibers most vulnerable to optical damage because they can collect dirt. At high powers, dirt on the end can absorb so much light energy that the resulting heat damages the surface. Minor damage can increase absorption, causing a cascade of further damage that may “blow off” the end.

These limiting effects depend on the peak power, not on the average power. Thus they are strongest for short laser pulses, which reach the highest peak powers. Continuous lasers are less vulnerable because they do not reach such high power levels.

9.3.7 Large-Core Fibers and Beam Combination

Using larger fibers allows operation at higher laser powers by spreading the light through a larger area in the fiber, reducing the intensity below the level that could cause nonlinear problems. Microstructured fibers can allow single-mode operation in larger cores. Fiber lasers also could be operated multimode in larger cores, which reduces beam quality but is adequate for much laser materials-working where the beams are directed to nearby targets.

Another approach to very high powers is to combine the output of many smaller fiber lasers into a larger-core multimode fiber, where spreading it through a larger area reduces the beam intensity to avoid nonlinearities and damage. In this way IPG Photonics built a 100-kilowatt continuous laser for research by a Japanese university on cutting and welding thick metal sheets for shipbuilding.

The development of laser weapons emitting several tens of kilowatts is more demanding because they require much higher beam quality to destroy targets at distances to a few kilometers away. Lockheed Martin has demonstrated an approach called *spectral beam combination*, in which narrow-line beams from about a hundred fiber lasers, each emitting 300 watts at a different wavelength, are combined to generate a single high-quality 30-kilowatt beam. This approach also has been scaled to higher powers, but details are classified.

9.4 RARE-EARTH-DOPED FIBER LASERS

So far we have described fiber lasers in general terms. Now it is time to look at the laser species used in fiber lasers—specifically, at the

rare earths that are the active species in most fiber lasers and amplifiers. The most important rare earths for fiber lasers are ytterbium, erbium, thulium, and holmium. Neodymium emits in the ytterbium laser band, and is included for comparison because of its wide use in other solid-state lasers, but it is rarely used in fibers.

As you learned in Section 8.2, rare earths have many possible transitions, and their behavior on those transitions also depends on the hosts that contain them. Rare earths share many properties that make them attractive for laser use, particularly unusual sets of closely spaced energy levels linked to a ground state, like those shown in Figure 9-8. Electric fields split both the upper and lower laser states into closely spaced sublevels, a process called the Stark effect. Pump light at left pumps atoms from the ground state into a sublevel of the upper laser band, and the atom then drops to the bottom of the upper laser level. Stimulated emission from that state drops the atom down to one of the sublevels above the base of the ground state, allowing the laser to emit across a range of wavelengths.

Only a few transitions look like this, but they work very well for lasers because most of the energy in the pump photon emerges in the stimulated emission. Ytterbium is the rare earth most widely




Figure 9-8. Upper and lower laser states in rare-earth elements may be split into many separate levels, allowing laser emission across a range of wavelengths. Spacing of sublevels is not uniform.

used in fiber lasers because it has a set of laser transitions like this which emit light from 1015 to 1120 nm. Erbium, thulium, and holmium have similar laser bands at longer wavelengths in the near infrared.

These interactions produce gain on transitions between sublevels of the upper state and sublevels of the ground state. If the transition drops to upper sublevels of the ground level, which are sparsely populated at normal temperatures, the system acts like a four-level laser. If the transition drops to a sublevel close to the ground state, which has a large population at normal temperatures, it acts like a three-level laser.

Such systems are called *quasi-three-level lasers* because they are not quite three-level or four-level lasers. Their gain is low when they act like three- or four-level lasers, but in between those extremes they have higher gain and properties that make them very attractive for use in fiber lasers. Erbium and ytterbium are good examples because they have strong pump bands close to laser emission lines.

One important consequence of this energy-level structure is that diode pumping can be quite efficient. If the pump diode excites the laser species directly to a state near the bottom of the upper of the two bands shown in Figure 9-8, a large fraction of that energy will reappear in the stimulated emission photon. For example, erbium can be pumped at 1480 nm and has peak gain around 1535 nm. Likewise, ytterbium can be pumped at 975 nm and has peak gain at 1030–1080 nm.

Another consequence is that gain is possible across a broad range of wavelengths covered by the spread among sublevels. Laser oscillation can be tuned to specific wavelengths in that range, or the laser cavity can be designed for broadband oscillation to produce short pulses. The gain bandwidth is particularly broad in ytterbium, which helps it work better in a diode-pumped fiber laser than neodymium, which has a narrow laser line.

Now let us look at the details of the key types.

9.4.1 Ytterbium-Doped Fiber Lasers

The most powerful fibers lasers are doped with ytterbium and emit pulsed or continuous waves between 1015 and 1120 nm, with peak output at 1030 to 1080 nm. The wavelength depends on

factors including the laser cavity length and optics, and how the laser is excited. Typically, single semiconductor lasers or arrays of semiconductor lasers pump along the length of the doped fiber in one or both directions. The strongest pump bands for ytterbium in the germanate–silicate glasses normally used in fiber lasers are broadband centered at about 915 nm and a narrow peak of much stronger absorption at about 975 nm.

Having pump bands so close to the output band means that the quantum defect of light from the pump photon lost in the output photon is unusually small, so the optical-to-optical conversion of pump diode light into ytterbium laser output is unusually high. Figure 9-9 compares pump and laser bands of ytterbium with that of neodymium, as described in Section 8.4.2. Ytterbium emits across a wide range of wavelengths, but in this comparison we assume it emits at the same 1064-nm wavelength as neodymium, a common practice for industrial applications originally developed for neodymium lasers. The quantum defect in this case is only 8.35% for ytterbium




Figure 9-9. Quantum defects in ytterbium and neodymium, compared for ytterbium emitting at the 1064-nm laser line of Nd-YAG, calculated using Equation 9-2.

pumped at 975 nm, compared with 24% for neodymium pumped at 808 nm. The quantum defect would be even less—5.3%—for ytterbium pumped at 975 nm and emitting at 1030 nm.

Ytterbium also has other advantages for fiber lasers. It lacks self-quenching effects that can reduce the efficiency of neodymium, and it can be used at higher concentrations, increasing power. Importantly, its pump bands at 915 to 980 nm match the emission band of highly efficient InGaAs diode lasers. Combining all these benefits makes it possible for high-power Yb-fiber lasers to have wall-plug efficiencies from 25% to as high as 50%. As shown in Table 9-1, the higher the efficiency, the less waste heat is generated. Excess heat can generate thermal gradients which degrade beam quality, so higher efficiency can improve beam quality as well as reduce power consumption. That, in turn, allows commercial Yb-fiber lasers to achieve single-mode power above 10 kilowatts and multimode power up to 100 kilowatts.

Careful design is required to achieve such extreme performance. Figure 9-10 gives an example of what IPG Photonics did to achieve a record-setting 10-kilowatt output from a single-mode fiber laser in 2009. The system started with a diode-pumped master oscillator that generated 1 kilowatt at 1070 nm. That output was directed into a 15-meter long fiber amplifier pumped by 45 Yb-fiber lasers each emitting 300 watts at 1018 nm, pumped by 975-nm diodes. Those fiber pump lasers delivered higher intensity input




Figure 9-10. Simplified view of IPG Photonics design for a 10-kilowatt single-mode fiber laser. Cavity optics of the oscillator stage were fiber Bragg gratings. Black ovals show side couplers and pump combiners.

than diodes, improving optical pumping efficiency. The fiber amplifier had a 30- μm core with effective mode area of 500 μm^2 , helping it reach 10-kilowatt output without damaging the fiber. Each pumping step had a quantum defect of 4%–5%, reducing the heat load that had to be removed from the high-power amplifying fibers.

The high efficiency and high power of Yb-fiber lasers has led to a tremendous expansion of industrial laser applications, from cutting and drilling to additive manufacturing, as described in Chapter 13. Ytterbium-doped fibers are by far the most commercially successful fiber lasers in industry, and typically what people mean when they say “fiber laser” in industry.

9.4.1.1 Ytterbium Bandwidth and Short-Pulse Operation

The unusually broad gain bandwidth of ytterbium-fiber lasers offers two important advantages over narrow-line lasers such as neodymium. One is sufficient bandwidth to generate ultrashort pulses, valuable for both industrial applications and research. The other is the ability to tune output across that range, although most Yb-fiber applications are not particularly sensitive to wavelength.

Mode-locked ytterbium-fiber oscillators can generate pulses as short as 50 femtoseconds. Shorter pulses can be generated by taking advantage of nonlinear effects which are relatively high in fibers because of the small core diameter. Using special high-nonlinearity fibers can increase the spectral bandwidth. As will be described in Sections 11.3 and 11.4, such nonlinear effects can stretch a fiber laser’s emission band across more than an octave—a factor of two in wavelength or frequency.

The refractive index of glass in the ytterbium-fiber band decreases as the wavelength increases, causing shorter wavelengths to pass more slowly through the glass. This disperses the light by color as it passes through the glass, with the lowest frequencies (the longest wavelengths) going the fastest, and the frequency rising with time in what is called a *positive chirp*. This stretches the duration of the pulse as the nonlinear effect stretches its bandwidth, and the further the pulse travels through the fiber, the more the pulse is extended. Stretching the pulse duration also decreases its peak power.

The chirp effect can be reversed by passing the light through external optics that delays the longer wavelengths more than the shorter ones, reducing the duration of the pulse and increasing its

peak power. That is a necessary step to reduce the pulse length after the bandwidth has been increased.

An external amplifier can be included to boost total energy in the pulse as well as peak power. These optical tricks are combined in a technique called *chirped pulse amplification*, which spreads the pulse bandwidth, stretches its duration, then amplifies it and compresses it to produce pulses that have shorter duration and higher energy and peak power than the pulses emerging from the laser oscillator.

9.4.1.2 *Harmonic Generation and Wavelength Conversion*

The high power and energy density available from fiber lasers allows the use of nonlinear optics to convert the fundamental output of the fiber laser to harmonics and other wavelengths.

The second harmonic is the most important for ytterbium. Typically this is at 515 nm, produced by doubling 1030 nm, but harmonics also are available at longer green wavelengths up to 540 nm, including that at the 532 nm line of neodymium.

Other wavelengths also can be generated externally, such as by Raman shifting or by using Yb-fiber lasers to pump Raman fiber lasers, described in Section 9.6.

9.4.1.3 *Femtosecond Fiber Lasers*

The ability of ytterbium-fiber lasers to generate sub-picosecond pulses has led to femtosecond fiber lasers, a new class of lasers with a wide range of uses. Previous generations of ultrashort-pulse lasers were complex, delicate and costly devices that required skilled laser operators and carefully controlled environments. All-fiber beam paths make fiber lasers much less sensitive to vibration than previous ultrafast lasers. One manufacturer highlighted this by mounting a femtosecond fiber laser on a vibration table at a trade show. A friend told me “you’ve got to see this” and took me to the booth, where I was duly amazed.

Such femtosecond fiber lasers are examples of how commercial lasers are evolving from sophisticated instruments designed for experts into highly automated tools for use by nonspecialists. Some are intended for research use in fields such as biology or chemistry, where femtosecond measurements can yield new insight into the properties of materials or organisms. These lasers are essentially black boxes, with a few controls on the outside, a computer or network interface of some sort, a power cord and an output port.

The user selects the laser output they want manually or through a computer interface, and the automated laser does the rest, sending relevant information to the user's computer network.

A more dramatic change is the growing use of femtosecond fiber lasers for machining or even surgery. Femtosecond pulses briefly achieve extremely high intensity but do not deliver much energy, so they can ablate thin layers of material without damaging the substrate. That is invaluable in delicate processes such as machining the special ruggedized glass used for smartphone screens so the glass retains its structural integrity without cracks after machining.

9.4.2 Erbium-Doped Fiber Lasers

Erbium-doped fiber lasers have a laser transition similar to ytterbium, but at lower energy, corresponding to wavelengths of 1520 to 1630 nm. Peak gain is from 1530 to about 1570 nm. As with ytterbium, the exact wavelength depends on laser design. By a fortunate coincidence, this band spans the wavelengths where glass optical fibers used in telecommunications are most transparent, leading to the use of erbium-doped fibers in optical amplifiers, described in Section 9.5. Erbium-fiber lasers were demonstrated first, but optical amplifiers were their first major application.

Erbium has two primary absorption bands for pump lasers at 980 and 1480 nm. The 980-nm absorption excites erbium atoms to a high energy state, which then drops to the upper laser level. The 1480 nm line excites erbium directly to the upper laser level, so that pump band has a quantum defect of only about 5% and can achieve higher optical-to-optical efficiency. It allows much higher efficiency than the 37% quantum defect when pumping with 980-nm pump diodes, but 980-nm diodes have higher electrical-to-optical efficiency, and can generate more power than 1480-nm diodes, so they are used more often. Adding ytterbium to the erbium-fiber core can increase 980-nm pump efficiency. Ytterbium absorbs more efficiently and can transfer the absorbed energy to nearby erbium atoms, which drop to the upper laser level and emit laser light.

Erbium-fiber lasers can reach continuous beams to 2 kilowatts, but cannot match the high power of ytterbium. However, the longer erbium wavelength is more attractive for some important applications. One is in telecommunication systems, where fiber lasers can offer performance not available from semiconductor laser sources. Wavelengths longer than 1500 nm also have an important eye safety

advantage because they are strongly absorbed by ocular tissue, preventing them from damaging the retina. Pulsed lasers in the 1- μm region pose serious eye hazards because stray pulses can damage the retina. (See Appendix A for more information on laser eye safety.)

Like ytterbium, erbium-fiber lasers have adequate bandwidth to produce ultrashort pulses and are available in femtosecond versions. Because glass has anomalous or negative dispersion in the 1500-nm range, erbium-doped glass fibers can generate special pulses called solitons, which retain their pulse shape and duration over long distances but are inherently limited in energy. Most other lasers lack that capability.

Harmonic generation with pulsed erbium-fiber lasers can generate 780 nm in the near infrared, with better beam quality than from semiconductor diode lasers.

Erbium has a strong laser line near 2.8 μm that cannot be used in conventional silica glass fibers because they absorb light very strongly in that range. However, fibers drawn from fluoride glasses have low loss at wavelengths beyond 2 μm , and they can be used to make erbium-fluoride lasers emitting near 2.8 μm .

9.4.3 Thulium-Doped Fiber Lasers

A third type of fiber laser offered commercially is the thulium-doped fiber laser, which has laser gain from about 1.9 to 2.07 μm and can generate laser power to a kilowatt. It is the most powerful laser source in that infrared band, which has attracted attention for applications in medicine, pollution measurements, and materials working on plastics. It also can be used to pump solid-state or fiber lasers emitting at longer wavelengths.

Thulium has an unusual energy-level structure that allows two types of optical pumping. One is with an erbium-doped fiber laser emitting at 1600 nm, which can raise thulium ions to an excited state that decays to the upper laser level and yields stimulated emission at 1.9 to 2.07 μm . Alternatively, a 793-nm diode laser can excite thulium atoms to a much higher energy level that drops into the upper laser state and spontaneously emits a second photon which can excite another thulium atom to the upper laser level. Thus a single 793-nm photon could provide energy to pump two separate thulium atoms into the upper laser level. This is unusual for lasers, but it can enhance the efficiency of thulium.

Thulium is near the long-wavelength limit for silica glass fiber lasers. Glass absorption increases rapidly at longer than $2\ \mu\text{m}$, limiting thulium laser oscillations to wavelengths shorter than about $2.07\ \mu\text{m}$.

9.4.4 Holmium-Doped Fiber Lasers

Holmium-fiber lasers have operated at wavelengths longer than $2.1\ \mu\text{m}$, the longest yet obtained from lasers in silica glass fibers. Those wavelengths are attractive for medical applications because water present in tissue absorbs strongly at $2.1\ \mu\text{m}$, so light penetrates only about 0.4 millimeter into tissue. Other potential applications include remote sensing, lidar, and research into other mid-infrared sources.

So far, holmium pumping has required a series of lasers pumping each other. For example, a diode-pumped erbium-fiber laser emitting at $1.55\ \mu\text{m}$ pumps a thulium-fiber laser emitting at $1.95\ \mu\text{m}$, which in turn pumps a $2.1\text{-}\mu\text{m}$ holmium fiber laser. Potential applications include prostate surgery and fragmentation of kidney stones. So far most operation has been in silica glass fibers, holmium-doped fluoride fibers having much lower loss have been demonstrated at the lab.

9.4.5 Other Fiber Lasers

The principle of the doped fiber laser is quite general and can be applied to fibers of other glassy materials doped with light-emitting elements other than rare earths. The key factors are that the host material must be suitable for drawing into fibers and the dopants must have suitable laser transitions. Their primary applications are in generating light at wavelengths where oxide glasses are not transparent, typically at wavelengths longer than $2\ \mu\text{m}$. So far, most of this work is developmental and remains in the laboratory.

The physics of drawing molten materials into fibers is quite complex, and most materials are not suitable. The material should form a glass when cooled rapidly, with the atoms arranged at random rather than arrayed in regular patterns to form a crystal. The liquid form of the material should be thick (viscous) so it holds together when stretched and cooled. Oxide-based glasses such as silica are the most familiar example of good fiber materials, but thick sugar syrup also forms fibers; we call them cotton candy. The fibers

also should be able to withstand normal environmental conditions and not dissolve when exposed to moist air or become brittle when exposed to air.

Some nonoxide glasses can be drawn into usable doped optical fibers. The most widely used to date is a blend of fluoride compounds of zirconium, barium, lanthanum, aluminum, and sodium fluorides-called ZBLAN ($\text{ZrF}_4\text{-BaF}_2\text{-LaF}_3\text{-AlF}_3\text{-NaF}$). The compound transmits well between 0.3 and 5 μm and has been used in rare-earth fiber lasers across that range, with several emitting more than a watt. The main interest in those materials is for wavelengths beyond 2 μm , where the best powers are over 20 watts with erbium at 2.8 μm and around 10 watts with holmium near 2.1 μm .

Another class of fiber host materials is compounds containing the elements sulfur, selenium, and tellurium in place of oxygen. These are called *chalcogenide* compounds. Sulfur generally forms the best glasses, followed by selenium and tellurium. Both groups can be used with rare-earth lasers emitting at wavelengths longer than 2 μm , where silicate glasses start absorbing light strongly. Elements including arsenic and antimony can be added to chalcogenide glasses to make fibers with desired properties.

9.5 RARE-EARTH-DOPED FIBER AMPLIFIERS

Optical amplifiers are essentially a gain medium without a resonant cavity, which amplifies light that enters it from one end and emits a higher power at the other, as described in Section 6.5. Optical amplifiers can be used as additions to laser oscillators that generate more power than the oscillator alone can produce, as the master oscillator power amplifier covered in Section 6.5.3. They also can be used to amplify signals transmitted in optical form in telecommunication systems, which this section covers.

9.5.1 Optical Amplifiers in Communications

Fiber optics are now widely used in communications networks because they have a high transmission bandwidth and a low transmission loss, so they can carry a lot of information (in bits per second) over a long distance. However, glass fibers are not perfectly transparent, so signals fade slowly with distance. After 50 to

100 km, the light signal becomes so faint that it must be amplified to travel further.

As described in Section 9.1.3, erbium-doped fiber amplifiers (sometimes called *EDFAs*) emerged to serve that need in the 1990s. The erbium gain band covers the 1550-nm band where optical fibers have their minimum loss, and erbium-doped fibers have many desirable properties, most importantly their ability to amplify many separate signals at closely spaced wavelengths in that band, the standard used for global communications networks.

A single optical fiber can carry optical signals at many different wavelengths, just as the air transmits radio signals at many different frequencies. As long as the signals are at distinct wavelengths or frequencies that can be separated optically from each other, they do not interfere with each other. In fiber optics this is called *wavelength-division multiplexing (WDM)* and it multiplies how much information a fiber can transmit.

Fiber amplifiers normally are used on systems carrying signals over 100 km or more. Typical data rates are 10 or 100 gigabits per second on each wavelength, with capacity increased by adding more wavelength channels that are evenly spaced (in frequency) from 1530 to 1565 nm. The standard spacing for *dense WDM* is 50 gigahertz, which can fit about 100 separate wavelengths into the 1530- to 1565-nm band. At 100 gigabits per second, the capacity of 100 wavelengths comes to a total of 10 terabits (trillion bits) per second, roughly 20 two-layer Blu-ray discs a second.

In practice, most fiber-optic links in the global communications network carry only a few wavelengths, so their individual capacity is lower. But capacity can be increased by lighting additional wavelength slots in the fiber as needed. Research continues on ways to increase total capacity.

9.5.2 Types of Optical Amplifiers

Virtually all modern long-haul high-speed communication networks operate in the standard 1550-nm band. Developers are studying the use of other wavelengths to further increase network capacity, which will require different types of fiber amplifiers.

Today's standard erbium amplifiers are designed to operate at 1530 to 1565 nm, called the *C band* in the communications industry. Erbium has a much broader gain bandwidth, so erbium amplifiers also can be designed to operate at 1565 to 1625 nm, called the *L band*

for long wavelength. The L band is not in wide use because erbium gain is lower in that range and fiber attenuation is higher. However, it is expected to be used more as bandwidth needs increase.

Many other types of fiber amplifiers have been demonstrated in the laboratory. Ytterbium-doped fibers can be used as amplifiers as well as for lasers, but fiber attenuation is much higher in the 1015- to 1120-nm ytterbium band, so it is not used for communications. Fibers doped with the rare-earth praseodymium emit in the 1300-nm band used for short- and medium-distance fiber-optic systems, but they do not match the performance of erbium amplifiers in the 1550-nm band, so they are little used.

In Section 9.6, you will learn about another type of fiber amplifier based on stimulated Raman scattering that is used in some communications systems, mostly to complement EDFAs.

In Chapter 10, you will learn about semiconductor optical amplifiers, which are resonator-less versions of semiconductor diode lasers.

9.5.3 Structures of Fiber Amplifiers

Figure 9-11 shows how an EDFA works in a communication system. A weak optical signal enters the left end of the erbium-doped amplifying fiber. Light from a pump diode enters from the right through a coupler to excite erbium atoms in the coiled fiber, producing a population inversion. The weak input signal stimulates emission at the wavelengths in the signal, which is amplified as it passes through the coil. The coupler at the right end of the coil




Figure 9-11. Erbium-doped fiber amplifier, as used in a communication system. Internal fiber structure is not shown for simplicity.

transmits the amplified signal toward the output fiber at the lower right.

Light being amplified makes only a single pass through the amplifying fiber, with stimulated emission increasing signal strength along the way. Filters at the ends of the amplifying fiber treat the pump and signal wavelength differently. The input filter at left transmits the signal wavelength from left to right, but either absorbs or reflects any pump light that reaches it from the right end. In practice, little pump light survives the round trip through the amplifier because virtually all of it is absorbed by erbium atoms.

The amplified signal light leaving the amplifier passes through a wavelength-selective coupler into the fiber transmitting the signal through the next span of the communication system. The second arm of the coupler delivers the pump wavelength into the amplifier fiber. The pump light is depleted as it excites erbium atoms in the amplifier; what little is reflected from the filter at the signal input is further depleted by passing a second time in the opposite direction through the amplifying fiber.

9.5.4 Fiber Signal Amplifier Performance Considerations

Fiber signal amplifiers must meet stringent performance requirements for state-of-the-art communication networks. The most important for our purposes are uniformity of the signal gain across the transmission band and signal-to-noise ratio.

The gain of any laser medium depends on the wavelength, and that variation is strongest when the input power is small, as shown in Figure 9-12. Erbium gain is strongest at 1530 to 1535 nm, levels off between 1540 and 1555 nm, then declines at longer wavelengths, as shown in the upper plot for low input power of about 10 microwatts. The gain decreases and becomes flatter at higher input, as shown by the lower plot for input of 1 milliwatt.

The variation of the gain across the spectrum poses a problem for long WDM systems because they carry many wavelengths through a series of amplifiers. If the highest gain on any channel was twice the lowest gain, and that the signal had to pass through five identical amplifiers, the weakest signal would end up only $(0.5)^5$ times the peak signal—a mere 3%. After 10 amplifiers—at best only spanning 1000 km—the difference between the strongest and the weakest channels is a factor of 1000. To avoid that, amplifiers for WDM systems are designed with filters and other optics to make




Figure 9-12. Gain of an erbium-fiber amplifier differs across the spectrum. The upper curve shows high but nonuniform gain at low input power (-20 dBm is $10\ \mu\text{W}$) and more uniform but lower gain at higher input power (0 dBm is 1 mW). (Courtesy of Thorlabs)

gain uniform across the spectrum. (Another alternative is balancing gain by using amplifiers with different gain spectra, such as the Raman fiber amplifiers described in Section 9.6.1.)

The signal-to-noise problem comes from the fact that laser amplifiers also produce some spontaneous emission when atoms in the upper laser level drop to the lower laser level and emit light before another photon arrives to stimulate emission. Most spontaneously emitted photons in a fiber amplifier leak out of the fiber, but a few are directed along the length of the fiber—and in that case they are amplified along with the signal light. This *amplified spontaneous emission* is noise that builds up along the length of the chain of fiber amplifiers, creating background noise that cannot be filtered out, limiting system performance.

9.6 RAMAN FIBER LASERS AND AMPLIFIERS

A separate family of fiber lasers (and amplifiers) is based not on rare-earth doped fibers but on an effect called *Raman scattering* that occurs when photons lose or gain a bit of extra energy when they bounce off atoms in undoped glass fibers.

Raman scattering, described in Section 5.6.5, is a relatively weak process in which an atom absorbs a photon at one wavelength

and then almost instantly emits a photon at a different wavelength, because of a change in vibrational energy. In theory, the atom could either convert some of the photon energy to vibration in the material or take vibrational energy from the material and add it to the energy of the emitted photon. However, in practice the scattered photon usually has less energy than the input photon.

At low light levels, only a very small fraction of photons are Raman scattered. But intense light illuminating a material can cause a nonlinear effect called *stimulated Raman scattering*. The brighter the light, the more photons are shifted in wavelength by Raman scattering—and these Raman-shifted photons are more likely to stimulate atoms they encounter to emit Raman-shifted photons, usually with less energy than the illuminating photons. The strength of stimulated Raman scattering increases nonlinearly with the power of the input light.

This process can become the basis of either optical amplification or a laser. It works well in fibers and also can work in bulk materials including liquids and gases (which in some cases may be used to fill a hollow-core fiber).

9.6.1 Fiber Raman Amplification

Fiber Raman amplification is attractive in telecommunication systems because the amplification can take place in the undoped core of the fiber transmitting the signals. In this case, the pump light travels in the opposite direction to the signal, starting from the receiver rather than the transmitter and amplifies the signal as it approaches the receiver. As mentioned in Section 9.5.4, gain in Raman amplifiers increases with longer wavelengths, so it can offset the decrease in gain with wavelength in erbium-fiber amplifiers.

Raman amplifiers work by transferring energy from a strong pump beam to a weak signal beam. The *cross-section* or probability for stimulated Raman emission is small in the silica (SiO_2) used in standard optical fibers. However, the small cores of single-mode fibers concentrate light in a small volume, increasing its intensity, and the probability of stimulated Raman scattering rises with the distance the light passes through the fiber. Thus the relatively weak amplification of the signal beam accumulates through kilometers of fiber where the pump beam is strong.

Gain in a Raman fiber amplifier comes from Raman scattering in the glass of an undoped fiber and is distributed along the length




Figure 9-13. Raman laser converts the 1100 nm pump laser beam to the longer 1178 nm wavelength, with gain peaking at the longer wavelengths, offsetting the peak at shorter wavelengths in an erbium-doped fiber amplifier.

of the fiber. The strength of the gain depends on the pump power, so Raman gain is strongest at the end closest to the pump laser.

The Raman amplification process shifts energy from the strong pump beam to amplify the weaker signal beam, which becomes stronger. For example, Raman gain in silica peaks at a frequency 13 terahertz (THz) lower than the pump laser frequency, so stimulated Raman scattering would shift energy from an 1100-nm pump to amplify a weak signal beam at 1178-nm, as shown in Figure 9-13. With enough pump power and hundreds of meters of fiber, more than half of the input energy can be converted to the shorter wavelength.

At first glance, Raman amplification might seem to be just another variation on a fiber amplifier. But it has two important advantages over other amplifiers. One is that it can be tailored to work at any wavelength as long as you can find a pump source shifted by 13 THz from the light you want to amplify. The second is that the gain spectrum of a Raman amplifier nicely complements that of the erbium-fiber amplifier in Figure 9-12. Erbium gain peaks at short wavelengths in the gain band, while Raman gain peaks at longer wavelengths, so combining the two can make gain more uniform across the amplified band of a WDM system.

9.6.2 Fiber Raman Laser Oscillation

Stimulated Raman scattering also can produce laser oscillation in a suitable resonant cavity. The trick to making a Raman laser oscillate is to use spontaneous Raman emission produced by an intense

pump beam as the starting seed for amplification. Stimulated Raman scattering then amplifies the initially weak spontaneous signal, building up stimulated Raman emission that resonates in the laser cavity.

With high-power lasers available at only a few wavelengths, the ability of stimulated Raman scattering to shift wavelengths can be valuable in developing new laser sources. This ability to shift wavelength is important because efficient, high-power lasers are available only at a few wavelengths. It is possible to cascade two or more Raman fiber lasers to generate additional wavelengths, with the 1178-nm output of the Raman fiber laser in Figure 9-13 pumping another Raman stage emitting at a longer wavelength. Although each step loses energy, the process can produce otherwise unobtainable wavelengths.

Likewise, the output of a Raman fiber laser can be frequency-doubled to produce visible wavelengths not otherwise available at reasonable power. The second harmonic of the 1178-nm Raman laser is at 589 nm in the yellow, where good lasers are scarce. That specific wavelength is needed to excite sodium atoms in the upper atmosphere for use in adaptive optics systems that allow giant ground-based telescopes to see clearly through the air.

9.7 WHAT HAVE WE LEARNED?

- Fiber lasers are solid-state lasers in which the light-emitting material is in the core of an optical fiber.
- Total internal reflection guides stimulated emission along the length of a fiber laser.
- Light in fiber lasers may stay within the fiber from the pump laser to the end of the beam-delivery fiber, greatly easing maintenance and operation.
- Wall-plug efficiency of fiber lasers can reach 50%, the highest of any nondiode laser.
- Fiber lasers usually are diode-pumped, but can be pumped by other fiber lasers.
- The tight confinement of light, efficiency, and the ease of dissipating waste heat allow industrial fiber lasers to produce steady beams to tens of kilowatts.
- Single-mode fiber lasers have excellent beam quality.

- Rare-earth elements used in fiber lasers include ytterbium, erbium, thulium, and holmium.
- Many fiber lasers have three layers, an inner core doped with light-emitting rare earths that emits laser light, an inner cladding carrying pump light, and an outer cladding keeping the pump light inside the inner cladding so it can pump the core.
- The inner cladding of a dual-clad fiber is sometimes called the outer core because it confines the pump light.
- Fiber Bragg gratings are layered structures like multilayer coatings that are fabricated inside optical fibers; they are used as cavity mirrors on fiber lasers.
- Microstructured fibers can enhance fiber laser performance.
- Ytterbium- and erbium-doped fiber lasers are diode pumped. They may be used to pump other fiber lasers emitting at other wavelengths where good pump diodes are not available.
- The high efficiency of fiber lasers comes from the efficient conversion of pump light from highly efficient diode lasers.
- Fiber laser power is limited by nonlinear effects.
- Special core structures that confine light in a larger area allow higher fiber laser power.
- Ytterbium-doped fiber lasers are the most powerful type. Their output is at 1015 to 1120 nm.
- Little of the energy in a photon pumping a ytterbium atom to the upper laser level is lost. This loss is measured as the quantum defect.
- Ytterbium has a broad emission bandwidth, so ytterbium fibers can generate ultrashort pulses.
- Erbium-doped fiber lasers emit at 1520 to 1630 nm and are diode pumped.
- Thulium-doped fiber lasers emit at 1.9 to 2.07 μm .
- Holmium-doped fiber lasers emit at wavelengths up to 2.1 μm .
- Erbium-doped fibers are used in amplifiers to boost signal strength in long-distance telecommunications because they amplify light at 1530 to 1570 nm where fibers have the lowest losses.
- Erbium-doped fibers can simultaneously amplify signals at many separate wavelengths, increasing the capacity of fiber-optic communications.
- Stimulated Raman scattering in optical fibers can complement erbium-doped amplifiers in boosting signals evenly across the spectrum.

- Fiber lasers can pump stimulated Raman scattering in an optical fiber in a resonator to make a fiber Raman laser at a different wavelength.

WHAT'S NEXT?

In Chapter 10, we will learn about the fast-moving technology of semiconductor diode lasers.

QUIZ FOR CHAPTER 9

1. Which of the following is not an advantage of fiber lasers?
 - a. High wall-plug efficiency
 - b. High surface-to-volume ratio helps dissipate heat
 - c. High beam quality
 - d. Ability to generate extremely high-energy pulses
 - e. Light path contained entirely within fiber
2. What is the function of an inner cladding in a dual-core fiber laser?
 - a. It collects and guides pump light along the fiber.
 - b. It collects and guides stimulated emission.
 - c. It confines the active species so the dopant does not leak out of the fiber.
 - d. It confines stimulated Raman emission.
 - e. Coolant flows through it.
3. Where is the light-emitting laser species in a dual-clad fiber laser?
 - a. All through the fiber to maximize light conversion
 - b. Contained entirely in the core
 - c. Contained in both the core and inner cladding
 - d. Entirely in the inner cladding
 - e. Filling the holes in microstructures running the length of the fiber
4. What is the most powerful rare-earth-doped fiber laser?
 - a. Neodymium
 - b. Erbium
 - c. Thulium
 - d. Yttrium
 - e. Ytterbium

5. What types of cavity optics are used in fiber lasers?
 - a. Multilayer dielectric coatings applied to the fiber ends
 - b. Fiber Bragg gratings
 - c. Reflective metal layers embedded within the fiber
 - d. No cavity needed because of total internal reflection within fiber
 - e. Fresnel reflection at end of fiber
6. Which other laser could be replaced with a ytterbium-fiber laser emitting at the exact same wavelength?
 - a. Ruby
 - b. Neodymium-YAG
 - c. Cobalt-ZnSe
 - d. Alexandrite
 - e. Erbium-YAG
7. What is the quantum defect when thulium-fiber laser emitting at 1950 nm pumps a holmium-fiber laser emitting at 2100 nm? (Use Equation 9-2)
 - a. 7.14%
 - b. 7.69%
 - c. 8.3%
 - d. 15%
 - e. 92.9%
8. You need a fiber laser that delivers 1 kW of light energy. How much input power could you save if you bought a laser with wall-plug efficiency of 40% than if you bought one with 20% efficiency?
 - a. 0.2 kW
 - b. 0.4 kW
 - c. 1 kW
 - d. 2.5 kW
 - e. 5 kW
9. What allows ytterbium-fiber lasers to generate pulses as short as 50 femtoseconds?
 - a. High energy efficiency
 - b. Excellent beam quality
 - c. Broad emission bandwidth
 - d. Can be switched off and on quickly
 - e. It is not possible
10. What limits the development of glass-based fiber lasers at wavelengths much longer than 2 μm ?
 - a. High absorption of silica glass
 - b. Lack of suitable light-emitting elements

- c. Lack of suitable pump diodes
 - d. Inability to fabricate dual-clad fibers
 - e. Lack of any applications for such fiber lasers.
11. What active species is used in fiber amplifiers for telecommunications?
- a. Neodymium
 - b. Erbium
 - c. Thulium
 - d. Yttrium
 - e. Ytterbium
12. What element must be added to a glass fiber for it to be used as a Raman amplifier?
- a. Ytterbium
 - b. Neodymium
 - c. Thulium
 - d. Erbium
 - e. None

DIODE AND OTHER SEMICONDUCTOR LASERS

ABOUT THIS CHAPTER

This chapter covers electrically powered lasers made from semiconductors. The most widely used are called *diode lasers* because of their electrical properties, but there are a few others as well. We will start by defining these types and describing the key optical and electrical properties of light-emitting semiconductors. Then we will cover the various types of semiconductor diode lasers and compare them to nonlaser light-emitting diodes. Our descriptions will include structures and materials, and will close with quantum cascade lasers and similar devices, which are electrically powered but are not diodes.

10.1 TYPES OF SEMICONDUCTOR LASERS

Technological progress has made it hard to sort lasers into simple and clearly defined categories. This is particularly true for semiconductor lasers, so let us take a brief but careful look at the variety of semiconductor lasers and their most important differences.

The term *semiconductor laser* refers to lasers made from materials with electrical properties intermediate between conductors and insulators, and thus are called semiconductors. Silicon is the best known semiconductor because of its wide use in electronics,

but it is a very poor light emitter. Semiconductor lasers are made of *compound semiconductors* which contain two or more elements, such as gallium arsenide (GaAs), gallium nitride (GaN), and indium–phosphide (InP). You will learn more about compound semiconductor materials in Section 10.3.

Most semiconductor lasers are two-terminal devices called *diodes* powered by an electric current that passes in one direction through the semiconductor. The key factor that identifies a diode is the presence of an internal *junction* between two regions of semiconductor with slightly different composition that conduct current in different ways, as described in Sections 10.2 and 10.3. Properly they are called *semiconductor diode lasers* but typically are known as *diode lasers*, *laser diodes*, or sometimes just *diodes*. Other common terms are *junction laser*, because lasing occurs at the junction between two types of semiconductor, and *injection lasers* because current is injected into the semiconductor to make them operate. In this chapter, we will just call them diode lasers or diodes, dropping the “semiconductor” like most people drop their middle names in daily life. Most of this chapter is about the various ways diode lasers can be made.

Another important type of semiconductor laser is powered by electric current but is not an electrical diode. The best known of these is the *quantum cascade laser*, described in Section 10.10. They are in this chapter because they are semiconductor devices powered directly by electricity.

A third type of semiconductor laser is the *optically pumped semiconductor laser (OPSL)* described in Section 8.5. They have their own distinctive internal structure and are powered by light from another laser, almost always by a diode laser. They are covered in Chapter 8 because they are optically pumped like the other solid-state (non-semiconductor) lasers covered in that chapter.

Semiconductor research is a hot area, and new variations are being developed in the laboratory.

10.2 DEVELOPMENT OF DIODE LASERS

Semiconductor lasers have much in common with the other lasers described so far. They emit a laser beam when spontaneous emission triggers a cascade of stimulated emission from a population

inversion inside a resonant optical cavity. As in a gas laser, the excitation energy comes from an electric current passing through the laser material, but the material in a diode laser is a solid fabricated so the current flowing through it produces a population inversion.

Light emission from semiconductors had been known for decades when the first diode lasers were made in 1962. Diode lasers had their roots in *light-emitting diodes* or *LEDs*. Henry J. Round made the first LEDs in Britain in 1907. In the 1920s, Oleg Losev in the Soviet Union independently rediscovered the idea and studied LEDs in more detail. Other physicists demonstrated LEDs as they learned more about semiconductor physics, but through the 1950s their emission was feeble. It changed in 1962 with the dramatic demonstration of bright infrared emission from gallium arsenide by Robert Rediker at the MIT Lincoln Laboratory. Lasers were a hot topic then, and within a few months, Robert Hall of General Electric in Schenectady, NY, had extended that to make the first diode lasers, and Nick Holonyak Jr., working for GE in Syracuse, NY, had made both the first bright visible LEDs and the first visible diode lasers.

The fundamental difference between the LEDs and the diode lasers was that diode lasers were placed between mirrors and pushed across the threshold for laser oscillation by driving them with a much higher current than used for early LEDs. The stronger current produced a larger population inversion and higher gain, yielding a cascade of photons.

The first diode lasers had to be operated at the cryogenic temperature of liquid nitrogen, 77°K (−196°C or −321°F). It took years to develop versions that could generate a continuous beam at room temperature and more time to improve them so they could operate for more than a few seconds or minutes at a time. But progress has been rapid and steady since the development of the first long-lived diode lasers in the mid-1970s.

Today diode lasers are a pervasive if often hidden technology. They are tiny semiconductor chips like the one shown enlarged in Figure 10-1. Life size, they look like tiny pieces of flat metallic confetti, each about the size of a grain of salt. Billions have been sold, making diodes by far the most common lasers. Red diode-laser pointers are cat toys. Diode lasers read product codes in stores, play video on DVDs or Blu-Ray discs, and transmit data through the fiber-optic backbone of the Internet. Diode lasers align walls for




Figure 10-1. Square diode laser chip near the eye of a needle. (Courtesy of Jet Propulsion Laboratory, government work not subject to copyright)

construction crews and power industrial fiber lasers. Behind the scenes, diode lasers power most fiber and solid-state lasers.

10.3 SEMICONDUCTOR BASICS

To understand how diode lasers work, you need to learn a little semiconductor physics. The field as a whole is beyond the scope of this book, but some fundamental aspects are crucial to the operation of diode lasers. This section will introduce you to semiconductor and the materials used in diode lasers to help you understand how they work.

10.3.1 Valence and Conduction Bands

Semiconductors get their name because electrons flow through them better than through an insulator like glass, but not as well as a conductor like copper. Current flow in a solid depends on how tightly atoms hold onto electrons in their outer shells. In insulators the electrons are bound tightly, often in covalent bonds between pairs of atoms, so it takes a high voltage to pull them loose. Conductors have only a weak hold on their outer shell electrons, letting them flow freely through the solid if a weak voltage is applied across it. A semiconductor holds onto its outer electrons tighter than a conductor but more loosely than a conductor.




Figure 10-2. Valence and conduction bands in conductor, semiconductor, and insulator.

Figure 10-2 shows the energy levels of electrons in the outer shell of atoms in three different types of solids. The lower-energy states are *valence bands* occupied by electrons with a low energy that bind to atoms in the solid. Electrons with enough energy to move around in the solid (conducting current) are in the higher-energy *conduction band*. In a conductor at left, the top of the valence band overlaps with the bottom of the conduction band, allowing those charge carriers to move easily in the solid. In an insulator at right, the conduction band is much higher than the valence band, and electrons lack the energy needed to cross the large *bandgap* to reach the higher-energy conduction band. In a semiconductor, in the middle, the bandgap between the lower valence band and the higher-energy conduction band is small enough for some valence electrons to escape to reach the conduction band. The bandgap usually is measured in electron volts, the number of volts needed to move an electron from the valence band to the conduction band.

The number of electrons N in the valence and conduction bands depends on the material's bandgap energy ΔE and the temperature T :

$$\frac{N_{\text{conduction}}}{N_{\text{valence}}} = \exp^{\Delta E/kT} \quad (10-1)$$

where k is the Boltzmann constant. This is the same formula that describes the relative proportions of atoms and molecules in a pair of different energy levels.

Plugging numbers into the equation shows that very few electrons are in the conduction band of a semiconductor at room temperature. In 100% pure silicon, where the bandgap is 1.1 electron volts, only 3×10^{-19} of the electrons are in the conduction band at room temperature. That is enough to carry a feeble current, but it gives pure silicon a high electrical resistance.

10.3.2 Electrons, Holes, and Doping

Semiconductor electronics require higher conductivity. All it takes is a few impurities. Looking at silicon will help you understand that, even though silicon is a very poor light emitter.

Silicon atoms have four electrons in their outer shell, and in silicon crystals each of those four electrons forms a bond with an adjacent silicon atom, as shown in Figure 10-3A. If the silicon crystal is perfectly pure, very few of its valence electrons can reach the conduction band at room temperature, as you saw above. However, actual silicon crystals include a few impurity atoms that occupy positions where silicon atoms otherwise would be. If one of those atoms was phosphorous, which has five electrons in its outer shell, four of those electrons would form bonds with adjacent silicon atoms and stay in the valence band, but the fifth would be unbound and free to move about through the solid and conduct current, as shown in Figure 10-3B.

If the impurity atom had only three outer electrons, such as aluminum or gallium, they would bond to three adjacent silicon atoms in the crystal. However a vacancy or *hole* would remain where the fourth outer electron of silicon normally would bond to an adjacent silicon atom, as shown in Figure 10-3C. Such holes can move if a nearby electron shifts to another location in the crystal to fill the vacancy, leaving a hole where the electron had been. This allows holes to move through the crystal as carriers of positive charge to complement the negative charge carried by electrons.

Adding such impurities intentionally to the semiconductor makes it more conductive and more useful for electronic devices. A semiconductor doped with atoms that donate electrons to create free electrons carriers, such as phosphorus in silicon, is called *n*-type because the electrons are *negative* current carriers. Semiconductors doped with atoms that produce holes (or electron acceptors) are called *p*-type because they contain extra holes that serve as *positive* current carriers. The degree of conductivity in *n*- and *p*-type

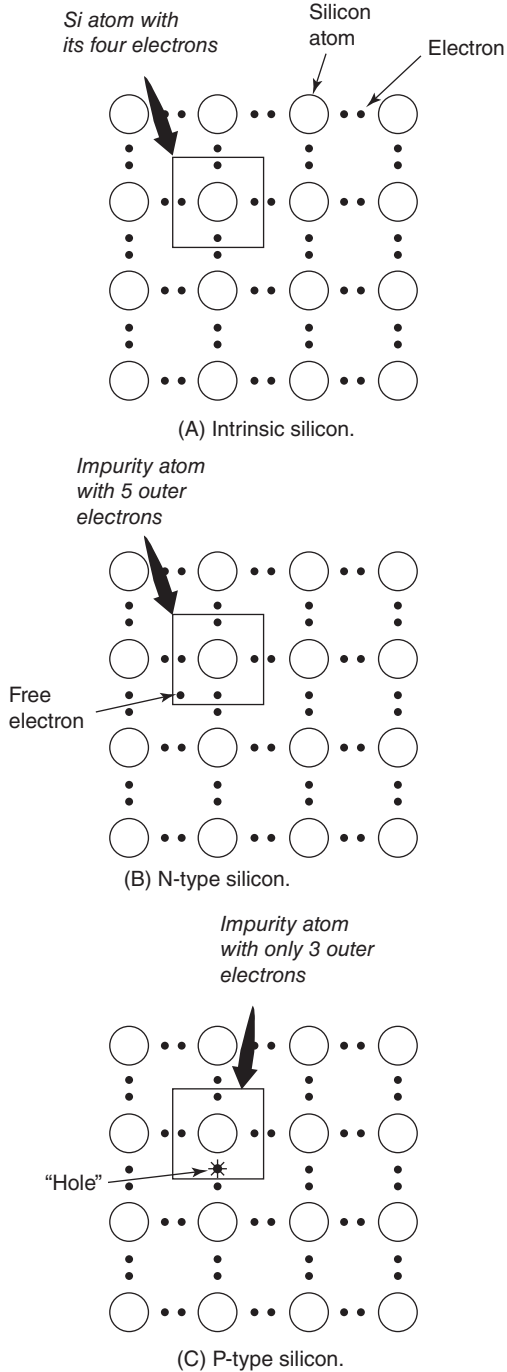


Figure 10-3. Bonding in pure and doped silicon crystals.

materials depends on the impurity doping. Undoped semiconductors are called *intrinsic* or *i*-type.

10.3.3 Diodes, Junctions, and Recombination

In the broader world of electronics, diodes are electronic devices with two terminals that transmit current easily (with low resistance) in one direction and poorly (with very high resistance) in the opposite direction. Diodes are called bipolar devices because their behavior differs depending on how you connect voltages to their two terminals. The first electronic diodes were vacuum tubes, now replaced by semiconductor diodes. The main electronic use of diodes is limiting current to flowing in only one direction through the device. In optoelectronics the main functions of diodes are to emit light (LEDs and diode lasers) and to convert light into electricity (solar cells and optical detectors). After describing semiconductor diodes in general, we will focus on their use as LEDs and diode lasers.

A semiconductor diode consists of regions of *p*- and *n*-type material with a thin *junction* layer between. Diodes usually are made by diffusing an excess of one type of dopant into semiconductor doped with the other type. For example, an electron acceptor such as aluminum or gallium can be diffused into a slab of *n*-type silicon, forming a top layer of *p*-type material in which the holes outnumber the electrons. A junction region where holes and electrons are nearly equal in number separates the *p*-type region from the *n*-type region. The junction typically is only 0.1 to 1 μm thick, and it is the place where the current flow changes and important things happen. Exactly what happens depends on the voltage applied across the junction and the diode's intended purpose.

With no voltage applied across the junction, charge carriers are distributed through the crystal in roughly the same way as impurities and not much happens. Near the junction, electrons from the *n*-type material can fall into holes in the *p*-type material, creating a zone where charge is distributed unevenly, but normally no net current flows. (Solar cells are exceptions, where light raises some valence electrons into the conduction band, generating electricity and causing a current to flow.)

No current also flows if the diode is *reverse biased* by applying a positive voltage to the *n* side of a junction and a negative voltage

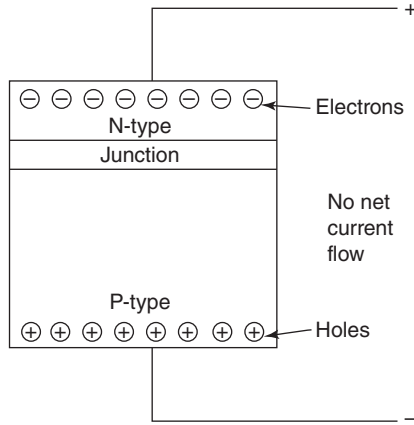


Figure 10-4. No current flows in a reverse-biased semiconductor junction.

to the p side. As shown in Figure 10-4, the positive electrode attracts electrons from the n -type material, and the negative electrode attracts holes from the p -type material. This draws carriers away from the junction, and virtually no current flows through the junction. Some current does leak through because the semiconductor's resistance is not infinite. Applying a high reverse voltage across the diode can pull valence electrons from atoms, causing electrical *breakdown* and current flow.

Current can flow through the diode when it is *forward biased* by applying a positive voltage to the p side and a negative voltage to the n side. The current starts to flow once the applied voltage exceeds the bandgap energy—typically 0.5 to a few electron volts. As shown in Figure 10-5, this attracts the p carriers to the n side of the device and vice versa, and they come together at the junction. There the electron and the hole *recombine* to form an electron-hole pair called an *exciton*, which releases energy equal to the gap between the conduction and valence bands when the electron falls into the hole. This produces a voltage drop at the junction equal to the bandgap energy.

The junction thus is the place where light is produced and much of the optical action happens in a diode laser, and for this reason it is often called the *active layer*.

The recombination energy normally winds up as heat in electronic applications. But Henry Round discovered over a century ago

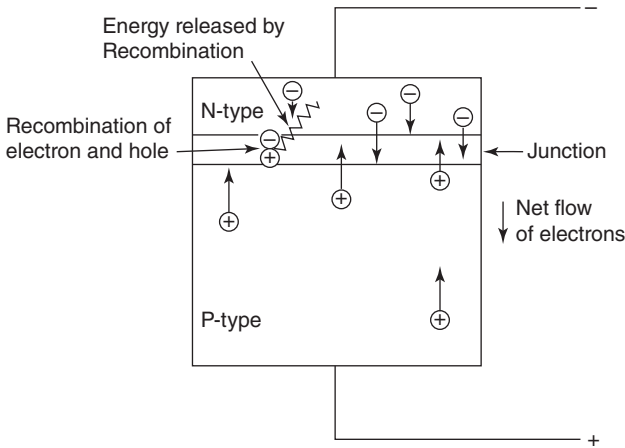


Figure 10-5. Forward-biasing causes electrons and holes to flow toward a junction in a semiconductor diode.

that some semiconductor diodes released some of the recombination energy in the form of light. That light became the basis of LEDs and diode lasers.

10.3.4 Indirect and Direct Bandgaps

Whether or not recombination produces light depends on the nature of the transition that releases the energy, and that depends on the semiconductor. Silicon emits essentially no light because an electron in its conduction band must interact with something else to change its momentum before it can drop into the valence band and release energy as light. This condition, called an *indirect bandgap*, delays the transition so long that something else releases the recombination energy before it can be released as light. Germanium and some other semiconductors such as gallium phosphide also have indirect bandgaps that make them poor light emitters.

Silicon LEDs have been demonstrated in the laboratory, but getting silicon to emit light requires special tricks to make its transitions behave differently. For example, nanostructures called quantum dots can confine electrons and holes on the scale of a few nanometers, making a conduction electron more likely to drop directly into a hole. But those silicon LEDs remain too inefficient for practical use.

What is needed to make efficient LEDs and diode lasers is a *direct bandgap* that lets electrons drop directly from the conduction band to the valence band without changing their momentum. The most important direct-bandgap semiconductors are compounds of elements from groups III and V in the periodic table, such as GaN, GaAs, and InP, known as *III–V compounds*. They are the type used for most LEDs and diode lasers.

The distinctions between direct and indirect bandgap compounds are not always sharp or obvious. GaAs has a direct bandgap and GaP has an indirect bandgap. Mix a little phosphorous with GaAs, and it remains a direct-bandgap semiconductor until the phosphorous level crosses a threshold.

Because this book is about lasers, we will focus on the compound semiconductors with direct bandgaps that allow light emission, making them useful in LEDs and semiconductor lasers.

10.3.5 Compound Semiconductors

Compound semiconductors are inorganic compounds containing two or more elements that form solids with the electrical properties of semiconductors. In principle, all types of semiconductor devices can be made from compound semiconductors but, in practice, silicon dominates the market for electronic devices. Compound semiconductors fill specialized niches and are particularly important for LEDs, lasers, and optoelectronic devices.

The most important compound semiconductors for laser applications are the III–V compounds, which contain equal amounts of atoms from group IIIa and group Va of the periodic table. The important elements are listed as follows:

Group IIIa

Aluminum (Al)
Gallium (Ga)
Indium (In)

Group Va

Nitrogen (N)
Phosphorus (P)
Arsenic (As)
Antimony (Sb)

The simplest of these materials are “binary” compounds containing two elements, such as GaAs, InP, and GaN. Each of these compounds has its own set of characteristics, including energy levels, bandgap, and atomic spacing in the crystalline lattice.

Other elements can be added to the compound as long as they maintain the balance of equal numbers of atoms from group III and group V, which have different valence. This can adjust properties of the semiconductor, particularly the size of the bandgap. For example, replacing some gallium in gallium arsenide with aluminum increases the bandgap energy in gallium–aluminum arsenide (GaAlAs). Such compounds containing three elements are called *ternary* and are written in the form $\text{Ga}_{1-x}\text{Al}_x\text{As}$, where x is a number between 0 and 1. This format indicates what we said above that the number of gallium atoms plus the number of aluminum atoms must equal the number of arsenic atoms.

Adding a fourth element to make a *quaternary* compound gives more flexibility and control over material properties. An example is indium–gallium arsenide–phosphide (InGaAsP), which is written $\text{In}_{1-x}\text{Ga}_x\text{As}_{1-y}\text{P}_y$, where both x and y are numbers between 0 and 1. In this case, the total number of indium and gallium atoms must equal the number of arsenic and phosphorus atoms. It is also possible to have three elements in one group and only one in the second, such as indium–gallium–aluminum phosphide (InGaAlP), written as $\text{In}_{1-x-y}\text{Ga}_x\text{Al}_y\text{P}$, where both x and y are numbers between 0 and 1 together adding to less than 1. In this case, the total number of indium, gallium, and aluminum atoms must equal the number of phosphorous atoms to form a semiconductor crystal.

Ternary and quaternary compounds are harder to fabricate into good crystals, but changing the composition can control the bandgap and lattice constant, which is invaluable in fabricating devices such as diode lasers. Practical concerns impose some limits on material characteristics. Ternary and quaternary compounds are hard to grow in bulk, so they normally are deposited on substrate wafers made of easier-to-grow binary compounds, such as GaN, GaAs, and InP. Successful deposition requires careful matching of the lattice spacing of the deposited compound with the substrate, limiting the blends that are easy to fabricate. This affects the composition and wavelengths of semiconductor lasers, as described later in this chapter.

Light emission has been demonstrated from two other families of inorganic compound semiconductors. One is silicon carbide (SiC), composed of two group IV elements with four valence electrons. The other family is called *II–VI compounds* because they contain elements with two and six valence electrons. They come mostly, but not entirely, from columns IIB and VI of the

periodic table. The most important elements in these compounds are as follows:

Column IIB	Column IVB	Column VI
Zinc	Tin	Oxygen
Cadmium	Lead	Sulfur
Mercury		Selenium
		Tellurium

The II–VI compounds fall into two broad groups. Compounds of zinc and cadmium with group VI elements have large bandgaps and can emit visible light. Compounds of lead and tin with group VI elements have small bandgaps and emit in the infrared. So far these technologies remain largely in the laboratory for LEDs and diode lasers, although some compounds are useful in making light-detecting diodes.

10.3.6 Organic Semiconductors

Organic LEDs (called OLEDs) are made from organic compounds that have electronic properties similar to inorganic semiconductors. They offer a number of potential advantages over inorganic semiconductors, including easy fabrication with inexpensive technology, compatibility with more substrate materials, and more freedom to tailor material properties including emitting wavelength. Recent progress in OLEDs has been impressive and is most evident in color displays.

Unfortunately, the progress has not translated into advances in organic diode lasers. The most dramatic claims of success were exposed as fraudulent and retracted in 2002. Research continues, but the difficulties are formidable.

10.4 COMPARING LED AND DIODE-LASER EMISSION

Henry J. Round was puzzled in 1907 when he saw light emission from an impure form of silicon carbide called carborundum. He knew the light arose from the junction between a metal conductor and the carborundum, which we now recognize as a semiconductor, but he did not understand what produced it. Figuring that out took other researchers decades.

Today, we know that electrons in the conduction band of a semiconductor emit *recombination radiation* when they drop into holes in the valence band. In Round's case, the junction was between a semiconductor and a conductor, similar to the point contact between semiconductor and metal in the first transistor. Modern LEDs and diode lasers emit light at a p - n junction between regions of p - and n -type semiconductor, which makes them more versatile, like the junction transistors that followed the point-contact type.

As you learned in Section 10.3.3, recombination releases an energy equal to the bandgap of the semiconductor, but the energy is released as light only if the material has a direct bandgap. Because the bandgap energy depends on the semiconductor composition, the emission wavelength also depends on composition. The wavelength also depends on the difference between LEDs and diode lasers.

10.4.1 Emission from LEDs

LEDs generate spontaneous emission across a range of wavelengths that depends on the bandgap energy of the compound in the junction layer. Table 10-1 lists the wavelengths of important LED materials. Ranges are given for compounds that are blends of varying composition; single wavelengths are given for specific compounds. Figure 10-6 shows the natural ranges of wavelengths from three

Table 10-1. Emission ranges of some important inorganic LED materials

Material/substrate	Peak wavelength or range (nm)	Status
AlGaIn/GaN	230–350	Developmental
InGaIn/GaN	360–525	Commercial
ZnTe/ZnSe	459 (blue)	Developmental
SiC	470 (blue)	Commercial
GaP	550–590 (green–yellow)	Commercial
GaAs _{0.15} P _{0.85}	589 (yellow)	Commercial
AlGaInP/GaAs	625–700 (red)	Commercial
GaAs _{0.35} P _{0.65} /GaAs	632 (red)	Commercial
GaAs _{0.6} P _{0.4} /GaAs	650 (red)	Commercial
GaAsP/GaAs	700 (red)	Commercial
Ga _{1-x} Al _x As/GaAs	650–900 (red and infrared)	Commercial
GaAs	910–1020 (infrared)	Commercial
InGaAsP/InP	600–1600 nm	Commercial




Figure 10-6. Bandwidths of typical commercial blue, green, and red LEDs, with curves drawn so all are about the same height. (Courtesy E. F. Schubert, www.lightemittingdiodes.org.)

typical visible LEDs. The typical range of a few tens of nanometers is narrow enough for the eye to see LED emission as a single color.

The wavelength range of an individual LED is much broader than that of an individual diode laser because the output is spontaneous emission. The amplification of light by stimulated emission in a diode laser selectively amplifies the strongest wavelength, narrowing laser output to a much narrower line.

LEDs radiate spontaneous emission in all directions from their junction layer, as shown in Figure 10-7. To get the most efficient output, the junction should be close to the surface to reduce absorption within the semiconductor material. Typically LEDs are packaged to concentrate light in one broad direction. Fabrication of a transparent lens on one output surface can concentrate emitted light, usually perpendicular to the surface. LEDs can be designed to emit in a variety of patterns to meet application requirements.

The lack of a resonant cavity means that LED output generally is less intense than a laser beam. However, the emitting surfaces of the white LEDs used for illumination can reach 10% of the solar intensity. The brightness comes from the concentration of light emitted through a small semiconductor chip to increase efficiency. The LEDs used for illumination are indium–gallium nitride devices emitting blue light at 465 nm, which is focused

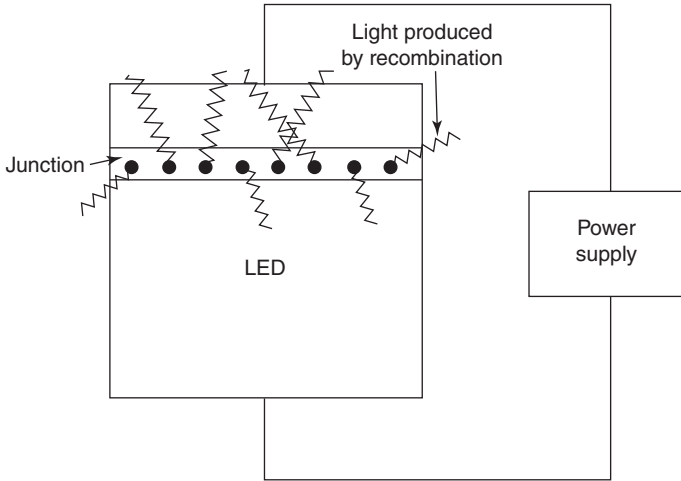


Figure 10-7. An LED emits light in all directions from its junction.

onto yellow-emitting phosphors that make the light look whitish to the eye. For indoor use, the light generally passes through frosted glass, diffusing the light over larger area at a lower intensity. But remove the frosted glass from an overhead fixture and the intensity of the bare surface is blinding and almost painful to the eye. Unfortunately, some municipal street lights do not use diffusers over the LED emitters, causing severe glare.

Most household LED lamps use blue LED emitters and a blend of phosphors that give colors similar to those of incandescent lamps. Industrial and workplace LEDs may use a bluer mix, more like a fluorescent lamps used in similar environments. Virtually all indoor LED lamps use frosted glass diffusers to reduce intensity to pleasant levels.

Color-tunable LEDs can be made by combining red, green and blue LEDs and adjusting the intensity of each color to get the desired color blend. The Philips Hue bulb is one example for household use. Buildings, bridges, and other architecture may be illuminated by tunable color lights at night to enhance their appearance. Colored LEDs also are used for indicator lights, traffic signals, and other applications requiring single colors.

Because this is a book about lasers, we will not go into detail on LEDs, but it is helpful to understand how they work and how they are used.

10.4.2 From LEDs to Diode Lasers

It was a surprisingly small step from the first efficient LEDs to the first semiconductor lasers. Both were demonstrated in 1962 in the same material, gallium arsenide, and the LED was almost overlooked in the race to make the first semiconductor laser.

The simplest diode lasers are structurally similar to LEDs. Both are small chips that generate light from recombination of electron–hole pairs at a forward-biased junction. Below the laser threshold, both generate spontaneous emission with an intensity that depends on the drive current. The big difference is that diode lasers have reflective surfaces on opposite edges of the wafer that create optical feedback by reflecting light back and forth in the plane of the pn junction, as shown in Figure 10-8.

The feedback has little impact at low drive current, when electron–hole pairs (excitons) release their energy by spontaneous emission in all directions, as in an LED. As the drive current increases, it produces more electron–hole pairs that emit light spontaneously, increasing the likelihood that a spontaneously emitted photon will encounter and stimulate emission from an exciton that has yet to release its extra energy. Once the drive current reaches a high enough level, it produces a population inversion between the exciton state (the upper laser level) and the atoms with the extra electron bound in the valence band (the lower laser level).

A population inversion is not enough to cross the threshold for laser oscillation. Feedback from the reflective edges must build up




Figure 10-8. Reflection from facets forms a laser cavity that oscillates in the junction plane of a diode laser.

enough for the amount of light bouncing back and forth between them to increase on each round trip. That means the drive current must increase well above the minimum needed for a population inversion.

In this simple diode-laser design shown in Figure 10-8, the best reflective edges are those perpendicular to the junction plane, so the reflected light will bounce back and forth along the length of the junction layer. Semiconductors have a high refractive index, so an uncoated solid-air interface reflects much of the stimulated emission back into the semiconductor to provide feedback. The large population inversion at high drive current makes gain high in semiconductor lasers, so cavities only a few hundred micrometers long can sustain oscillation. One facet may be coated to reflect all incident light, leaving the remaining stimulated emission to emerge from the other end.

Figure 10-9 shows how output power changes as the drive power increases. At low drive current, spontaneous emission dominates and the light output increases linearly at a slow rate. As the input current approaches the *threshold current*, the output current




Figure 10-9. Light output of a diode laser rises sharply above the laser threshold current.

risers faster, soon reaching a higher linear rate of increase, where the diode operates as a laser, converting a larger fraction of input electrical power into light, as shown by the steeper slope.

Threshold current is an important measure of diode laser performance. Light emission efficiency is low below threshold, but increases dramatically at and above threshold. The energy not turned into light becomes heat that must be dissipated in the laser. So does the fraction of above-threshold current that is not converted into light. The heat produced is not just wasted power; it also can degrade laser performance and shorten its lifetime. High currents also stress the laser by putting highly concentrated power through the junction; this is measured as *threshold current density* (threshold current divided by the junction area operating as a laser) rather than the total threshold current.

This should give you general idea of how a simple diode laser works. However, so far we have only described progress made during the early years of development, to a point where diode lasers could not operate for more than a fraction of a second at room temperature without self-destructing. The next sections will describe the details leading to today's high-power, high-efficiency and high-performance diode lasers.

10.5 CONFINING LIGHT AND CURRENT

The progress of diode lasers owes much to the development of structures to confine and control the flow of light and electrons in semiconductors. The better the confinement and control, the stronger are the interactions among light, current, and the semiconductor material. Those enhancements increase diode-laser efficiency and performance. Tighter control of light, current, and structures also enables the faster modulation needed for high-speed photonic switching and data transmission.

This section describes the important types of structures, how they are made, and how they work. It starts with a brief description of laser diode fabrication and then covers layered structures parallel to the junction and structures fabricated in the device plane. Then it describes the two primary classes of laser diode emitters, those emitting from their edges in the junction plane and those emitting vertically, perpendicular to the junction.

10.5.1 Photolithography and Semiconductor Fabrication

Semiconductor electronics are fabricated by depositing a series of layers of semiconductor onto a semiconductor substrate. The basic processes are common to all semiconductor electronic devices, but details vary considerably depending on the materials and device types.

For laser diodes, the substrates are wafers of binary semiconductors that can be fabricated in bulk, usually GaAs, GaN, and InP. The thin layers of compounds with slightly different compositions are deposited one at a time using a process called *epitaxy*. Atomic spacing in each layer, called the *lattice constant*, must be close to that in layers above and below to avoid producing flaws in the crystal. Typically this means making layers from a ternary or quaternary compound with lattice constant close to that of the substrate.

After each layer is deposited, a pattern can be etched into it using a process called *photolithography*. First the existing surface is coated with an emulsion sensitive to light. Then light from an external source is focused through a mask to write a pattern on the emulsion. Washing away the exposed areas of the emulsion uncovers the surface below. Then another chemical is spread across the whole layer to etch away the areas no longer covered by the exposed emulsion, leaving a pattern in the surface. Then the surface is cleaned, leaving the etched semiconductor layer, and another semiconductor layer is applied to the surface. The process is repeated again and again to complete the device. Insulating layers are added to control current flow. Semiconducting layers have impurities added to produce *n* or *p* type material or to change the material's refractive index and bandgap.

These processes are the key to the extraordinary improvements in semiconductor diode lasers over the past several decades. Let us examine what they can accomplish, starting with layering, then turning to in-plane structures and designs of laser cavities.

10.5.2 Layers and Confinement in Diode Lasers

Refinements in diode-laser structure beyond the simple design of Figure 10-8 have increased efficiency and operating lifetime by confining both the drive current and the light energy to small regions of the semiconductor, particularly the active layer. The

better the confinement, the lower the threshold current, the longer the lifetime, and the better the laser operation. We will start with a central concept that earned Herbert Kroemer of the USA and Zhores Alferov of Russia the 2000 Nobel Prize in physics, the semiconductor heterostructure.

10.5.2.1 Heterostructures and Layering of Devices

The first diode lasers consisted essentially of two layers of the same compound, typically GaAs with a junction between them. One was doped with electron donors to make an n -type semiconductor, while the other was doped with an impurity that added holes to the valence band to create a p -type material. These devices were called *homostructure* or *homojunction* lasers because all their layers were made from the same compound semiconductor. They emitted light when electrons recombined with holes at the junction, but the homojunctions did little to confine the electrons at the junction, so many drifted away without recombining to emit light.

Homojunction diodes worked, but they did not work very well and required cooling even to produce pulses of light because they were not able to confine the locations of either current carriers or light very well. In 1963, just a year after the first diode lasers were operated, Kroemer and Alferov independently suggested that diode lasers would work better if they contained two layers of different composition, which they called a *heterojunction* or *heterostructure* to highlight the difference.

They reasoned that a junction between layers with different bandgaps would trap conduction electrons in the material with lower bandgap because they lacked the energy needed to move into the layer with a larger bandgap, as shown in Figure 10-10. For example, substituting aluminum for some gallium in a gallium-arsenide diode laser increases the bandgap energy, so an electron in the GaAs layer would not be able to cross into an AlGaAs layer with a higher bandgap energy. Confining electrons in low-bandgap material near the junction would enhance the likelihood that they would recombine with holes and generate light, improving diode-laser efficiency.

Figure 10-11 compares electron confinement among diode lasers with homojunctions, one heterojunction, or a pair of heterojunctions above and below the junction layers. A *single heterojunction* between two materials improves electron confinement enough




Figure 10-10. Energy levels at a heterojunction allow conduction electrons to drop to the low-bandgap side, but not to return to the high-bandgap side.

for a diode laser to operate in pulsed mode at room temperature. Note how the electrons and the light they generate by recombination occupy a smaller area in the single-heterojunction laser than in the homojunction laser.

If one heterojunction is good, two should better confine light and electrons, as you can see at the bottom of Figure 10-11. In this case, the active layer is GaAs, which has a lower bandgap energy than the GaAlAs layers above and below it. This design is called a *double-heterojunction* or *double-heterostructure* laser, and it was the first diode laser capable of continuous-wave operation at room temperature. Modifications to the double-heterostructure design are the basis of today's diode-laser industry.

Refractive index also varies with material composition, and a difference in refractive index at a material boundary can also confine light. As you learned in Section 2.5, total internal reflection confines light inside the low-index core of an optical fiber if the light is directed along the length of the fiber. So if light is directed along the plane of the junction layer of a diode laser, total internal reflection could confine it in the junction layer if the layers above and below it had higher refractive index. However, higher-bandgap materials which confine electrons often have lower refractive indexes, so heterojunctions which confine electrons do not necessarily also confine light by total internal reflection.

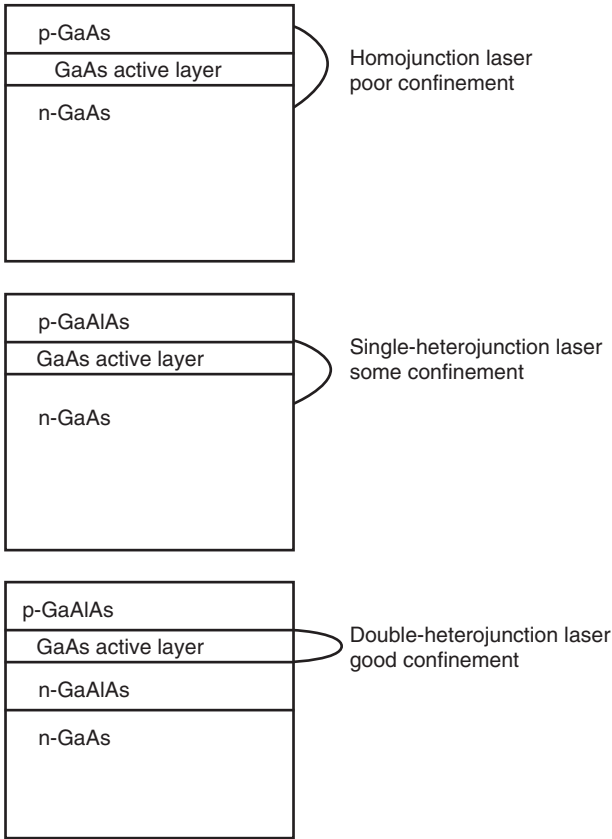


Figure 10-11. Confinement in homojunction, single-heterojunction, and double-heterojunction lasers. (Thickness of active layer is exaggerated to fit label.)

10.5.2.2 Quantum Wells, Quantum Wires, and Quantum Dots

The active or junction layer in a double-heterojunction laser typically is 0.1 to 1 μm thick, so it does not display quantum effects. However, reducing the thickness to 50 nanometers does induce quantum behavior in electrons. Electrons, like photons, have a dual existence as waves and particles, and 50 nm is roughly the “wavelength” of an electron in a semiconductor.

Earlier you learned that double heterostructures can confine electrons so they can move in only two dimensions. When the double heterostructure is made thinner than 50 nanometers, its internal energy levels change, so the layer becomes a *quantum well*, which confines electrons more tightly than bulk semiconductor

layers. Stacking two or more quantum wells inside the active layer in a multiple quantum-well laser increases the quantum-well volume in which stimulated emission can occur. Photons have wavelengths of hundreds of nanometers in a semiconductor, so they are confined on a larger scale than electrons.

As Figure 10-12 shows, confinement in semiconductors can go even smaller. *Quantum wires* can be made to confine electrons within a linear structure thinner than an electron wavelength, so they can only move in one dimension, along the quantum wire. However, *quantum dots*, which confine electrons in all three dimensions on a scale smaller than the electron wavelength, have proved more interesting.

Quantum dots can be made in sizes from 2 to 50 nm and have become a hot area in nanotechnology. They can be embedded in the junction layers of diode lasers or LEDs and excited by electric current. The tight confinement of electrons improves performance, reducing temperature sensitivity and allowing operation at higher temperatures and lower threshold currents. Their size controls their wavelength, so control of size can give precise colors both in lasers and LEDs. You can think of quantum dots as a tool to enhance performance of light emitters, some of which now are available commercially.

A quantum dot is much smaller than the wavelength, so it cannot be a resonant cavity. The resonant cavity in a quantum-dot laser is a region in the junction layer that contains a number of quantum dots and is bounded by cavity mirrors. Nanoscale lasers are in the research realm and covered in Section 14.8.

10.5.2.3 Lattice Matching and Strain

A key to success in making diode lasers is the ability to deposit layers with atomic spacing that matches that of the underlying material. This poses a challenge because the atomic spacing depends on the semiconductor composition and that composition must differ between layers for a diode laser to function properly. Moreover, ternary and quaternary compounds required for functional layers must be deposited on substrates made of binary compounds with different composition. Failure to match lattice spacing can lead to defects in the crystal, which degrade its optical, electronic, or mechanical properties.

How much the lattice constant varies with composition depends on the semiconductor. Adding modest amounts of

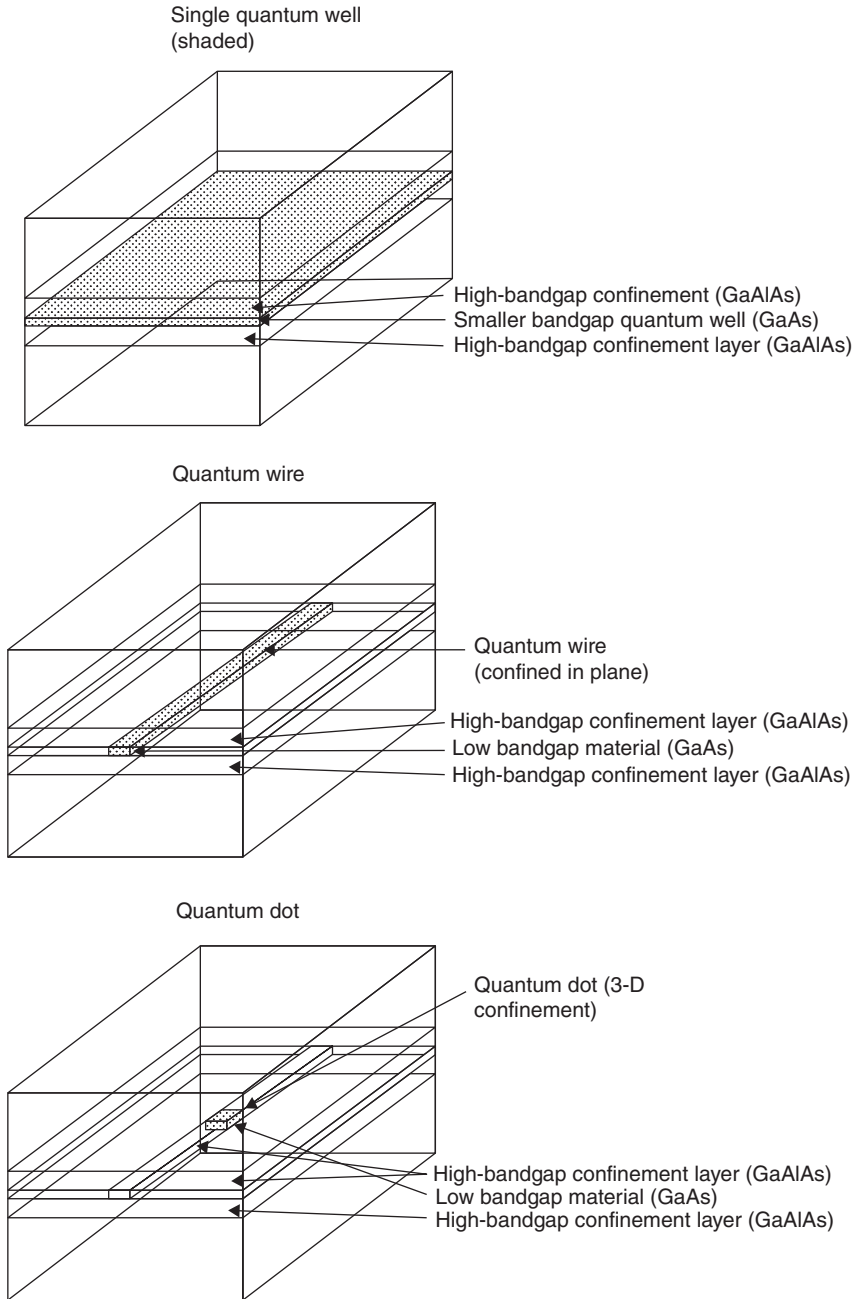


Figure 10-12. Confinement in quantum wells, quantum wires and quantum dots is down to the scale of electron wavelength, shown in comparison with bulk material.

aluminum to replace gallium atoms in gallium arsenide produces little change in the lattice constant, so heterostructures were first used in GaAlAs diodes deposited on GaAs substrates. Lattice constants of indium–phosphide compounds vary more widely with composition, so developers turned to quaternary compounds, adding both gallium and arsenic to InP to produce a blend of InGaAsP with a lattice constant to match InP substrates.

Compound semiconductors can tolerate small differences in atomic spacing, as long as the strain produced by the difference is too small to damage the crystal. That makes it possible to accommodate larger differences between lattice constant by depositing buffer layers with intermediate spacing, which distribute the strain through more of the crystal. These buffer zones are called *strained-layer superlattices* and broaden the range of materials usable in diode lasers. They were important in developing high-power diode lasers of indium–gallium arsenide in the 900–1000 nm range.

10.5.3 Confinement in the Junction Plane

The first diode lasers produced recombination energy and light across the entire junction plane, with the edges of the semiconductor chip reflecting some light back into the laser medium and emitting other light through the edge. Concentrating the drive current and stimulated emission to a narrow stripe in the junction plane, shown in Figure 10-13, was an important step in development of these *edge-emitting* diodes, which were the first practical diode lasers and are still widely used. That design also added a totally reflective mirror on one edge of the laser, so all the output emerged




Figure 10-13. Stripe-geometry diode laser.

from one side. Confining drive current so that it passed through a narrow stripe in the junction layer increased power density and thus improved power-conversion efficiency and laser performance. It was the first in a series of improvements in control of current and light in edge-emitting diode lasers.


10.5.3.1 Gain- and Index-Guided Lasers

Stripe width can be limited in two ways: by restricting the flow of current to a narrow stripe or by fabricating stripes of material with different refractive indexes in the junction plane. Both can be done in the same diode laser. Typical stripes are a few micrometers wide and extend across the length of the chip, a few hundred micrometers to a millimeter.


In a *gain-guided laser* such as shown in Figure 10-14A, insulating regions at the top of the laser chip block current from flowing to either side in a complex double-heterojunction laser. The only path for the current is through a narrow stripe at the middle, which runs the length of the chip between the two cavity mirrors. Current carriers recombine to produce a population inversion only in the narrow stripe through which the current flows, so only that zone has gain. Because there is no gain at the sides, those regions do not emit light, even though no physical boundary separates the stripe from the rest of the active layer, but some light may spread into that region.

Index-guided lasers add another level of confinement by surrounding the stripe in the active layer with a material of a lower refractive index, so total internal reflection confines diode emission in the stripe. As shown in Figure 10-14B, the resulting structure can be complex. In this case, the laser has been etched during fabrication to leave only a narrow stripe or mesa containing the GaAs junction layer that runs the length of the chip. Then *n*-type GaAlAs was deposited on either side of the stripe and covered with an insulator before depositing a metal contact. The insulator confines current flow through the mesa; the boundary with the GaAlAs confines light in the GaAs active layer the same way a double heterojunction confines light in the central active layer. The design shown in Figure 10-14B is an example of a *buried heterostructure* laser, in which the light-emitting stripes are buried entirely by other materials, except at the light-emitting facets.

Index-guided lasers are used for most diode-laser applications because they confine light better, producing better beam quality. Gain-guided lasers are simpler to make, and their poorer confinement of light can be an advantage in generating high powers because



(A) Gain-guided laser.



(B) Index-guided laser.

Figure 10-14. End views of gain-guided and index-guided lasers.

spreading light over a larger area reduces the chance of optical damage to the emitting surface.

10.5.3.2 Broad-Area Lasers for High Power

The power available from narrow-stripe diode lasers is limited by the generation of waste heat and by the damage that extremely concentrated light power can cause to optical surfaces. Diode

lasers decrease in efficiency and age more rapidly as temperatures increase, so waste heat must be removed from lasers operating at high power. However, even active cooling cannot eliminate the danger of optical damage as the power per unit area becomes very high in the small emitting area of a diode laser.

These limitations can be avoided by making the diode-laser stripe 100 or 200 μm wide, much larger than in low-power stripe-geometry diodes. This produces gain in a larger volume and spreads the emission over a wider area on the edge of the active layer, reducing power density at the surface. Continuous output of broad-area lasers can exceed 20 watts.

Their main advantage is delivering high power from a single emitting aperture. The beam quality is not as good as that from a narrow-stripe laser, although the wide emitting aperture means that the beam does not diverge as rapidly. The high power from a single emitting area is an important advantage for diode pumping of solid-state and fiber lasers.

10.5.3.3 Multistripe Laser Bars and Stacks

Laser power from a single laser stripe is inherently limited, even from a broad-area laser. To obtain higher powers, many parallel laser stripes can be fabricated on the same monolithic semiconductor substrate and their outputs coupled together, as shown in Figure 10-15. The combined output power of such a diode laser *bar* or linear array can reach kilowatt class, much higher than that from any single diode laser, although the beam quality is inevitably lower.




Figure 10-15. Monolithic diode-laser array has many parallel stripes.




Figure 10-16. Output from stack of five planar monolithic arrays is collected and focused by an optical system for laser pumping or other application. (Courtesy of Power Photonic)

Output power can be further increased by stacking together many multistripe bars or wafers into a *stack* to produce a two-dimensional array of laser emitters. Figure 10-16 shows how several bars are stacked on top of each other, with light emerging from the right. Space between the bars allows coolant to flow between the bars to remove waste heat. The optics to the right of the bars focuses the light from individual strips and from the bars into a light-guiding fiber or other optics at lower right. Laser stacks can emit continuous powers in the kilowatts, and repetitively pulsed arrays can generate even higher average powers. The most efficient diode arrays convert up to 70% of the input electrical power into light, but that still leaves 430 watts of heat for every kilowatt of laser output in the small volume of the array, so some cooling is needed.

Multistripe lasers are used to generate raw optical power. Some are used in industrial applications that do not require high beam quality because of their high efficiency and high power. Others are used for pumping solid-state or fiber lasers, as described in Chapters 8 and 9.

10.6 EDGE-EMITTING DIODE LASERS

The optical cavities of edge-emitting diode lasers are in the plane of the junction layer and emit light from the edge of the chip. You have

already seen examples in Figures 10-8 and 10-13. Edge emitters are the oldest diode lasers, and many variations have been developed over the years, some of which you read about in Section 10.5.3. Their gain per unit length is high, and typical stripe-geometry lasers can extract laser light from an area a few micrometers wide and a few hundred micrometers long in the active layer. However, edge emitters have poor beam quality and high beam divergence, and the chips have to be diced to expose their emitting edges for testing. Several varieties of edge emitters are in use.

Surface-emitting diode lasers use the same semiconductor materials and emit at the same wavelengths, but otherwise are quite different devices, which you will learn about in Section 10.7.

10.6.1 Fabry–Perot Lasers

The simplest diode-laser resonator is the Fabry–Perot cavity introduced in Chapter 3, a linear cavity with flat mirrors on the two ends. In a semiconductor laser, the linear cavity is the active stripe in the junction layer, and the flat mirrors are the facets on the edges of the chip. Typical diode-laser resonators are only 300 to 500 micrometers long, but diode lasers have high gain, so milliwatt powers are easy and much higher powers are possible.

The cavity mirrors are formed by the boundary between the semiconductor crystal and the air. The refractive index of GaAs is 3.34 at 780 nm so according to Equation 5-4 roughly 30% of the light hitting the boundary is reflected back into the semiconductor, providing adequate feedback for laser oscillation, with the remaining 70% emitted.

A diode laser with such a simple cavity can emit light from both the facets. In practice, one facet is coated to reflect all or most of the incident light and the beam emerges from the other facet, as you saw in Figure 10-8. Often lasers are designed with the rear facet coated to let a little light emerge so it can be monitored to control laser output in devices such as fiber-optic transmitters.

Narrow-stripe lasers typically emit from an area less than a micrometer high and only a few micrometers wide, producing a single transverse mode, with a central intensity peak in the emerging beam. Wide-area lasers oscillate emit from a much wider stripe, up to about 100 μm , in multiple transverse modes. You will learn in Section 10.8.1 that these odd thin emitting areas strongly affect beam quality.




Figure 10-17. Mode hopping in a Fabry–Perot diode laser near 1550 nm.

The resonators of Fabry–Perot diode lasers are short compared with those of gas lasers, so cavity modes are spaced more widely in diode lasers—about 0.2 nm for an 800-nm GaAlAs laser and 0.7 nm for a 1550-nm InGaAsP laser. That means that Fabry–Perot lasers typically emit most of their light at one wavelength at any one time, but may have side bands. However, the gain curve of diode lasers is broad and the refractive index varies with temperature, which itself is a function of drive current. That means that a Fabry–Perot laser might *mode hop*—shift to another wavelength—when its intensity is modulated, as shown in Figure 10-17. Mode hopping is undesirable because it can generate noise.

Fabry–Perot lasers are widely used where low cost is important and precise wavelength is not critical, such as in CD players, laser pointers, and short-distance, low-speed communication links.

10.6.2 Distributed-Feedback Lasers

Applications such as high-speed data transmission in fiber optics require limiting laser emission to a narrower range of wavelengths than possible with a Fabry–Perot cavity. This requires adding optics to limit the range of wavelengths emitted by the laser. The leading approach is to fabricate a diffraction grating in the base of the active layer to make a *distributed-feedback* or *DFB* laser, shown in Figure 10-18A. The regularly spaced grooves scatter light back into the active layer at a narrow range of wavelengths, so only those wavelengths can be amplified in the laser cavity.

The distance D between the grating lines selects the oscillating wavelength λ according to a formula that also depends on refractive




Figure 10-18. (A) Distributed-feedback laser and (B) distributed Bragg-reflection laser.

index n and an integer m (in practice, 1 or 2) that denotes how light is being scattered by the grating:

$$D = \frac{m\lambda}{2n} \quad (10-2)$$

Plugging in the numbers for a 1550-nm DFB laser made of InGaAsP ($n=3.4$), we find that grating spacing D is 228 nm for $m=1$ and 456 nm for $m=2$. Although oscillation wavelength shifts slightly because refractive index depends on temperature, DFB resonators maintain the laser in a stable single longitudinal mode, limiting the range of emitted wavelengths and preventing mode hopping. Both are critical for high-speed fiber-optic transmission.

An important variation on the DFB laser is placing the grating in a part of the active layer where there is no laser gain, as shown in Figure 10-18B. This has the same effect of limiting amplification

to a narrow range of wavelengths, but the physics differ in detail. The result is called a *distributed Bragg reflection* or DBR laser. The physics are similar to those of distributed Bragg gratings used as resonator mirrors in fiber lasers.

10.6.3 External-Cavity and Tunable Lasers

Diode-laser emission also can be limited to a narrow linewidth by placing the diode in an external cavity with suitable tuning optics. For example, one facet may be coated to transmit nearly all light emerging from the active layer to external optics that select the wavelength, as shown in schematically Figure 10-19. Actual external cavities are longer relative to the small diode laser, limiting output to a narrow range of wavelengths. Adjusting the optics tunes the wavelength by selecting which wavelengths are reflected back into the active layer of the diode laser, an important advantage of external-cavity diode lasers.

Another advantage of external-cavity diode lasers is the ability to narrow their free-running bandwidth from the hundreds of gigahertz for conventional diodes to as little as a few kilohertz. This concentration of output in a narrow band can be combined with amplification to allow harmonic generation, which can be difficult with conventional diode lasers.

Other types of tunable diode lasers have also been developed, largely for fiber-optic systems for which it is desirable to be able to switch emission wavelengths. One example is the *sampled grating, distributed Bragg reflector (SG-DBR)* laser. Like the example in Figure 10-18B, it has gratings on both sides of the active region of the




Figure 10-19. External-cavity semiconductor laser.

laser, but in this case the spacings are different on the two sides. Slight adjustments of the two gratings tune them relative to each other, producing a disproportionately large change in the wavelength at which the laser oscillates.

10.6.4 Semiconductor Optical Amplifiers

A semiconductor diode can serve as an optical amplifier as well as a laser. Like a diode laser, a semiconductor optical amplifier is powered by a current flowing through the device, producing recombination at the junction layer. The difference is that the device amplifies light from an external source, rather than oscillating on its own. This requires coating the diode facets to suppress surface reflection so light makes a single pass through the diode rather than oscillating between reflective surfaces. Semiconductor optical amplifiers are less expensive and available at more wavelengths than the high-performance fiber amplifiers described in Section 9.5.1. However, they are subject to more noise and other problems.

10.7 SURFACE-EMITTING DIODE LASERS

Diode lasers can emit from their surfaces, which offer some important advantages over edge emitters.

Edge emitters must be diced to expose their edges and then packaged before they can be tested. Surface emitters can be tested on the wafer before it is diced, and are easier to package—a key advantage because packaging is the biggest cost in diode-laser production. Surface emitters also can be integrated more easily on a single substrate.

The emitting areas on edge emitters are thin and wide, producing rapidly diverging beams that are hard to focus. Surface emitters can be designed with larger, round emitting areas, which produce better quality beams with much lower divergence. The larger emitting areas can also handle more power, although the gain volume is much smaller in vertical cavity lasers because the active layer is thin.

Two designs have emerged for surface emitters: one with a vertical cavity perpendicular to the active layer and another that resonates partly in the active layer. Each deserves a separate description.

10.7.1 Vertical-Cavity Surface-Emitting Lasers

Vertical-cavity surface-emitting lasers (*VCSELs*) get their name because their resonant cavities are vertical, perpendicular to the active layer, as shown in Figure 10-20. Mirror layers are fabricated above and below the junction layer, with the beam usually emerging through the substrate.

VCSELs differ greatly from edge emitters. Instead of oscillating along the long dimension of a long, narrow, and thin active layer, VCSELs oscillate perpendicular to the surface of a thin active layer. VCSEL cavities also are shorter. These structural differences make VCSELs behave rather differently than edge emitters.

Overall gain within a VCSEL cavity is low because light oscillating between the top and bottom mirrors passes through only a thin slice of active layer. Although the gain per unit length is high, the active layer is so thin from top to bottom that the total gain in a round trip of the VCSEL cavity is small. To sustain oscillation, resonator mirrors on top and bottom of the VCSEL must reflect virtually all the stimulated emission back into the laser cavity, like a helium–neon laser.

The high-reflectivity mirrors needed for the laser cavity are fabricated in the semiconductor itself, by depositing many




Figure 10-20. Vertical-cavity semiconductor laser.

alternating thin layers of two different compositions of semiconductor which have slightly different refractive indexes. This structure forms a multilayer interference coating, described in Section 5.3.3, which can strongly reflect a particular wavelength. The reflector on the substrate side transmits a small fraction of the cavity light; the reflector above the active layer normally reflects all the light back into the cavity.

Single VCSELs with circular output apertures of a few micrometers are mass-produced for applications using a milliwatt or less, a market they dominate. Their design gives them a low threshold power, making them more efficient than edge emitters at similar output. Their low drive current enhances their lifetimes. Their circular emitting areas give them a circular beam with better quality and lower divergence than edge emitters. Those attractions led to their highly publicized use in Apple's iPhone X for face recognition and camera focusing.

The short length of the VCSEL cavity has another important consequence. Recall from Chapter 4 that an integral number N of wavelengths λ fit into a laser cavity with length L and refractive index n according to the formula:

$$2nL = N\lambda \quad (10-3)$$

The shorter is the cavity, the larger is the difference between resonant wavelengths. That means that VCSELs are much less likely to hop to modes oscillating at different wavelengths than edge emitters. This improves their performance and direct modulation at tens of gigabits per second by varying drive current, important for fiber-optic data transmission.

Single VCSELs can be made up to several hundred micrometers wide, enough to generate watt-class power. Monolithic VCSEL arrays can emit much higher powers. VCSEL efficiency decreases with aperture, so arrays work best if made of many smaller VCSELs. Multiple VCSEL arrays can combine to generate over a kilowatt.

In principle, VCSELs can be made from any direct-bandgap III–V semiconductor using standard semiconductor manufacturing process to deposit the mirror layers as well as the p – n junction structure. Gallium arsenide VCSELs were developed first because the refractive index of GaAlAs varies considerably with aluminum content, providing the refractive-index contrast needed for the multilayer mirrors. That led to early use of low-power 850-nm VCSELs

in short fiber-optic communication systems, which require little power where limited power was not a problem. Now InGaAsP VCSELs are available for the 1300- and 1550-nm fiber-optic windows. InGaN VCSELs also have been developed.

VCSELs are fabricated in arrays of many devices on a single wafer, but most of them are packaged individually rather than used in arrays. Because VCSELs emit from the wafer surface, they can be tested before the wafer is diced into many individual devices and packaged. Edge emitters cannot be tested until the wafer is scribed and diced into individual components, raising the costs of testing and fabrication. This leads to lower testing and packaging costs for VCSELs, and lower prices.

10.7.2 Vertical External-Cavity Surface-Emitting Lasers

Two variations of VCSELs have been developed with external optical cavities in which one mirror is physically separated from the semiconductor structure. Rather than having multilayer cavity mirrors on the top and bottom of the semiconductor VCSEL element, these lasers have one mirror physically separate from the semiconductor structure. These two types differ in important ways from VCSELs and from each other and are important to distinguish.

One type is an electrically powered device called a *tunable VCSEL* with the external mirror separated a small distance from the semiconductor structure so it can be moved to tune the wavelength resonant in the vertical laser cavity. The best developed approach had the movable mirror suspended above the semiconductor element and changed the resonant wavelength by moving the mirror up and down to change the length of the cavity.

The second type is an optically powered device that you met earlier in Section 8.5 as an *OPSL*. Because it is optically pumped rather than powered electrically by drive current, it works more like the solid-state lasers of Chapter 8 than like the diode lasers of this chapter, so it is covered there. However, the light comes from a thin multilayer semiconductor structure based on the VCSEL, and some of its developers call it a *vertical external-cavity surface-emitting laser*. The different names came from different companies, so there is a bit of commercial conflict behind the confusion. They are described in detail in Section 8.5.

10.7.3 Horizontal Cavity Surface Emitters

Another interesting technology tries to combine the best features of edge and surface emitters in a diode laser with a resonant cavity that is bent 90 degrees. Most of the length of the cavity is in the active layer, where laser light bounces between a total reflector perpendicular to the junction and a mirror or grating that directs light upward at 90 degrees to an output mirror on the surface. The idea is to collect as much energy as possible from the relatively large volume of the junction layer and then direct up through the surface to get the beam quality of a surface-emitting diode. Several groups have studied these lasers, but they are not widely used.

10.8 OPTICAL PROPERTIES OF DIODE LASERS

The optical properties of diode lasers depend on their semiconductor composition, their cavity structures, and whether they are single emitters or part of an array. Edge emitters and VCSELs share some features, but differ in other ways because of their distinct cavity structures.

10.8.1 Beam Shape and Divergence

As you learned in Section 4.3.5, a beam emitted through a small aperture D has a relatively large divergence angle θ , because divergence is proportional to the wavelength λ divided by the aperture D . That is, the smaller the aperture is relative to the wavelength, the more rapidly the beam diverges. For circular apertures, many wavelengths across, the divergence angle is

$$\theta = \frac{K\lambda}{D} \quad (10-4)$$

where K is a constant near 1 and depends on the beam profile.

Stripe-geometry, edge-emitting diode lasers emit from the edge of the active layer, which is a fraction of a micrometer high and several micrometers wide. This shape is quite different from the round emitting areas of most gas and solid-state lasers, so the formula for beam divergence is only an approximation. Nonetheless, measurements confirm that the beam spreads much more rapidly in




Figure 10-21. Beam divergence from an edge-emitting laser.

the vertical direction, where the emitting aperture is narrow, than in the horizontal direction, where the emitting aperture is wider, as shown in Figure 10-21. Typical values for beam divergence are 10 degrees (0.17 radian) in the direction parallel to the active layer and 40 degrees (0.70 radian) in the direction perpendicular to the active layer. These spreading angles are larger than the beam from a good flashlight, making the beam quite different from the tightly focused beam of a helium–neon laser or a solid-state laser.

Fortunately, external optics can correct for this broad beam divergence. A cylindrical lens, which focuses light in one direction but not in the perpendicular direction, can make the beam circular in shape. Collimating lenses can focus the rapidly diverging beam from an edge emitter so it looks as narrow as the beam from a helium–neon laser. Those lenses are hidden inside standard red or blue diode-laser pointers, so they produce tightly focused spots on the screen or wall. You can see how these optics look in Figure 10-16, which shows the optics used for focus is a diode stack.

VCSELs are a different matter because they emit from circular apertures generally at least 5 μm across, so their beams diverge symmetrically. Typical divergence for a small VCSEL is about 10 degrees (0.17 radian). The diffraction limited for a 50- μm VCSEL is around 1 degree (0.017 radian). The addition of collimating optics can reduce beam divergence, but cylindrical optics are not needed because VCSEL beams are circular.

10.8.2 Direct Modulation, Modes, and Bandwidth

An important feature of diode lasers is that they can be directly modulated by controlling the drive current. The potential bandwidth depends on the modulating electronics and the size of the diode lasers and can reach tens of gigabits per second in small diodes.

Output modes depend on the diode design. Broad-area diode lasers and arrays of many diode-laser stripes oscillate in many modes. Narrow-stripe diode lasers typically oscillate in only one transverse mode and may be called “single-mode lasers,” but Fabry–Perot diodes can oscillate in more than one longitudinal mode and hop between modes when operating conditions change. The modes are 0.2 to 0.7 nm apart, which is fine for many applications but unsuitable for most telecommunication applications.

Edge-emitting DFB, DBR, and external-cavity lasers limit laser oscillation to a single stable longitudinal mode, giving them bandwidth narrow enough for moderate data rates. However, they may require external modulation for higher speeds.

The short cavity of VCSELs leads to widely separated longitudinal modes and allows direct modulation of small-area VCSELs at speeds to tens of gigabits. Large-aperture VCSELs oscillate in multiple transverse modes.

10.8.3 Power and Efficiency

Often the most important optical property of diode lasers is their high efficiency in converting electrical energy into light. VCSELs can generate milliwatts of light with a low threshold. The most efficient high-power diode lasers can convert more than 70% of the input electrical energy into laser output.

The high efficiency of diode lasers makes it practical to use them to pump other lasers optically. Diode pumping has opened up new applications for solid-state lasers and made fiber lasers practical and efficient.

10.9 DIODE-LASER MATERIALS AND WAVELENGTHS

So far, this chapter has focused mostly on the properties of optoelectronic semiconductors in general and the structures used in diode

lasers. The operation of diode lasers also depends strongly on the composition of the semiconductor. Now that you have learned how diode lasers work, it is time to look in more detail at the semiconductors used and how they affect diode-laser performance.

10.9.1 Bandgaps and Compositions

The bandgap is one of the most important parameters in determining how a semiconductor material performs in diode lasers. This, in turn, depends on the composition of the semiconductor, which also determines the atomic spacing or lattice constant, described in the next section.

Table 10-2 lists important types of semiconductor lasers and the wavelengths at which they operate as lasers. All the diode lasers listed are made from compound semiconductors, most of them blends of elements in groups III and V of the periodic table (sometimes called groups 13 and 15). A few other diodes are compounds of groups II and VI (sometimes called groups 12 and 16). The exact wavelength depends on the mixture of elements used. Except where noted, all the lasers in the table are diodes.

Table 10-2. Major semiconductor laser materials and wavelengths. All are diodes except where noted

Material	Wavelength range	Comments
AlGaIn	350–400 nm	Developmental at shorter wavelengths
GaInN	375–525 nm	Commercial; UV to green
ZnSSe	447–480 nm	Laboratory
ZnCdSe	490–525 nm	Laboratory
AlGaInP/GaAs	620–680 nm	Commercial; red
Ga _{0.5} In _{0.5} P/GaAs	670–680 nm	Commercial; red
GaAlAs/GaAs	750–900 nm	Commercial; near-IR
GaAs/GaAs	904 nm	Commercial
InGaAs/GaAs	915–1050 nm	Strained layer; commercial
InGaAsP/InP	1100–1650 nm	Commercial
InGaAsSb	1.8–4.2 μm	Some commercial
PbCdS	2.7–4.2 μm	Requires cooling
Interband cascade	3–4 μm	Not a diode
Quantum cascade	4–50 μm	Not a diode
PbSSe	4.2–8 μm	Requires cooling
PbSnTe	6.5–30 μm	Requires cooling
PbSnSe	8–30 μm	Requires cooling

Electrons and holes recombine in the junction layer, where electrons are trapped in zones with a lower bandgap energy than surrounding material. These low-bandgap zones are called double heterostructures or quantum wells. Bandgap energy is often given in electron volts, the energy an electron needs to move through a one-volt potential. To calculate the wavelength λ (in nanometers) divide 1240 by the bandgap energy E in electron volts:

$$\lambda = \frac{1240}{E} \quad (10-5)$$

10.9.2 Lattice Constants and Compound Semiconductors

Atomic spacing is a second crucial parameter for compound semiconductors, because the lattice constant of adjacent layers needs to be closely spaced to avoid flaws in the crystal, as you learned in Section 10.5.2.3. This can be tricky because controlling the flow of current and light in the semiconductor requires using different compounds in the layered structure, and strain must be minimized between the layers to avoid flaws that lead to failure of the laser.

Figure 10-22 shows the bandgap energy, and lattice spacing varies for different blends of III–V semiconductors usable in diode lasers. For example, a nearly vertical line in the middle of the graph connects dots representing the bandgap and lattice spacing of GaAs and AlAs. Add a bit of one compound to the other, and you move to an intermediate point. For example, a mixture of 20% AlAs and 80% GaAs has a higher bandgap than pure GaAs. The formula for such a compound can be written as $\text{Ga}_{0.8}\text{Al}_{0.2}\text{As}$ or simply GaAlAs. (The order of the two elements that replace each other—aluminum and gallium—is arbitrary; some people write AlGaAs.) By good fortune, the lattice constant changes very little as AlAs is added, which makes ternary GaAlAs diodes relatively easy to make.

Things get more complex if the compound semiconductor contains four elements. Depending on the relative concentrations of the elements, such “quaternary” semiconductors can have properties somewhere in a broad range. For example, as shown in Figure 10-22, InGaAsP (more precisely $\text{In}_x\text{Ga}_{1-x}\text{As}_{1-y}\text{P}_y$) can have lattice spacing and bandgap within an area defined by four possible binary compounds: InAs, InP, GaAs, and GaP. Because you can vary either the In/Ga ratio or the As/P ratio, this is like having two “knobs” to adjust to match the lattice spacing and maintain the desired bandgap energy to match the lattice spacing of the substrate.

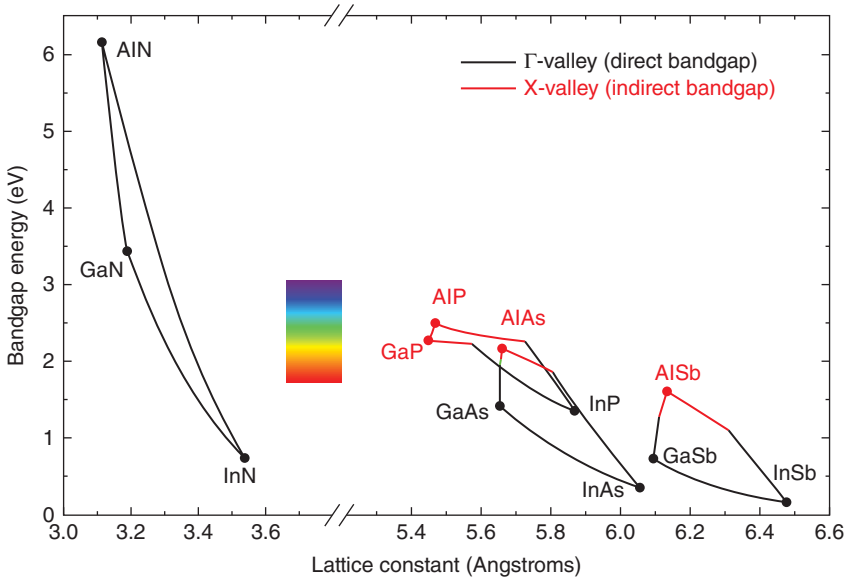


Figure 10-22. Lattice constant versus bandgap for III–V semiconductors in the ultraviolet, visible, and near infrared, including compounds with direct and indirect bandgaps, as marked. Note discontinuity in lattice-spacing scale. (Courtesy of Colin Humphreys, University of Cambridge)

Another approach to lattice matching is by depositing a series of thin layers with slightly different composition. This approach, called a *strained layer superlattice*, spreads the strain from the lattice mismatch through many layers, each of which has to withstand only a small strain. This structure is used to fabricate InGaAs lasers on GaAs substrates, which otherwise would be likely to fail because of strain in the structure.

Note that Figure 10-22 plots bandgap energy and lattice spacing of the compounds, not the wavelengths available in commercial diode lasers. For example, although nitride semiconductors have bandgaps from about 200 to 620 nm (6.2 to 2 eV), nitride diode lasers have operated from only about 370 to 540 nm, and commercial diode lasers are limited to a smaller range of 375 to 525 nm.

10.9.3 Nitride Diodes from Ultraviolet to Green

The family of nitride compounds at the left side of Figure 10-22 are used in diode lasers with output from the ultraviolet to the green

Table 10-3. Bandgap energies and wavelengths of pure nitride compounds

Binary nitride	Bandgap energy	Bandgap wavelength
InN	0.7 eV	1771 nm
GaN	3.4 eV	362 nm
AlN	6.0 eV	206 nm

part of the spectrum. The best developed compound is gallium–indium nitride (GaInN), which is used in commercial diode lasers emitting from 370 nm in the near ultraviolet to 525 nm in the green. The lasers often are called GaN because the indium is a minor constituent added to GaN, which has a bandgap of 3.4 eV corresponding to a 362-nm wavelength. Adding indium decreases the bandgap and increases the output wavelength; pure InN has a bandgap of 0.7 eV and wavelength of almost 1800 nm, as shown in Table 10-3.

All three nitride compounds and their alloys have direct bandgaps that span nearly an order of magnitude in energy. However, the range of wavelengths at which laser diodes can be made has been limited by difficulties in fabricating good materials and electrical contacts. AlN has a small lattice constant that tends to produce flaws in areas rich in aluminum. Fabrication of *p*-doped layers becomes harder as indium concentrations increase above a certain level. Thus efforts to expand the range of direct diode-laser wavelengths beyond the 370- to 525-nm have largely stalled out, although LEDs can reach much shorter and longer wavelengths.

Recently the quest to make deeper-ultraviolet diode lasers has turned to frequency doubling and quadrupling of output of longer-wavelength diode lasers. This can be achieved by narrowing their linewidth with external cavities and amplifying their output.

GaInN lasers are the brightest visible diode lasers, with the highest power available at blue and violet wavelengths. The most common wavelengths in the bright bands are 405 nm in the violet and 445 to 460 nm in the blue. Green and ultraviolet outputs generally are lower.

GaInN lasers at 405 nm originally were developed to provide high-density data storage for Blu-Ray video disk players. Their high power in the blue has led to their use in projection displays. Blue LEDs are combined with yellow phosphors to provide bright white light for use in indoor and outdoor lighting. Violet lasers also

can be combined with phosphors for illumination; the laser light is scattered so the beam itself does not enter the field of view. Green GaInN lasers are used in some projection displays, together with blue and red lasers to provide full color.

Blue and violet diode-laser pointers are readily available on the Internet, but you should take care in buying and using them. Some are advertised as “pointers” despite delivering power well above the 5-milliwatt maximum considered safe for laser pointers. (See Appendix A for a brief description of laser safety.) Some pointers advertised as “blue” actually emit violet light at 405 nm. The eye responds weakly to violet light, but many materials fluoresce brightly when exposed to 405 nm. That bright flash of fluorescence should be a reminder that what looks like a faint violet beam when pointed at some objects packs much more power than your eye senses.

Green laser pointers generally are not diodes but actually are diode-pumped, frequency-doubled solid-state lasers emitting at 532 nm, which pose their own safety concerns.

10.9.4 Red Diode Lasers

Red diode lasers have become commonplace as laser pointers and cat toys. The first to reach the market in the late 1980s emitted at 670 nm and had active layers of $\text{Ga}_{0.5}\text{In}_{0.5}\text{P}$ surrounded by layers in which aluminum replaced some of the gallium to raise the bandgap and improve confinement.

Most red diode lasers today emit at 635 or 650 nm. The human eye is 10 times more sensitive to 635 nm than to 670 nm, so 635 nm is used where human visibility is important. The 650 nm wavelength is the standard for standard-definition video and data storage on DVD drives, where it offers the best balance between the higher power at longer wavelengths and the higher data recording density of shorter wavelengths. Both have active layers of AlGaInP, in which the number of aluminum, gallium, and indium atoms together equal the number of phosphorous atoms, a blend that can be written as $\text{Al}_y\text{Ga}_x\text{In}_{1-x-y}\text{P}$. Like 635-nm diodes, they are grown on GaAs substrates.

Red laser pointers and DVD playback lasers emit milliwatt class beams. Red diodes also can be fabricated in arrays delivering up to watts of red light. Generally, diodes emitting longer wavelengths produce higher power and have longer lifetimes.

10.9.5 Gallium Arsenide Near-Infrared Lasers

The oldest and best developed family of diode lasers have active layers of GaAlAs or GaAs and are fabricated on GaAs. They are usually called “gallium arsenide” lasers. Pure GaAs active layers nominally emit at 904 nm, but replacing some of the gallium with aluminum increases the bandgap to generate shorter wavelengths. Fortunately, aluminum and gallium atoms are nearly the same size, so GaAlAs can be easily lattice-matched to a GaAs substrate. However, flaws are more likely to develop at high aluminum concentrations, so the shortest wavelengths from GaAlAs lasers typically are around 750 nm in the near infrared.

GaAlAs lasers emitting at 780 nm are widely used in instrumentation, CD players, CD-ROM drives, laser printers, and some short fiber-optic data links.

High-power GaAlAs lasers are used extensively to pump solid-state lasers. The most widely used types are 808-nm lasers for pumping neodymium-doped lasers. High-power pump diodes also are produced at 830 and 850 nm.

10.9.6 InGaAs Pump Lasers

Obtaining longer wavelengths from lasers in the GaAs family requires adding indium to replace some of the gallium and reduce the bandgap. Fabricated on GaAs substrates, InGaAs lasers can generate high powers for pumping solid-state or fiber lasers. InGaAs emitting at 915 nm is widely used as a pump for ytterbium-doped fiber lasers. InGaAs emitting at 980 nm is widely used for pumping erbium-doped fiber amplifiers and erbium-fiber lasers. InGaAs emitting at 940 nm is used for pumping erbium-doped and erbium-ytterbium-codoped fiber lasers and amplifiers.

Adding indium to increase the wavelength of the laser also increases the lattice constant, requiring fabrication of thin strained layers between the GaAs substrate and the InGaAs layers to prevent dislocations in the laser structure.

10.9.7 InGaAsP Lasers for Fiber-Optic Systems

The quaternary (four-element) semiconductor InGaAsP is a versatile laser material grown on InP substrates. The sum of the number of indium and gallium atoms equals the total number of arsenic

and phosphorous atoms. The chemical formula can be written as $\text{In}_x\text{Ga}_{1-x}\text{As}_y\text{P}_{1-y}$, with the values of x and y varying independently. Adding either element changes the lattice constant, so successfully growing the compound on an InP substrate requires balancing the composition to match the lattice constant to the substrate, as shown in Figure 10-22.

InGaAsP active layers can emit at 1100 to 1650 nm when lattice-matched structures are grown on InP substrates. The major applications of InGaAsP lasers are in fiber-optic communications, so most lasers are manufactured for the fiber-optic bands around 1310 nm and from about 1480 to 1600 nm.

The type of InGaAsP lasers used depends on the number of wavelengths being transmitted through the fiber-optic system. Wavelength tolerances are ± 20 nm or so for systems transmitting only a single wavelength, which can use uncooled Fabry–Perot lasers. Systems transmitting several wavelengths about 20 nm apart (coarse wavelength-division multiplexing) also use Fabry–Perot lasers, but more care must be taken to see that their output falls into the proper wavelength channel.

High-speed networks pack wavelengths much closer together to increase capacity; this is called dense wavelength-division multiplexing. These wavelength slots are only a fraction of a nanometer wide, so they require narrow-line DFB or DBG lasers kept at a stable temperature.

10.9.8 Other Diode Lasers

As you can see in Table 10-2 and Figure 10-22, diode lasers can be made from a variety of other semiconductor compounds.

III–V compounds of indium, gallium, arsenic, and antimony (InGaAsSb) have smaller bandgaps than InGaAsP and can be deposited on substrates of GaSb. Diode lasers made of these materials emit at wavelengths between 1.8 and 5 μm . Lasers at some of these wavelengths are available commercially; their main applications are in instruments.

Diode lasers have been demonstrated in compounds containing elements from groups II and VI of the periodic table (also known as columns 12 and 16). These II–VI compounds include zinc sulfide, zinc selenide, zinc telluride, cadmium selenide, and cadmium sulfide. However, GaN lasers proved more practical for short-wavelength diode lasers.

A family called “lead salt” diode lasers emits from $2.7\ \mu\text{m}$ to about $30\ \mu\text{m}$ in the infrared, depending on composition, as listed in Table 10-2. Made of compounds containing lead, sulfur, selenium, and other elements, they require cooling to cryogenic temperatures. Their prime applications were in research and precision measurement of the infrared properties of materials.

10.9.9 Integrated Photonics, Silicon Photonics and “Silicon Lasers”

As you learned in Section 10.3.4, silicon, the semiconductor that dominates electronics, has an indirect bandgap. Much to the annoyance of the silicon industry, many years of effort and untold millions of dollars have not succeeded in overcoming this barrier to making silicon diode lasers that can be integrated with electronics on silicon wafers.

Instead, the industry is developing a hybrid technology called *integrated photonics* or sometimes *silicon photonics* that combines diode lasers made from compound semiconductors with silicon electronics on a silicon substrate. Figure 10-23 shows how an optical chip based on indium–phosphide or another III–V material can be bonded to a silicon wafer. Adding optical waveguides and electronic circuits to process information creates an integrated photonic circuit that combines the strength of silicon electronics for processing data and the strength of compound semiconductors in transmitting signals as light at much higher data rates than possible over wires.

Integrated photonics is important to mention because it promises to overcome data-transmission bottlenecks in places like




Figure 10-23. A hybrid InGaAsP/Si waveguide laser generates light in the compound semiconductor and transfers it to the silicon waveguide. (Courtesy of Intel Corp. and University of California Santa Barbara.)

cloud computing centers and data processing centers. However, it is more of an evolutionary technology than a revolutionary one because it combines the existing strengths of integrated electronics and compound semiconductor optoelectronics.

10.10 QUANTUM CASCADE LASERS AND RELATED TYPES

Quantum cascade lasers (QCLs) are semiconductor lasers which differ fundamentally from diode lasers. Instead of extracting energy from electrons and holes recombining at a central junction, a quantum cascade laser extracts energy from electrons as they pass through a series of quantum wells in the semiconductor. The transition in a diode laser is between two bands, the conduction band and the valence band. The transition in a quantum cascade laser is between two subbands in the conduction band, which releases less energy and thus emits infrared light. This means that only electrons carry current in QCLs, and all electrons pass through the stack of quantum wells in the same direction, emitting photons as they drop to lower subbands.

A voltage is applied across the entire stack of quantum wells in a quantum cascade laser. An electron that enters a quantum well is first trapped in an upper subband of the conduction band in quantum well. In that state it can be stimulated to emit a photon and drop to a lower substate in the quantum-well conduction band. Then the electron can tunnel through the valence band in barrier layer to the next quantum well, where it is trapped in the upper energy substate until it emits a photon, drops to the lower substate, and tunnels through another barrier to the next quantum well. This produces a quantum cascade of photons, all at the same wavelength. Typically it involves tens of stages.

Figure 10-24 shows an electron moving from left to right through a series of quantum wells separated by barrier layers. The quantum wells are the low parts of the slanted lines, which represent the electric voltage applied across the quantum cascade laser. The barrier layers are the high parts, where bandgaps are high. The barriers are thin enough for the electron to slip through, landing in an upper substate (horizontal line) in the quantum well. Stimulated or spontaneous emission drops the electron to a lower-energy substate in the quantum well, where it repeats the process. The slant




Figure 10-24. Electrons cascading through a series of quantum wells separated by barriers emit light in a quantum cascade laser.

comes from the variation of the electric field along the laser. The electric potential creates a gradient in energy levels that essentially pulls the electron through the structure, like a marble rolling down stairs. Ideally, the electron could be stimulated to emit energy each time it drops to a new quantum well, so one electron could emit many photons, but it is not that easy. Figure 10-24 is a simplified view, so each step shown actually represents a series of thin layers that combine to function as one level. The figure does not show the

laser cavity, which includes reflectors at each end and a waveguide through the semiconductor that guides light through the cavity.

10.10.1 QCL Output Bands and Powers

The conduction-band substates in the quantum wells are tailored by selecting the composition and thickness of the quantum wells. In practice, the subband transitions have less energy than the recombination radiation released in diodes lasers. QCLs emit in two infrared bands at longer wavelengths than emitted by most diode lasers. One is from 3 to 30 micrometers, important because it covers vibrational transitions of many molecules. A second band is at wavelengths of 60 to 500 μm , which is called the *terahertz band* because it corresponds to frequencies of 0.6 to 5.0 terahertz. The terahertz band is important because it includes rotational transitions of many molecules. (InP and GaAs, the semiconductors used in most QCLs, suppress emission at 30 to 60 μm .)

QCLs use optical cavities similar to those in diode lasers. Early versions used Fabry–Perot cavities with simple mirrors at the ends, like edge-emitting diodes. Distributed-feedback cavities offer narrower bandwidth. External-cavity lasers enhance tunability.

Tunability is important because QCLs are one of the few laser sources available for sensing and spectroscopy at wavelengths of 5 to 14 μm . Typically several lasers would be needed to scan across the entire range, but developers are trying to enhance tunability. Output power usually is from milliwatts to tens of milliwatts, varying across the band. Watts have been generated at wavelengths between 4 and 5 μm , where high powers are sought to blind heat-seeking missiles that look for those wavelengths. QCLs can operate pulsed or continuously, but their internal dynamics make it impractical to generate short pulses by Q switching or mode-locking. Active development of the technology continues, particularly in the terahertz range and in improving output efficiency.

10.10.2 Interband Cascade Lasers

A hybrid approach that combines features of diode and QCLs offers higher power at wavelengths of 3 to 4 μm . This is approach is called an *interband cascade laser* because it produces a cascade of emission from a series of quantum wells where the transitions are between the conduction and valence band rather than between

subbands of the conduction band. This requires both electrons and holes to be carrying current through the structure.

Interband cascade lasers consist of alternating layers of semiconductors with lower bandgap energies than those used in QCLs. For example, carriers cascading through the system pass between conduction band minima in indium–arsenide layers and valence-band maxima in adjacent layers of gallium–indium antimonide. Because of their complex structure, they use several layers rather than the tens of layers in QCLs. The technology remains in development.

10.11 WHAT HAVE WE LEARNED?

- Diode lasers are the most important types of semiconductor lasers.
- Diode lasers can convert up to 70% of the input electrical power into laser light.
- Diodes are two-terminal electronic devices that conduct current in one direction.
- Electrons in a semiconductor may be bound to atoms in the low-energy valence band, or able to move freely in the solid in the conduction band.
- The difference in energy between the valence band and the conduction band is called the bandgap and depends on the nature of the semiconductor.
- Holes are vacancies in the valence band of a semiconductor. A hole moves when an electron from another atom fills the empty space in the valence band.
- *n*-type semiconductors are doped with elements that release electrons to carry current; *p*-type semiconductors are doped with elements that form holes.
- Current flow through a forward-biased diode causes electrons to recombine with holes at the junction layer between *n*- and *p*-type semiconductors.
- An electron–hole pair called an exciton exists briefly before the electron drops into the hole and releases its energy as recombination radiation.
- Recombination radiation at the junction layer of an LED produces spontaneous emission at the bandgap energy of the semiconductor.

- Semiconductors need a direct bandgap to emit light efficiently. Silicon has an indirect bandgap so it does not emit light efficiently.
- III–V compounds such as GaAs, InGaN, and InGaAsP are called compound semiconductors and mostly have direct bandgaps.
- Diode lasers have resonant cavities and produce stimulated emission when the drive current exceeds laser threshold.
- Room-temperature continuous-wave emission requires a double heterojunction, with a junction layer of one composition sandwiched between layers of other composition.
- Confining emitted light to a narrow stripe in the junction plane improves diode-laser performance.
- Confining current flow to a narrow stripe in the junction layer constrains laser gain to that region and is called gain guiding.
- Surrounding the active stripe with material having a lower refractive index is called index guiding and relies on total internal reflection.
- High-power diode lasers require either a broad stripe or arrays of many parallel stripes.
- A simple edge-emitting diode laser oscillates in the junction plane, with the edges of the chip serving as the resonator mirrors. This design is called a Fabry–Perot laser, because it uses two parallel surfaces as a resonator.
- The emitting area of a narrow-stripe laser is a few micrometers wide and a fraction of a micrometer high.
- Distributed-feedback and distributed Bragg-reflection lasers confine diode-laser oscillation to a narrow range of wavelengths by controlling feedback in the laser.
- External-cavity lasers confine light to a narrow range of wavelengths that can be tuned by adjusting the optics.
- Semiconductor optical amplifiers resemble diode lasers, but reflection from edge facets is suppressed so they can amplify light but not oscillate.
- A VCSEL is a diode laser that oscillates in a vertical cavity, perpendicular to the junction plane and emits light from its surface. VCSELs are inexpensive to fabricate and package and have low thresholds currents and high efficiencies.
- A quantum well is a double heterojunction thinner than 50 nm which confines electrons tightly in a layer with lower bandgap energy than the surrounding layers.
- Edge-emitting diode lasers have a large beam divergence, but external optics can focus the light into a narrow beam.

- Defects occur in diode lasers if the atomic spacing in successive layers is not matched or managed carefully.
- GaInN diodes emit blue, violet, and ultraviolet light.
- Diode lasers with active layers of AlGaInP emit at 635 nm in the red.
- GaAlAs lasers emit in the near infrared; their main use is in CD and CD-ROM players. They also can emit high power for pumping other lasers at 808, 830, and 850 nm.
- InGaAs lasers emitting at 980 nm are widely used as pump lasers.
- InGaAsP diode lasers emit at wavelengths of 1100 to 1650 nm; their main applications are in fiber-optic communication systems.
- Packaging determines the function of a diode laser; it often costs more than making the laser itself.
- Diode lasers are simple to modulate by changing their drive current.
- A quantum cascade laser extracts energy in a series of steps from electrons passing through a series of quantum wells. It emits at infrared wavelengths longer than 3 μm .

WHAT'S NEXT?

In Chapter 11, we will describe types of lasers that are in development or do not fit into the major categories of Chapters 7 through 10. These include tunable organic-dye lasers, free-electron lasers, and extreme-ultraviolet lasers.

QUIZ FOR CHAPTER 10

1. In order to make a good diode laser, a material should be
 - a. A semiconductor with an indirect bandgap
 - b. A semiconductor with a direct bandgap
 - c. An insulator with an indirect bandgap
 - d. An insulator with a direct bandgap
 - e. A conductor
2. Which of the following is a quaternary III–V semiconductor?
 - a. InGaAsP
 - b. PbSnSSe
 - c. GaAlAs
 - d. GaAs
 - e. NSbAsP

3. The active layer of a diode laser is made of $\text{Ga}_{0.8}\text{Al}_{0.2}\text{As}$. What percent of the atoms in the active layer are of each element?
 - a. 80% gallium, 20% aluminum, 100% arsenic
 - b. 80% gallium, 20% aluminum, 0% arsenic
 - c. 40% gallium, 10% aluminum, 50% arsenic
 - d. 20% gallium, 80% aluminum, 100% arsenic
 - e. 33% gallium, 33% aluminum, 33% arsenic
4. An exciton is
 - a. A free electron in the conduction band
 - b. A free electron in the valence band
 - c. The energy released by the combination of an electron and a hole
 - d. An electron–hole pair in an excited state because the electron has not combined with the hole
 - e. A molecule that exists only in the excited state.
5. What gives a double-heterostructure laser better efficiency than a homostructure laser?
 - a. Reverse biasing
 - b. Better confinement of conduction electrons
 - c. Restricting of current flow to the active layer
 - d. Lower levels of spontaneous emission
 - e. A homostructure laser is more efficient
6. Which type of laser can be made from GaAlAs?
 - a. Fabry–Perot edge emitter
 - b. VCSEL
 - c. External cavity
 - d. Distributed-feedback edge emitter
 - e. All of the above
7. Which kind of laser has the shortest optical cavity?
 - a. Helium–neon gas
 - b. Edge-emitting diode
 - c. External-cavity diode
 - d. VCSEL
 - e. Impossible to tell
8. You are building a video player for use with high-definition disks. What kind of laser do you need?
 - a. InGaN diode
 - b. GaAlAs edge emitter
 - c. GaAlAs VCSEL
 - d. InGaAsP distributed feedback
 - e. Quantum cascade

9. A semiconductor has a bandgap of 1.5 electron volts. At what wavelength will it emit light if it can operate as a laser?
 - a. 1500 nm
 - b. 1000 nm
 - c. 827 nm
 - d. 667 nm
 - e. 333 nm
10. Which types of semiconductor lasers emit a single longitudinal mode?
 - a. Homojunction
 - b. Buried heterostructure
 - c. Monolithic arrays
 - d. Distributed feedback
 - e. Those built for compact disc players
11. What family of diode lasers emits the shortest wavelength?
 - a. GaAlAs
 - b. AlGaInP
 - c. InGaAsP
 - d. GaInN
 - e. AlGaN
12. What type of diode laser is most likely to be used in a red laser pointer?
 - a. GaAlAs
 - b. AlGaInP
 - c. InGaAsP
 - d. GaInN
 - e. AlGaN
13. You need a laser with a very low threshold, high efficiency, and a good quality beam without extra optics. Which is the best choice?
 - a. Distributed feedback
 - b. Edge emitter
 - c. Monolithic array
 - d. Quantum cascade
 - e. VCSEL
14. You need a laser that can generate a large amount of power efficiently, and does not require high beam quality. Which is your best choice?
 - a. Broad stripe single edge emitter
 - b. Narrow-stripe edge emitter
 - c. Monolithic array

- d. Stack of multiple arrays
 - e. VCSEL
15. You know a photon with energy of 1 electron volt has a wavelength of 1240 nm, and you know the product of the energy times the wavelength equals a constant hc . You know the product of the photon energy and the wavelength is a constant, but you do not remember the constant and the Internet is down. Calculate the energy of a 400-nm photon.
- a. 0.31 eV
 - b. 0.496 eV
 - c. 1.42 eV
 - d. 2.48 eV
 - e. 3.1 eV

OTHER LASERS AND LASER-LIKE SOURCES

ABOUT THIS CHAPTER

Not all lasers fit into the categories covered in Chapters 7 to 10, and some light sources generate laser-like light but not by stimulated emission in an oscillator. This chapter describes these lasers and laser-like sources and explains how they work.

Many of these sources have particular niches in the laser world. Dye lasers and parametric sources can be tuned broadly in wavelength. Free-electron lasers and high-harmonic sources can extend into the far ultraviolet and X-ray bands, and free-electron lasers are broadly tunable and can attain very high peak powers as well. Femtosecond frequency combs and supercontinuum sources generate wide ranges of wavelengths. Superradiant sources are emitters in the hazy boundary between LEDs and diode lasers. We also cover some emerging laser concepts that may find important applications in the future.

11.1 TUNABLE DYE LASERS

The ability to tune the wavelength of laser light is invaluable for laser research, development, and measurement. Today laser users have many options, but in the early years of laser development their

first good option was optical pumping of organic dyes dissolved in liquid solvents, which remains in use today.

Dye lasers are based on a family of large organic molecules that get their bright colors from transitions between complex sets of electronic and vibrational energy levels. The vibrational energy levels create many sublevels of the electronic states, as in solid-state vibronic lasers, so transitions can occur across an unusually wide range of energies, and laser gain can occur across a range of wavelengths. Although each dye has gain over a limited range, many different dye molecules are available, and together they cover the near-infrared, visible, and near-ultraviolet.

Laser dyes require optical pumping with light from an external laser or another bright light source like a flash lamp. Normally, the dyes are dissolved in an organic solvent or other liquid. Laser action has been reported in vapor-phase dyes, but they do not work well. Solid dye lasers also have been demonstrated, but have problems dissipating heat and replenishing molecules, which break down during laser operation.

11.1.1 Properties of Laser Dyes

Dye molecules fluoresce, absorbing a short-wavelength photon and spontaneously emitting a longer-wavelength photon. Upper-state lifetimes are typically a few nanoseconds, and bright pump light is needed to produce a population inversion. The laser transition is between a level in an electronically excited band and a vibrationally excited sublevel of the ground electronic state, as shown in Figure 11-1. Dye sublevels are so closely spaced that they form a continuum, so the wavelength can be tuned continuously.

The bandwidths of laser dyes can range from 10 to 70 nm, with typical ranges 20 to 40 nm. Tuning across wider ranges requires switching dyes, and some dye lasers have multiple dye cells to make this easier. The tuning range of a particular dye may differ between light sources.

Most dyes belong to one of several families of organic compounds, and seemingly countless variations are possible just by adding an atom or two at different places on the molecule. For example, there are about 100 dyes in the coumarin family of dyes shown in Figure 11-2. The question marks show where atoms can be added to make new molecules.

Organic dyes have important limitations. They degrade after tens or hundreds of hours of use, even when the dye solution is




Figure 11-1. Energy levels in a typical laser dye.

flowed through the cell to reduce exposure to intense pump light. Most do not dissolve readily in water and require organic solvents which often are flammable and toxic. Dye solutions are generally classified as hazardous waste both because of the dye and the solvents.

11.1.2 Dye Laser Structure

Dye lasers focus pump light onto a dye solution in a resonant cavity that includes wavelength-tuning optics as well as the usual cavity




Figure 11-2. Chemical structure of coumarin laser dyes.

mirrors. In low-power dye lasers, the dye solution may be sealed in a closed cell, but higher powers require pumping the dye solution through the pumping region to avoid rapid degradation. Pulsed pump sources deliver high peak power to reach laser threshold. Continuous dye lasers require tight focusing of the pump beam to achieve high power densities in an unconfined jet of dye solution that flows through the pump volume.


Some dye lasers have cavities that allow simultaneous oscillation at a wide range of wavelengths, but most cavities are tunable and select a narrow range of oscillating wavelengths. Figure 11-3 shows three types of tunable cavities used in dye lasers.

The laser's wavelength is tuned by turning prisms or diffraction gratings, which direct light at different angles as a function of its wavelength, or by adjusting cavity mirrors or other optics. In Figure 11-3A, the grating is turned to change the wavelength, but in Figure 11-3B the prism stays fixed while the rear cavity mirror moves to select a particular wavelength refracted by the prism. In both cases, the selected wavelength resonates in the cavity, and other wavelengths are deflected out of the cavity.


In Figure 11-3C, the cavity can be tuned in two ways, coarsely with a wedge filter, or more precisely with an etalon, which functions as a miniature optical cavity to limit laser oscillation to a narrow range of wavelengths. Moving the etalon changes how far light in the oscillator travels between its reflective surfaces, which selects the wavelength. The figure shows both types of tuning, etalons with two parallel surfaces that are turned and wedge filters that are moved up and down.

Simple prisms and gratings can limit dye laser linewidth to less than 0.01 nm. Adding an etalon can further restrict the linewidth, and adding further equipment and control systems can limit emission range to much less than one part in a billion (10^9). Tunable laser cavities can also take odd shapes such as the ring cavity shown in Figure 11-4. Internal optics in a ring laser restrict oscillation to one direction, and other components limit oscillation to a narrow range of wavelengths. This design can be mode-locked to generate a series of extremely short pulses.


Note that Figures 11-3 and 11-4 do not show everything needed to operate a dye laser; the drawings are simplified to help you understand their operational principles. Similar cavities can be used with other tunable lasers, such as the titanium–sapphire lasers described in Section 8.6.2.



(A) Grating-tuned dye laser (pumped by pulsed laser).



(B) Prism-tuned dye laser (pumped by continuous-wave laser).



(C) Etalon-tuned dye laser (pumped by continuous-wave laser).

Figure 11-3. Three examples of wavelength-tunable laser cavities used with dye lasers.

11.1.3 Dye Laser Pumping

Dye laser pumping requires high pump intensities to efficiently excite dye molecules to the upper laser level. Either lasers or flash lamps can provide the pump energy.

Many types of pulsed lasers can pump dye lasers, offering high peak power at specific wavelengths that can be matched to dye absorption bands. The second, third, and fourth harmonics of neodymium lasers are widely used today, as are nitrogen and excimer gas lasers.




Figure 11-4. A ring dye laser.

Continuous dye pumping is possible with harmonics of neodymium or other solid-state lasers, but CW dye lasers have fallen out of favor because most researchers find tunable solid-state lasers easier to use. CW dye lasers can be mode-locked to generate chains of short pulses.

Flashlamps also can generate intense pulses of visible light suitable for exciting many laser dyes. Linear flashlamps can be focused onto a tube containing dye solution in an elliptical cavity, as for lamp-pumped solid-state lasers. An alternative is using a coaxial flashlamp with dye solution flowing through the center. However, flashlamp-pumped dye lasers are not widely used today.

11.2 OPTICAL PARAMETRIC SOURCES

Tuning across a much broader wavelength range is possible with a family of nonlinear devices called *optical parametric sources*. They are based on a process called *three-wave mixing* in which photons in an intense pump beam passing through a nonlinear material essentially split into two photons, the signal and the idler, which together have the same total energy as the pump wave: $\nu_{\text{pump}} = \nu_{\text{signal}} + \nu_{\text{idler}}$. The two most important optical parametric

sources are *optical parametric amplifiers (OPAs)* and *optical parametric oscillators (OPOs)*.

Optical parametric sources were originally developed at radio frequencies, but their principles were applied to optics within a few years after invention of the laser. They are attractive because they can generate a wide range of wavelengths which retain the coherence and beam quality of the laser pump source. Early versions were complex and required highly skilled operators, but automation and control systems have made them more practical sources of broadly tunable laser-like output.

11.2.1 Optical Parametric Amplifiers

An OPA starts with a strong pump beam from a laser source at a frequency ν_{pump} and a weak signal beam at a frequency ν_{signal} entering a nonlinear medium. The two beams follow the same path through the nonlinear medium and three-wave mixing takes energy from the pump beam to amplify the signal beam, and in the process generates a third wave, the idler at a frequency $\nu_{\text{idler}} = \nu_{\text{pump}} - \nu_{\text{signal}}$. Normally the process shifts most of the energy in the pump beam to the signal and idler beams, leaving only a weak pump beam, as shown in Figure 11-5. The pump and signal beams retain the coherence and other properties of the laser beam, but the output is not considered a laser because it is generated by a nonlinear process, not stimulated emission.

Like other optical amplifiers, OPA is a single-pass device without mirrors or a resonant cavity. The signal and pump beam start




Figure 11-5. In optical parametric amplification, photons at the strong pump frequency ν_{pump} passing through a nonlinear medium amplify the signal frequency ν_{signal} and produce the idler frequency ν_{idler} . This converts most energy from the pump beam into the two lower frequencies that add together to pump frequency: $\nu_{\text{pump}} = \nu_{\text{signal}} + \nu_{\text{idler}}$.

together at the input, three-wave mixing occurs in the nonlinear material, and the output beams emerge together, as shown in Figure 11-5. Typically a weak residual pump beam remains as well as the signal beam and the idler, both of which can be used to generate wavelengths from about 200 nm to 16 μm , a wider range than available from OPOs.

OPAs can be used to amplify the output of OPOs. The very wide bandwidth of OPAs makes them particularly well-suited for *chirped pulse amplification*, a technique described in Section 6.5.4 that was developed to amplify very short pulses to extremely high peak powers without damaging the amplifier. Normal amplification of such pulses would generate extremely high peak powers that could damage the interior of the amplifier or its output window. Chirped pulse amplification avoids such damage by stretching (chirping) the pulse duration by passing it through optics with a refractive index that increases or decreases with wavelength, so some wavelengths fall behind the others. Then the stretched pulse can be amplified, producing a long pulse with higher energy that can then be passed through another optical system to compress the pulse back to its original length by delaying the wavelengths that got ahead in the original stretching. This has become an important tool for producing pulses that are only a few waves long and have extremely high peak power.

OPAs were demonstrated as early as the 1960s, but their applications then were limited by their complexity and need for delicate adjustments. The development of advanced control systems has made OPAs much easier for nonspecialists to operate, and they have become standard commercial products.

11.2.2 Optical Parametric Oscillators

An OPO is based on the same physical effect, but works somewhat differently. A strong laser pump beam at a frequency ν_{pump} is focused into a nonlinear crystal housed inside an optical cavity with a resonance at a lower frequency ν_{signal} called the *signal*. As shown in Figure 11-6, the only input to the OPO is the pump beam, but background waves are always present at the signal and idler frequencies, so the high-power pump beam frequency mixes with them in the nonlinear material, converting some pump energy to the signal and idler frequencies. This is the same three-photon interaction as in an OPA, splitting the energy of a pump photon




Figure 11-6. In optical parametric oscillation, the only input is a strong pump beam at ν_{pump} which enters a resonant cavity containing a nonlinear material. Background waves at the signal frequency ν_{signal} and the idler frequency ν_{idler} mix with the pump and oscillate in the resonant cavity, converting the pump beam to the two other frequencies $\nu_{pump} = \nu_{signal} + \nu_{idler}$.

between one photon at ν_{signal} and a second photon at the *idler* frequency, ν_{idler} .

It is functionally equivalent to gain at those frequencies, but it is not stimulated emission. If the gain in the resonant cavity from three-wave mixing exceeds the loss, the OPO begins to oscillate at the resonant frequency as pump power reaches a threshold, as in a laser. Because OPOs emit light produced by three-wave mixing rather than by stimulated emission, it is not considered to be a laser source, although it oscillates in a resonant cavity and is pumped by a laser.

The oscillation wavelength depends on parameters of the resonant cavity, including the orientation and temperature of the nonlinear material. Generally, only one of the two frequencies produced by three-wave mixing oscillates in the cavity at any one time. The other produced frequency and the remaining pump power are dumped. OPOs can emit pulsed or continuous beams. Their output is coherent, and the spectral width of a continuous-output OPO can be very narrow, making it attractive for research. An OPO pumped by an ultraviolet laser can be tuned across the whole visible spectrum and well into the near-infrared. The tuning range can reach 200 to 2600 nm with the addition of optics for generating the second harmonic and sum frequencies.

An OPO can be used for downconversion if the signal and idler waves are tuned to the same wavelength, so $\nu_{pump} = 2\nu_{output}$. This converts one input photon into two output photons at twice the

wavelength. The two photons have correlated quantum properties, so they can be used in quantum mechanical experiments.

11.3 SUPERCONTINUUM SOURCES

When intense light passes through a highly nonlinear medium, nonlinear effects can combine to spread the wavelengths of the light across a broad range both above and below the wavelength of the pump laser. The higher the peak power and the stronger the nonlinearity of the medium, the broader will be the range of wavelengths. For example, a short-pulse 1064-nm laser pumping a highly nonlinear glass fiber can generate wavelengths from 380 to 2400 nm, limited on the long-wavelength end by glass absorption. A different laser wavelength can produce a supercontinuum spanning a different range. The physics are complex, and the properties of the supercontinuum depend on details including the wavelength and duration of the pump laser pulse and the properties of the nonlinear medium, generally a highly nonlinear optical fiber.

In some cases the output power can be relatively uniform across the spectrum; in others it may vary considerably. Spatial coherence can be high, particularly for a supercontinuum generated in a single-mode optical fiber. Temporal coherence may be low, but can be high if the supercontinuum is generated from a series of short periodic pulses.

Commercial supercontinuum sources may be called “white lasers” because their output is spread across such a broad band, but like OPOs and OPAs, the output is not stimulated emission. Instead, it is laser light modified by nonlinear processes. Typically output is repetitively pulsed at kilohertz to megahertz, with total average power ranging from a fraction of a watt to tens of watts, with 10% to 25% in the visible spectrum.

11.4 FREQUENCY COMBS

One of the most elegant concepts to emerge from research on laser spectroscopy and time measurement in the 1990s was the femtosecond frequency comb, which earned John Hall of the National Institute of Standards and Technology in Colorado and Theodor Hänsch of the Max Planck Institute for Quantum Physics in Germany shares

of the 2005 Nobel Prize in Physics. Its elegance comes from the way it transforms the series of pulses produced by a mode-locked laser into a rainbow of equally spaced continuous wavelengths of light.

That transformation comes from a mathematical technique called *Fourier analysis* that engineers use to analyze how a signal varies over time. A signal that is periodic in time can be broken down mathematically into a series of sine and/or cosine waves at *harmonics*, multiples of the signal frequency. A sine wave is the simplest form; its Fourier transform is a single sine wave at the sine wave's own frequency. The transform of a square wave—a light turning on and off—is the sum of sine waves at odd harmonics of the square wave frequency, with the intensity decreasing at higher harmonics. That means that if you add those sine waves together and look at how the sum of those waves varies with time, you will see a square wave.

If you have a regular series of very short pulses in time, as shown at top in Figure 11-7, analyzing the range of the frequencies present in the signal over time yields a very interesting result—a series of sine waves spaced uniformly in frequency, separated by an




Figure 11-7. A series of ultrashort pulses separated uniformly in time (top) corresponds to a series of uniformly spaced frequencies that are graded in amplitude (bottom).

increment equal to the repetition rate of the laser pulses, shown at the bottom of the figure. This is called an optical frequency comb, and the shorter the pulses are, the wider the range of frequencies they span. Each frequency component is a continuous sine wave in a very narrow band of frequencies. For a series of very short pulses, the amplitude of the series of frequencies varies roughly as shown.

Frequency combs can be produced in a number of ways. For example, pulses from a Ti-sapphire laser mode-locked at hundreds of megahertz can be spread further to produce a supercontinuum, yielding a series of spectra lines separated by the laser's repetition rate. Other approaches include four-wave mixing and coupling a laser into a microresonator with a series of evenly spaced resonant frequencies. Frequency combs can be spread to span more than an octave (a factor of two in frequency) to aid measurements across their output range.

Frequency combs are very attractive for spectroscopy and measurement applications because they contain many evenly spaced spectral lines, promising very accurate measurements of time and wavelength.

11.5 EXTREME ULTRAVIOLET SOURCES

Laser operation becomes increasingly difficult as wavelength decreases and photon energy increases. One reason is that population inversions become harder to produce and excited-state lifetimes decrease with transition energy. Another is that the absorption of conventional optical materials increases with photon energy and even air becomes essentially opaque at wavelengths shorter than about 200 nm.

These changes reflect the wide ranges of photon energy in the band we call the *ultraviolet*, at wavelengths shorter than the 400-nm limit of human vision and extending to 10 nm, usually considered the start of the X-ray band and measured in terms of photon energy ranging from 3.10 to 124 electron volts. It is a factor of 40 between the ends of the band, compared with a factor of 1.75 between the 400- and 700-nm ends of the visible spectrum. (The eyes of birds and fish can see light over an octave of wavelength, from about 700 nm to beyond 350 nm in the ultraviolet.) Because of these differences, the ultraviolet usually is divided into several bands, as listed in Table 11-1. Three of the bands listed are standards recommended

Table 11-1. Selected definitions of ultraviolet bands. UVA, UVB, and UVC were defined by the International Standards Organization and reflect health hazards. Some commonly used names are not standardized, so their definitions may vary.

Name	Wavelengths (nm)	Photon energy (eV)	Notes
Ultraviolet A (UVA)	315–400	3.10–3.94	Long-wave UV, little absorption by air or ozone. It causes skin tanning
Ultraviolet B (UVB)	280–315	3.94–4.43	Medium-wave, largely absorbed by ozone layer. It causes sunburn and skin damage
Ultraviolet C (UVC)	100–280	4.43–12.4	Short-wave, germicidal, absorbed by ozone, shorter than 200 nm strongly absorbed by air
Near-ultraviolet (NUV)	300–400	3.10–4.13	Visible to some animals
Middle ultraviolet (MUV)	200–300	4.13–6.20	
Far ultraviolet (FUV)	122–200	6.20–12.4	
Vacuum ultraviolet (VUV)	10–200	6.20–124	Strongly absorbed by oxygen in air, often used only in vacuum
Extreme ultraviolet (EUV)	10–200	6.20–124	Ionizing radiation, completely absorbed by air
Deep ultraviolet	less than 300 nm		Widely used but not well defined

by the International Standards Organization for wavelengths from 100 to 400 nm; the other names are commonly used, but not formal standards.

A number of lasers have fundamental wavelengths or second through fifth harmonics at ultraviolet wavelengths longer than 200 nanometers. However, laser sources are very few and far between at shorter wavelengths variously called the *extreme ultraviolet* or *vacuum ultraviolet*, the latter because wavelengths much shorter than 200 nm are so strongly absorbed by air that experiments often are done in vacuum. Only two important gas lasers emit in that range, argon fluoride excimer lasers at 193 nm and excimer-like molecular fluorine lasers at 157 nm. Free-electron lasers can emit in the X-ray band, but as you will learn in Section 11.6, these are large

and complex instruments based on electron accelerators. This section focuses on other extreme ultraviolet sources.

11.5.1 Why Extreme Ultraviolet Sources are Hard

Extreme-ultraviolet lasers operate on electronic transitions, but they are very different than the electronic transitions of conventional lasers. Transitions at visible wavelengths involve electrons in the outer shell of an atom, which are shielded from the strong attraction of the atomic nucleus by electrons in the inner shells. Extreme-ultraviolet transitions involve electrons in inner shells, which are more closely bound to the nucleus. Inner-shell transition energies are tens or even hundreds of electron volts, not the 1 to 5 eV typical for outer-shell transitions, as Figure 11-8 shows for selenium. Inner-shell transition energies E typically are given in electron volts; you can convert them to wavelength in nanometers using the formula

$$\lambda(\text{nm}) = \frac{1240}{E(\text{eV})} \quad (11-1)$$

The high-energy/short-wavelength end of the extreme ultraviolet is not rigidly defined. Some specialists call wavelengths of 1 to 100 nm *extreme ultraviolet*. Others call the region from 1 to 10 nm (or sometimes from 1 to 100 nm) *soft X-rays*, so extreme-ultraviolet lasers in this range have also been called X-ray lasers. Wavelengths shorter than 1 nm are called *hard X-rays*. So far the only lasers to have operated at such wavelengths are the laser-pumped sources in Section 11.5.5 and the free-electron lasers in Section 11.6.

Whether you call 1–100 nm range extreme ultraviolet or soft X-rays, it behaves quite differently from visible light. It has enough energy to strip electrons from atoms or molecules, so it is called *ionizing radiation*, and can damage biomolecules and living cells. Soft X-rays can see through things that look opaque to our eyes, like soft tissue, but they are blocked by bone and other materials containing heavy elements. Soft X-rays also cannot penetrate through long lengths of air. Another important difference is that most materials have refractive index very close to one in the extreme ultraviolet, so they can neither reflect nor refract light well, so different types of optics are needed.

Electron dropping from a high energy level into a hole in this inner shell emits a soft x-ray photon




Figure 11-8. Electron falling into an inner shell of selenium emits an extreme-UV photon.

11.5.2 High Harmonic Generation

Focusing high-power pulses from visible or near-infrared lasers into certain gases can generate laser-like light in the extreme ultraviolet by producing high-order harmonics. This requires focusing near-infrared laser pulses to such high intensity that peak power is 10^{13} to 10^{16} watts per square centimeter, producing the electromagnetic fields strong enough to ionize gas atoms. The resulting free electrons

oscillate back and forth in the strong electromagnetic fields from the light pulse, generating high-order odd harmonics of the input laser frequency.

Steady refinements have pushed this approach to wavelengths well below 10 nm, corresponding to wavelengths beyond the 100th harmonic of the input laser pulse. High harmonic generation began with titanium–sapphire lasers, but pumping the gas with longer-wavelength pulses can produce higher-order harmonics because the electric fields of the pump light change more slowly at lower frequencies, giving them more time to accelerate electrons. The resulting high-frequency pulses are low in power, but because they contain a very wide range of frequencies, they can also be extremely short, down to the attosecond (10^{-18} second) regime, important for research.

11.5.3 Pulsed Discharge-Pumped Extreme-Ultraviolet Ion Lasers

Continuous laser action is very difficult at short wavelengths because excited-state lifetime declines with increasing photon energy. But high-voltage pulses can produce short-lived population inversions by stripping several electrons from gas atoms to produce highly excited ions. Experiments at Colorado State University produced Ar^{+8} ions which lased at 47 nm when free electrons recombined with the argon ions. But those lasers have yet to prove practical.

11.5.4 Laser-Produced Extreme-Ultraviolet Plasma Sources

Intense laser pulses focused onto suitable targets can generate incoherent light in the extreme ultraviolet, and this technology has emerged as the leading candidate for a new generation of electronic chip production. Decades of progress in semiconductor electronics have been driven by steady improvements in manufacturing that has reduced transistor sizes so that more could be squeezed onto a chip. Continuing that progress requires shorter-wavelength light sources to achieve higher resolution than the current generation of systems using 193-nm ArF lasers.

To achieve that higher capacity, a new generation of photolithographic systems has turned to plasmas produced by focusing intense pulses from a CO_2 laser onto it. The vaporized tin plasma emits brightly at 13.5 nanometers in the extreme

ultraviolet, a wavelength attractive for extremely high-resolution photolithography. Although a laser powers the EUV emission, the source itself is not a laser because the emission is spontaneous.

11.5.5 Laser-Pumped Extreme-Ultraviolet or X-Ray Lasers

Focusing even more intense laser pulses onto suitable targets also can produce stimulated emission in the extreme ultraviolet or X-ray bands. As with the incoherent EUV source described in Section 11.5.4, the pump pulses vaporize the solid target and ionize atoms by removing many outer-shell electrons. In this case, however, the system is designed to produce a population inversion on short-wavelength transitions, as free electrons drop into vacant inner electron shells.

First demonstrated in 1984 with selenium by Dennis Matthews at the Lawrence Livermore National Laboratory in California, this technique produces single-pass amplification by stimulated emission because resonator mirrors are not available for extreme-ultraviolet wavelengths. It is nonetheless considered to be a laser. Researchers have progressed to shorter wavelengths since the first demonstrations.

Figure 11-9 shows how the Livermore laser worked; later experiments are similar in concept. A short, intense pulse from a powerful laser built for fusion experiments was focused into a long, thin line on a thin foil of selenium. The pulses, lasting about a nanosecond and reaching peak powers of terawatts (10^{12} W), vaporized and ionized the selenium, producing a linear plasma. Electron recombination produced stimulated emission at 20.6 and 20.9 nm, with photons passing along the length of the thin plasma being amplified to form a beam without a resonant cavity. Further experiments have extended the technique to shorter wavelengths and improved results.

11.5.6 The Bomb-Driven X-ray Laser

During the 1980s, Livermore also conducted highly controversial experiments, which sought to use the burst of X-rays from a nuclear explosion to stimulate emission of a highly directional beam of short-wavelength X-rays. Edward Teller, the father of the American hydrogen bomb, envisioned this as a “third-generation” nuclear weapon that could focus its energy to destroy attacking nuclear




Figure 11-9. A small, hot plasma generated by a laser pulse emits an X-ray laser pulse. (Courtesy of Lawrence Livermore National Laboratory.)

missiles, and it became part of the “Star Wars” system of the 1980s. However, the project was eventually cancelled after further tests failed to support early claims.

11.6 FREE-ELECTRON LASERS

Free-electron lasers are a unique type in which the stimulated emission comes from a beam of “free” electrons, unattached to any atom, as they pass through an array of magnets. The magnets create a magnetic field that varies periodically along the path of the electrons in ways that coax the electrons to release some energy as light. The energy of the electrons and the shape and strength of the magnetic field determine the output wavelength. The beam from a free-electron laser behaves much like an ordinary laser of the same wavelength, although it is generated in an entirely different way.

Free-electron lasers can operate across a wide range of electromagnetic spectrum, from microwaves to X-rays, depending on the electron energy and the magnetic field. No single free-electron laser

can be tuned across the entire range, but individual free-electron lasers can be tuned across a wide range compared with other lasers that depend on fixed transitions. John M. J. Madey invented the free-electron laser in the 1970s, when he first demonstrated it in the infrared. It has now become the leading source of coherent X-rays.

11.6.1 Structure of a Free-Electron Laser

A free-electron laser has three essential components: a powerful electron accelerator that accelerates a beam of electrons to high speeds, an array of magnets called a *wiggler* or *undulator*, and a set of cavity optics. The type of cavity optics depends on the wavelength; where good reflectors are available, they form a resonator, but at short wavelengths where reflection is low, the beam is generated on a single pass.

Electron accelerators are massive but fairly standard pieces of scientific equipment; what matters in a free-electron laser is what happens once the electrons have been accelerated, as shown in Figure 11-10. It is easiest to visualize the magnetic field as an array of permanent magnets aligned in alternating directions, but electromagnets also can be used. If the north pole of the first magnet is above the electron beam, as shown in the figure, the south pole of the second magnet must be above the electron beam, then the north pole of the third, and so on.

The electron beam enters the magnet array at right at a slight angle above the axis of the laser. The array is called an *undulator*




Figure 11-10. Electron beam passing through a wiggler generates a laser beam in a free-electron laser. (Courtesy of the University of California at Santa Barbara Quantum Institute.)

or *wiggler* because the magnetic field inside it varies periodically along its length, so it bends electrons in the beam back and forth. The magnetic field from the first magnet bends the electron beam in one direction and then the opposite-polarity field from the second magnet bends it back in the other direction. This process repeats until the electron beam passes out the other end of the array.

Every time the magnetic field bends the path of the electrons, the electrons release some energy as light. The electrons reabsorb most of the light when they are bent in the other direction. However, the electrons can amplify light if its wavelength λ is close to a resonance that depends on the velocity v of the electrons along the laser axis, the speed of light c , and the period of the wiggler magnet p , defined as the distance from one north pole to the next along the array:

$$\lambda = \frac{p}{2[1 - (v^2/c^2)]} \quad (11-2)$$

The $[1 - (v^2/c^2)]$ term also appears in Einstein's theory of special relativity and is needed because the electrons are moving at relativistic speeds.

Equation 11-2 highlights the basic physics of a free-electron laser, but if you want to calculate the laser wavelength λ , it is more practical to use a different version of the formula based on the wiggler period, the electron's accelerated energy E (measured in million electron volts), and its rest mass, 0.511 million electron volts:

$$\lambda = \frac{0.131p}{(0.511 + E)^2} \quad (11-3)$$

This tells us that the wavelength gets shorter as the magnet period decreases and electron energy increases.

11.6.2 Tunability of Free-Electron Lasers

Free-electron laser wavelength depends on two quantities that may be adjusted: the period of the wiggler magnet and the energy of the electrons passing through it. Changing either the magnetic field or the electron energy can tune the wavelength.

So far, we have considered the wiggler magnet as a stack of permanent magnets with alternating polarity, but the key element

of the wiggler is not the physical magnet but the magnetic field, which varies periodically along its length. That means you do not need a physical magnet, but just a periodic magnetic field that could come from any source, including a light wave. Thus, a light wave could create a wiggler field for a free-electron laser with periodic variations down to the wavelength of light, much shorter than the smallest magnets you could build. With such a short period, you could produce very short wavelengths from a free-electron laser.

Varying the energy of injected electrons also changes the free-electron laser, with high-energy electrons producing shorter-wavelength emission. Electron energy typically is easier to vary than the spacing of magnets.

11.6.3 Types of Free-Electron Lasers

Although the physical principles of free-electron lasers are the same across their wide operating range, the actual devices differ considerably, and individual devices are built to operate over limited parts of the spectrum.

Long-wavelength free-electron lasers operate in the Raman or “collective” regime, where the interactions are best described as among many particles. In this case, the electron energy is relatively low (under 0.005 electron volt), the current densities are relatively high, and the output wavelengths are longer than around 100 μm . Gain is high and output power can be high.

Interactions involving individual particles are more important at shorter wavelengths, which are generated using electrons with higher individual energies. This requires bigger accelerators, although they produce weaker currents. Important examples are:

- Vanderbilt University’s free-electron laser in Tennessee, which emits pulses from 2.1 to 9.8 μm , but was shut in 2008.
- The free-electron laser at the Thomas Jefferson National Accelerator Facility in Newport News, Virginia, produces continuous power to 14 kW in the infrared, making it the world’s most powerful tunable infrared laser. It can operate from the terahertz band to the vacuum ultraviolet.
- The Duke University Free-Electron Laser Laboratory in North Carolina runs on a storage ring capable of handling electrons to 1.2 billion electron volts and emits from 400 to 190 nm in the ultraviolet.

- The LINAC Coherent Light Source at the SLAC National Accelerator Laboratory in California was the world's first hard X-ray laser, firing up to 120 pulses per second that last femtoseconds, allowing ultrafast measurements of wavelengths from 10 nm to less than 0.1 nm. An upgrade is in progress.
- The Free-electron Laser at Hamburg operates at DESY, the German Electron-Synchrotron Laboratory in Hamburg. It uses billion-electron-volt electrons to produce femtosecond pulses at wavelengths from 6.5 to 47 nm.
- The SACLA X-ray free-electron laser at Riken in Japan, firing 30 pulses per second at 0.3 to 0.08 nm.
- The European X-ray free-electron laser in Germany, which first produced X-rays in 2017. It will be the world's largest and brightest X-ray free-electron laser emitting femtosecond pulses at 3.1 to 0.1 nm.

The new generation of X-ray free-electron lasers has given scientists new insight into atomic-scale chemical and biological processes on femtosecond time scales. Their strength is the immense energy they can apply to small areas in space and small intervals in time. Yet that immense sophistication comes at a high cost, and for the foreseeable future, X-ray free-electron lasers will remain research tools available only in a few specialized facilities.

11.7 WHAT HAVE WE LEARNED?

- Complex organic dyes in liquid solutions fluoresce across a range of wavelengths, making them useful in tunable lasers.
- Laser cavities developed for tunable dye lasers can be used with other tunable lasers.
- Organic dye lasers require optical pumping with light from an external laser or a flashlamp.
- Optical parametric light sources are broadly tunable and based on nonlinear three-wave mixing involving three photons. The sum of the signal and idler frequencies equals the pump frequency.
- OPAs take energy from a powerful pump beam to amplify a signal beam and produce an idler beam. They can be tuned between 200 nm and 16 μm .
- Optical parametric oscillators operate in an optical cavity with a resonance at the signal frequency.

- Sending an intense beam through a highly nonlinear medium spreads out the range of wavelengths to produce a supercontinuum.
- A frequency comb is a series of uniformly spaced frequencies produced by a series of short pulses of light at regularly intervals. Each is the Fourier transform of the other.
- Frequency combs are very useful for spectroscopy and measurement.
- Extreme-ultraviolet sources emit wavelengths shorter than about 200 nm, where the atmosphere is nearly opaque.
- High-order harmonics of visible or near-infrared pulses, generated in gases exposed to extreme laser intensities, can reach wavelengths less than 10 nm.
- Intense pulses from CO₂ lasers can vaporize tin to produce extreme ultraviolet pulses at 13.5 nm for photolithography.
- Laser ionization of thin films of metal can remove more than a dozen electrons from metal atoms, and when those atoms capture electrons to fill the holes in those energy levels, they can produce stimulated emission deep in the extreme ultraviolet.
- Free-electron lasers extract energy from a beam of electrons passing through a magnetic field that varies regularly in strength. They can be tuned across a broad range of wavelengths by changing the electron energy or the period of the magnetic field.
- Free-electron lasers draw on well-developed technology for particle accelerators.
- Big free-electron lasers have become important sources of X-ray laser light.

WHAT'S NEXT?

The final three chapters explore the applications of lasers first in low-power applications, then in high-power applications, and, finally, in research.

QUIZ FOR CHAPTER 11

1. What property gives organic dye lasers their wavelength tunability?
 - a. The many vibrational sublevels of electronic transitions
 - b. The use of organic liquids as solvents

- c. The high gain of the optically pumped dyes
 - d. Photodissociation of the dyes under intense pump light
 - e. The liquid nature of the laser medium
2. Typical tuning bandwidth of an individual laser dye is
 - a. 1 nm
 - b. 5 nm
 - c. 5–10 nm
 - d. 10–20 nm
 - e. 20–40 nm
 3. Typical linewidth of the output of a dye laser is
 - a. 1 nm
 - b. 0.01 nm
 - c. 1 MHz
 - d. 0.1 MHz
 - e. determined by the nature of the laser cavity and tuning optics
 4. An OPA is pumped by a 1064-nm solid-state laser and has a signal wavelength of 1300 nm. What is the wavelength of the idler beam?
 - a. 236 nm
 - b. 764 nm
 - c. 1600 nm
 - d. 2364 nm
 - e. 5845 nm
 5. What wavelength is an optical parametric oscillator tuned to?
 - a. Pump wavelength
 - b. Signal wavelength
 - c. Idler wavelength
 - d. Sum of all three wavelengths
 - e. Sum of all three frequencies
 6. Which light source has the broadest tuning range?
 - a. Tunable dye laser
 - b. Optical parametric oscillator
 - c. Optical parametric amplifier
 - d. Titanium–sapphire laser
 - e. Ytterbium-doped fiber laser
 7. What light source has been developed for photolithography at shorter wavelengths than are possible with 193-nm argon fluoride lasers?
 - a. Carbon dioxide lasers
 - b. KrF lasers
 - c. High harmonic generators

- d. Tin metal vaporized by pulses from CO₂ lasers
 - e. Bomb-driven X-ray lasers
8. What emits light in a free-electron laser?
- a. Free electrons in a beam passing through an array of magnets
 - b. Electrons conducting current in the coil of an electromagnets
 - c. Photons passing through a high-energy particle accelerator
 - d. Gas struck by electrons emerging from a particle accelerator
 - e. Atoms moving at the speed of light
9. What bands can a free-electron laser emit?
- a. X-rays
 - b. Ultraviolet
 - c. Infrared
 - d. All of the above
 - e. None of the above
10. How is energy extracted from free electrons in a free-electron laser?
- a. Their paths are bent back and forth when they pass through a periodic magnetic field.
 - b. The free electrons induce currents in the magnets they pass, and the magnets emit light.
 - c. The electrons release energy after they evaporate from the magnets.
 - d. The energy is extracted when electrons bounce off a resonator mirror and lose energy.
 - e. The electrons transfer energy to helium atoms in the air, which then transfer the energy to another species.

LOW-POWER LASER APPLICATIONS

ABOUT THIS CHAPTER

The final three chapters of this book cover laser applications, which are many and diverse. To keep the topic manageable, we will break them into three broad groups: those requiring little power (such as playing CDs, DVDs, or Blu-Ray disks), those requiring high power (such as cutting and welding), and those in scientific research, such as studying the properties of atoms and molecules. The division between low and high power is somewhat arbitrary, based on how the lasers affect the materials they illuminate. The intent is not to list everything, but to list the most important, interesting, and illustrative applications to help you understand laser applications in general.

This chapter covers applications that require low laser power, typically well under a watt. A laser with such little power does not have a dramatic effect on the objects it strikes, but it may change materials designed to be light-sensitive, such as exposing photographic film, and could damage the human eye. As we will see, there are many types of low-power laser applications. In some, the laser is little more than a high-performance light bulb, but others depend on special features of laser light, such as coherence or tight beam collimation. In most cases, the laser light transmits or processes information or makes measurements, such as playing video from a DVD, displaying a holographic image, transmitting

data through a fiber-optic network, drawing a straight line in space for surveying, or printing information on a laser printer.

12.1 ADVANTAGES OF LASER LIGHT

Theodore Maiman's first laser made headlines in the summer of 1960. As word of the laser spread, it seemed that every scientist and engineer wanted their own laser. Many had no clear purpose in mind, but the laser seemed like a neat toy, and curiosity lurks deep in the souls of scientists and engineers. They built their own lasers or bought lasers from a handful of little companies that sprang up to make them. Then they zapped just about everything they could find. They informally measured laser power in "gillettes"—how many razor blades a laser pulse could pierce. It was a fertile time for experiments, new ideas, and accidental discoveries, but, inevitably, many of the experiments failed and for a while the laser became "a solution looking for a problem," a description that Maiman's assistant Irnee D'Haenens coined at the dawn of the laser age.

Much has changed since the initial wave of laser enthusiasm. Laser technology has advanced dramatically, as have other technologies, particularly digital electronics, control systems, computing, and sensors. Lasers have been combined with other new technologies to make new things possible. Let us look at how low-power lasers have come to play crucial roles in consumer products, the Internet, and the emerging technology of self-driving cars, and then consider what makes lasers cost-effective solutions to important technological problems.

12.1.1 What Is "Low" Power?

The difference between "low" and "high" power from a laser is subjective in the sense that it depends on what you are trying to do. Generally reading, writing, or transmitting information requires powers from the milliwatt level to a fraction of a watt, and they are considered "low" power. Cutting, welding, drilling, or otherwise altering materials, such as metals, plastics, glass, or performing surgery generally is considered high power.

However, it is not as simple as defining a power level or a change on the surface as the dividing line. A milliwatt beam can expose photographic film, and modest milliwatt powers can write

data on the surface of an optical disk by forming tiny holes. Because those materials are made to be sensitive to low levels of light, their applications are considered low power.

The definition also differs in applications involving the human eye. The light-sensitive retina at the back of the eye is extremely vulnerable to damage from laser beams. The eye focuses the parallel rays in a laser beam onto a tiny spot on the retina, so a few milliwatts can be focused to an intensity high enough to injure the retina if you stare into the beam—just as staring at the sun can injure the retina. You will learn more about laser eye safety in Appendix A, but the basic rule to remember even for low-power lasers is not to look into the beam.

12.1.2 Lasers in Consumer Audio and Video

The laser emerged at a time when the cutting edge in consumer audio and video was stereo phonograph records, four-track reel-to-reel tape, and analog broadcast of color television. Compact audio cassettes and eight-track audio tape emerged in the early 1960s as alternatives to bulky phonograph records for consumer audio, but the electronics industry did not have a viable alternative for home video players. A number of companies began considering home video disk players based on lasers or electronic media.

The laser's attraction was its ability to focus a beam onto a spot roughly a wavelength wide to read information from spiral tracks pressed into the disk. The laser made a small spot possible by emitting coherent light of essentially a single wavelength, well matched so that it was all focused onto a smaller spot than light from any other source. Only milliwatts of laser power were needed to read the disks, and red helium–neon gas lasers could readily deliver that power. Engineers spent years developing systems that could record and play back about two hours of color video on a two-sided 12-inch (30-cm) “LaserDisc.” That was long enough to hold most movies, which entertainment companies thought would be the big market. After considerable haggling, the media company MCA and the electronics company Phillips merged their efforts, developed a standard format, and haggled laser manufacturers to cut the price of a helium–neon laser from \$100 to \$25 by about the time they introduced the first players in 1978.

By then, the first video cassette recorders (VCRs) had reached the market, and media-savvy consumers had begun recording

television programs on tape. Rental movies soon followed. Few showed much interest in buying laser video disks, although small markets emerged in Japan and among affluent videophiles. VCRs captured the mass market. RCA introduced a capacitive electronic video disks player in 1981, but it became an expensive failure.

The real breakthrough for the optical disk was the introduction of the audio compact disk in 1982 by Phillips and Sony. A series of breakthroughs in semiconductor lasers during the 1970s made 780-nm diode lasers practical and inexpensive light sources for digital audio players that could store over an hour of music on a 12-cm reflective disk. The disks were much smaller than phonograph records, did not wear out or scratch easily, and gave crisp, clear sound without the background hiss of tapes. Compact disks soon became the audio medium of choice.

Several factors helped laser audio players succeed. The technology developed for laser video disks could easily be transferred to the smaller audio disks. By 1980, diode lasers and digital electronics had become mature products, so prices dropped as manufacturing volume increased. The audio market also was ripe for a change, and most consumers found CD audio quality an improvement.

Those strengths helped propel the DVD to success in digital video in the 1990s. By improving optics, using new digital electronics, and switching to a red 650-nm diode laser, it shrank the size of data spots so much that it could squeeze a whole movie onto a 12-cm disk. The next-generation Blu-Ray player was able to squeeze a whole HD movie onto a 12-cm disk by further improvements and shifting to violet diode lasers. However, by the time it reached the market in 2006, digital delivery of video over the Internet was becoming competitive, leading to rapid growth of streaming video.

12.1.3 Lasers in Communications and the Internet

The idea of optical communications came before the laser. The amount of information that waves can carry increases with their frequency. Light waves have frequencies tens of thousands of times higher than the gigahertz frequencies of microwaves, giving light tens of thousands of times more information capacity. As soon as communication engineers had lasers to test, they started sending the beams through the atmosphere—and received a rude shock. The air is not as clear as we like to think; fog, haze, smoke, clouds,

precipitation, and air turbulence get in the way of the beam. So engineers tried to build hollow light pipes that would keep the path clear for the light.

Hollow light pipes did not work well, but in 1966 Charles Kao came up with a different approach—sending light through hair-thin optical fibers. Then the Corning Glass Works invented a way to purify glass fibers so they were incredibly pure and amazingly transparent. What they needed then was a tiny light source that could be turned off and on at very high speeds to send signals. LEDs worked over short distances, but the pure single wavelength output of a diode laser could be switched on and off even faster and was much more directional, so it fit better with small-core fibers, which caused less distortion of high-speed signals.

The first fiber-optic communication systems carried hundreds of digitized phone calls simultaneously through about 10 km of fiber. Rapid improvements in laser and optical technology steadily increased data rates. The invention of fiber amplifiers made it possible to boost signal strength every 50 to 100 km on land and under the sea. Narrow-line lasers and precision optical filters sliced the transmission spectrum into up to 100 separate bands, each of which can now carry up to 100 gigabits per second. Add it all up and a single fiber can carry over 10 trillion bits per second, the equivalent of roughly 20 two-layer Blu-Ray disks a second, up to 10,000 km, as described in Section 9.5.1. And system designers are crying for more capacity to carry the signals needed by autonomous cars, virtual reality, and other emerging technologies.

Laser technology has evolved to fill the changing needs of communication systems. The first fiber-optic systems used diode lasers at 810 nm, but fiber researchers found they could transmit more bits per second at 1300 nm. So laser developers turned to new semiconductors to make lasers for that wavelength. When the best fiber amplifiers turned out to be at 1550 nm, lasers were developed for that wavelength as well.

The rapid progress in fiber-optic communications has combined with the dramatic advances in digital electronics to revolutionize our information society. Since 1987, the capacity of fiber-optic transmission has kept pace with the Moore's law advancement in digital electronic circuits. Lasers and fiber optics have carried the information; digital electronics have processed it. Those two complementary technologies joined to transform global

communications and information handling and, in doing so, have changed the world.

12.1.4 Laser Radar for Self-Driving Cars

An important emerging new market for lasers is laser radars or *lidars* for autonomous cars. The principle of lidar is the same as microwave radar—a laser transmitter fires a pulse of light at an object, and a receiver detects the return signal to measure how long it takes for light reflected from the object to return. Light takes 300,000 km/second, so if you fired a pulse of light through the air and detected a reflection in 1 microsecond, the pulse would have traveled 3.3 km in that time, which would be twice the distance to the object 1.65 km away.

Self-driving cars need more information than the distance to one object. They need a three-dimensional map of the space around them, to help the car's computer steer it so that it does not collide with other cars, buildings, trees, light poles, pedestrians, bicycles, or wildlife. The lidars fire rapidly, measuring the direction of the point and the distance to whatever reflected it, building up a 3D cloud of points on nearby objects, such as shown in Figure 12-1. They repeatedly scan the local area, so they can identify objects and track their motion.

The first generation of car lidars use inexpensive diode lasers emitting pulses at 905 nm, which can be detected by low-cost silicon detectors. However, those lasers can fire only low-power pulses to avoid risk to the eyes of nearby pedestrians. Those low-power pulses limit the lidars to detecting objects no more than 30 or 40 feet away, and that is a serious limitation because cars moving at highway speeds can travel 33 m in a second, and it takes longer than that for the lidar and the car's computer to identify objects and stop the car.

However, a new generation of car lidars are being developed that use lasers operating at 1550 nm using diode lasers and infrared detectors developed for telecommunications. The advantage of that wavelength, you learned in Section 8.4.4, is that it cannot reach the retina, the part of the eye most sensitive to optical damage. Laser safety rules recognize that 1550 nm light cannot reach the retina and allow output power 40 times higher than permissible at 905 nm, which can stretch lidar range to 200 m, a range that car




Figure 12-1. Point map of a city street showing streets, pedestrians, cars, and buildings. Image in color in electronic edition. (Courtesy of Luminar.)

makers believe should be adequate for the car to recognize a hazard and stop safely.

12.1.5 Advantages of Lasers as Light Sources

These three examples show how lasers solve problems. Think of lasers as better light bulbs. Their light is coherent, with waves marching along like soldiers on parade. Most lasers last longer than and are far more efficient than the old-fashioned incandescent lamps we think of as “light bulbs.” Another advantage is that laser light is better controlled. It can be concentrated in a narrow, monochromatic beam and focused onto a tiny spot to read or write data. Laser light can be modulated very quickly and very effectively. Modern communication systems do not just turn the laser light off and on; they also shift its phase and polarization to multiply the number of bits per second the light can carry through an optical fiber. Lasers can be tailored to emit a very narrow spectrum of light, so that hundred separate lasers can transmit separate signals through a single optical fiber, adding up to 10 trillion bits per second. Lasers can fire millions of pulses a second to measure the position of nearby objects tens of times a second, keeping track of their changing positions from a moving car. Table 12-1 lists some major advantages of lasers and some important limitations.

Table 12-1. Laser advantages and disadvantages.

Advantages	Limitations
Well-controlled light	Do not emit white light
Concentrate light in a beam	Beam poorly suited for illumination
Precisely focusable to high power density	Coherent light shows speckle effects
Beams do not press on objects they contact	Efficient lasers not available at many wavelengths
Emit coherent light	Available wavelengths may not match needs
Emit monochromatic light	Beam may pose eye hazard (depends on power and wavelength)
Long lifetime	
High efficiency in generating light	
Very small size of diode lasers	
Can produce very short pulses	
Can emit steady beams	
Low cost of some types	

You can get a feeling for these advantages and disadvantages by comparing a laser to a light bulb. A light bulb illuminates a room with uniform white light; a laser beam illuminates a single spot with a single color. If you want to illuminate the whole room, you want a light bulb. If you want to illuminate a single spot, you want a laser. If you want a bright searchlight or spotlight, either a laser or a light bulb might work, depending on your application.

In many cases there are tradeoffs between lasers and light bulbs. For example, BMW has developed laser-based automotive headlights to provide long-range illumination for drivers, but they do not aim raw laser beams along the road. Instead, bright blue diode lasers illuminate phosphors, and optics focus incoherent glow from the phosphors along the road. This provides good color balance and avoids excessive brightness and the speckle patterns produced by coherent light. But the laser headlights are much more expensive than conventional types.

More is coming as technology advances, both for lasers and for lamps. LED bulbs turn more electrical power into light than incandescent bulbs, and laser efficiency also is increasing. Sophisticated control systems make once-fussy types of lasers and laser-like sources, such as optical parametric oscillators, much easier to operate than they were a decade or two ago. Physicists have learned how to make laser pulses shorter and shorter, so they contain only one or two waves. Many jobs now require lasers. The rest of this chapter will describe some of the many uses of low-power lasers.

12.2 READING WITH LASERS

Low-power lasers are widely used to read printed symbols. The laser spot scans across a surface, and a detector measures changes in the reflected light as the beam moves. Computers decode the pattern of reflected light to interpret the symbols.

In principle, lasers can read any pattern printed so that the symbols reflect different amounts of light at the laser wavelength. That means you have to pick the color of your ink carefully. Red ink reflects red light, so a red laser beam cannot read red ink on white paper because both reflect the laser beam. A milk company in the Boston area learned this the hard way when they printed milk cartons with red ink, making the striped product codes invisible to the red laser beam used to read them.

In principle, this scanning system can be used to read any suitable printed symbols, including text, as long as suitable software is available to decode the changes in reflected light. In practice, the main applications are reading symbols designed especially for use with laser scanners. The main examples are bar codes consisting of stripes of different widths and two-dimensional codes that are patterns of light and dark zones. Laser scanners are not the only way to read these symbols; your cell phone camera can read 2D bar codes.

12.2.1 Bar-Code Scanners

Most bar codes used with laser readers are variations on the Universal Product Code originally developed for automated checkout in supermarkets. That system was developed to read codes printed on food packages moved quickly past a scanning window at any angle, although some care is needed to make sure the code is visible to the scanner.

Figure 12-2 illustrates the workings of an early laser supermarket scanner. A beam deflector—here, a rotating mirror—scans the laser beam repeatedly in a carefully controlled pattern across the zone where the bar code appears. The optical system collects reflected light, filters out light other than the laser wavelength, and focuses the remaining light onto a detector, which generates an electrical signal proportional to the amount of light reaching it. That raw signal then is processed to interpret the information encoded by the bar code.

Our diagram does not show other lighting, but it cannot be ignored in the real world. In a supermarket, for example, bright overhead lights shine down onto the checkout counter, illuminating packages at the same time as the laser. However, the laser light is concentrated at a single red wavelength, and the room light is spread across the visible spectrum. Covering the sensor with a filter that transmits the laser wavelength but blocks other ambient light helps make the bar code readable without interference.

Supermarkets were the first to adopt laser scanners for automated checkout after the Universal Product Code was developed in the 1970s. Now, most retail stores use similar systems. Supermarkets often install the laser system under the counter, with the beam emerging through a window to illuminate packages as they pass along the counter. Most other stores use hand-held readers equipped with red lasers or LEDs that clerks scan over the bar code




Figure 12-2. A laser bar-code reader.

on the package. The bar code identifies the product and the size of the package; the store's computer then looks up that item and adds its price to your bill. Supermarkets print special bar codes that indicate the price and weight of meat and produce packaged at the store.

The UPC symbol originally was designed to be read by a 632.8-nm helium–neon laser, but modern readers use red diode lasers close to that wavelength. Variations on the pattern are now standard on many products including books and magazines, and on many airline luggage tags. Other kinds of bar codes are used for other applications, such as tracking packages and checking out library books.

12.2.2 One- and Two-Dimensional Bar Codes

The UPC code and many others are one-dimensional codes made of bars of equal height that encode information by their spacing and width. Some other bar codes have uniform spacing but differ in



Figure 12-3. (A) Two-dimensional QR code for this book's website (<http://www.understandinglasers.com>) (B) A one-dimensional UPC symbol from a magazine. (C) Maxicode for this book's website (<http://www.understandinglasers.com>).

height, notably the code the US Postal Services uses for letters and magazines. Typically, these codes carry a limited amount of information, such as a series of digits that identify a product or a postal code.

Two-dimensional codes are coded in both directions, like the familiar QR code designed to be read by smartphones. They can carry considerably more information than a simple bar code. Figure 12-3 compares a standard one-dimensional UPC symbol from a magazine with a 2D QR code which encodes the website for this book. Some two-dimensional codes do not use bar patterns; Figure 12-3C shows the Maxicode developed and put into the public domain by United Parcel Service.

12.2.3 Other Laser Readers

Lasers can read many other types of printed symbols. For example, some voting machines use laser readers to interpret black marks put on ballots by voters. Lasers also are used to scan text for optical character recognition.

12.3 OPTICAL DISKS AND DATA STORAGE

The biggest single application of lasers in terms of the number of lasers sold has been for playing and/or recording information on optical disks. This covers audio CDs, data CD-ROMs, video and data DVDs, recordable CDs and DVDs, and high-definition/data Blu-Ray disks. Billions of diode lasers have been in use since Philips and Sony introduced the first diode-laser players for digital audio, and mass production has brought prices down for the lasers used in all three standards, 780 nm for CD drives, 650 nm for DVD drives, and 405 nm for Blu-Ray.

Optical drives played an important role in the personal computer revolution by providing a format for distributing data and software with much higher capacity than the floppy disks used in early PCs. However, sales of optical drives for audio, video, and data have declined in recent years as users have shifted to streaming audio and video, downloading software and data from the Internet, and most recently cloud computing. Optical disks are not obsolete, but they have become something of a niche market or an option for people with limited Internet bandwidth. Nonetheless, they are an important example of laser applications.

12.3.1 Optical Disk Technology

The common concept at the root of all optical disk technology is shown in Figure 12-4. An optical head focuses light from a semiconductor laser onto the surface of a rapidly spinning disk. The head moves back and forth as the disk spins beneath it.

The original concept was to sell prerecorded digital audio disks (CDs) pressed from master recordings, as was done for vinyl phonograph records. In the optical system, laser would record digital data by burning small pits in a light-sensitive material on a master disk, a mold would be made from the master, and read-only CDs would be pressed from the molds using a process similar to but much more precise than pressing vinyl phonograph records. Recording companies distributed and sold CDs through the same channels as phonograph records. A compatible data-recording version, called CD-ROM, was soon developed to distribute information and software for computers. The early read-only audio and data players were illuminated by a milliwatt-class 780-nm diode laser and the reflected light was focused onto a detector in the laser head, which decoded the reflected light to yield audio or data signals.




Figure 12-4. Optical disk system.

The next step was the development of a standard for a writable version of the CD that could be played back on an audio CD player or CD-ROM drive. This required both a new disk material, which could be written as well as read by a diode laser, and drives with two modes—a low-power mode for reading disks, and a higher-power laser mode that could write as well as read. Two lasers could also be used, one for reading, the other for writing. An audio CD can store up to 72 minutes of music; a digital CD-ROM or writable disk can store about 700 megabytes.

12.3.2 DVD and Blu-Ray

Optical disk storage capacity depends on how densely data spots can be written and read. The density is inversely proportional to the square of the wavelength; reducing wavelength by a factor of two would multiply storage capacity on a 12-cm disk by

a factor of four. The development of shorter-wavelength diode lasers and of optics that can focus the beam to tighter spots has led to two generations of optical disks with greater storage density.

Red diode lasers emitting at 650-nm are used in DVDs, which originally were developed to distributed prerecorded digital video and later expanded to include computer data. (DVD originally stood for *digital video disk*, but when the writable version was developed for data storage, it became *digital versatile disk*.) DVDs record data digitally, but the video they store is a digital version of standard-definition analog video that ceased broadcasting in the USA at the end of 2010. DVD videos look much better than broadcasts or VCR recordings. DVDs can store about two hours of video; the nominal capacity of a writable DVD for data storage is listed as 4.7 gigabytes.

High-definition digital television requires squeezing several times more data on a single disk. To meet that requirement on a 12-cm disk, developers shifted to 405-nm diode lasers for the third-generation Blu-Ray disk. They can record 25 gigabytes on one layer and twice that on a dual-layer disk.

12.3.3 Writing on Optical Disks

Writing on optical disks requires the disk to be covered with a light-sensitive material. Blank CDs and DVDs are manufactured with an internal layer of light-sensitive organic dye that changes its reflectivity after being heated by a laser beam. Read-only drives lack the laser power to heat the material. The higher-power lasers in read-write drives are operated at high power to write and at lower power to read.

A number of writable optical disk formats have been developed over the years. The current types of disks that allow writing only once, called CD-R (CD-Recordable) and DVD-R disks, are the descendants of disks called WORMs, for write-once, read many. Early WORM disks came in a variety of sizes and formats, but computer users and manufacturers settled on the 12-cm CD size and retained that size for DVDs. Recordable Blu-Ray disks also are available.

Rewritable optical disks have also been developed, using other materials that change their phase when heated by a laser beam, changing their reflectivity in ways that can be erased and

overwritten many times. A second type of rewritable optical disks used a combination of magnetic and optical technology; called magneto-optical, they require special drives and were never widely used.

12.4 LASER PRINTING

The first laser printers introduced in the 1970s were based on gas lasers and designed for use on mainframe computers. Their attraction was the ability to print at much higher speeds and larger volumes than earlier impact printers. Much less expensive laser printers based on diode lasers were introduced in the 1980s for use with personal computers, and laser printers have been coming down in price ever since. Once limited to black and white, laser printers now can also print in color.

Laser printing relies on a process called *raster scanning*, illustrated in Figure 12-5. Input information controls the power emitted by the laser at left, turning the light off and on as the beam is scanned back and forth, producing an image line by line as the page moves up, so the written points eventually cover the whole page. Raster scanning earlier was used to produce images on cathode-ray tube televisions by scanning an electron beam across the screen.

Similar laser printing processes are used to print computer output and to write plates used for printing newspapers and magazines.

12.4.1 Laser Computer Printers

Laser printers for computer output do not write directly on paper. Instead, they write on a rotating drum coated with a light-sensitive

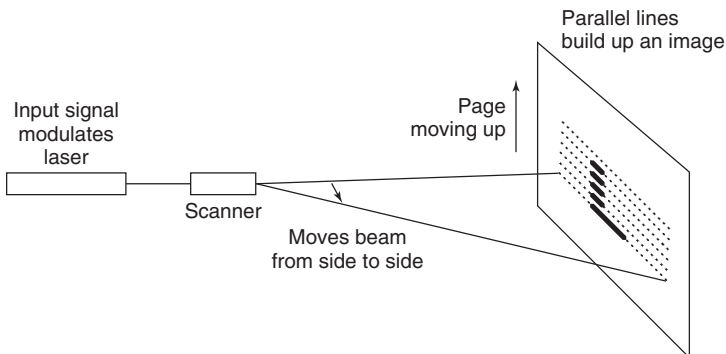


Figure 12-5. Raster scanning a laser beam across a page.

material that carries a static electrical charge, like those used in photocopiers. A copier illuminates the page to be copied with a bright light and focuses the reflected light onto the rotating drum, and the light discharges the electrical charge in bright areas. The drum then rotates through a region where parts of the drum that still carry an electric charge attract a dark fine-grained material called a “toner,” forming a positive image of the original page, with toner in the areas where the original was dark. The drum then transfers the toner image to a blank sheet of paper, reproducing the original page. A final stage bonds the fine dark toner to the paper.

A laser printer contains a similar drum arrangement, as shown in Figure 12-6, in which a scanning laser beam writes the page directly onto the drum. In this case, the laser-illuminated areas lose their charge, so they do not pick up the toner that produces the dark areas in the final copy. This means that the laser actually writes the white areas, not the dark areas we think of as the printed area in the final image on the paper.

With typical resolution of 600 to 1200 dots per inch, the output of a black-and-white laser printer looks almost as good as a




Figure 12-6. Elements of a laser printer.

conventionally printed page. In fact, the copy for many books, including all text and art, often is produced using a high-resolution laser printer.

Laser printers once printed only black toner on white paper. Color versions are now readily available, which work by making three passes using different colored toners, one for each of the three colors needed to generate the full range of colors for reflected images. Because they are more complex, color printers are more expensive than black-and-white ones, and comparable models have somewhat lower resolution.

12.4.2 Laser Prepress Equipment

Laser systems also play important roles behind the scenes in the publishing industry. Computer-to-plate or computer-to-press systems use lasers to write final versions of pages to be printed. The laser may write on photographic film that is used to transfer an image to the plates used in actual printing or may write directly on the printing plates.

Another laser application in printing is to prepare color photographs for printing in magazines or books. Color printing processes create full-color images by running pages through a press three times, each time adding one of the three colors needed to build up a full-color image. Printing that way requires a process called color separation, which analyzes the colors in the original photograph and uses that information to prepare separate plates for the three colored inks that together make a color image. Lasers can scan the original picture to produce those color separations.

12.5 LASERS IN FIBER-OPTIC COMMUNICATIONS

As you learned in Section 12.1.2, lasers generate the signals transmitted through the fiber-optic backbone of the Internet, which has had a tremendous impact on our society. Diode lasers are a natural match with fiber optics because the small size of the diode's output aperture can be matched to small light-guiding core in the fiber, and virtually all lasers used in optical communications are diodes.

Diode lasers also have other advantages. They can be modulated at high speeds and generate enough light to travel through 50 to 100 km of high-grade fiber. Laser output can be modulated

separately in amplitude, in phase, and in polarization, multiplying the amount of information a laser can send through a fiber. Lasers can be made to emit a very pure single color of light; so many lasers can send separate signals through the same fiber at once. Let us look briefly at how fiber-optic communications work and why these properties are important.

12.5.1 How Fiber-Optic Communications Works

Figure 12-7 shows a simplified view of a fiber-optic data link. The input signal arrives in an electronic form at a transmitter, which amplifies that signal and uses it to modulate the output of a diode laser. The simplest type of modulation is pulsing the laser output off and on, producing light carrying a digital code of 0s and 1s. That light then is coupled into the light-guiding core of an optical fiber, which delivers the signal to an optical detector at its destination. The detector converts the optical signal into electronic form and electronics in the receiver amplify the electronic signal and translate it into standard formats.

Lasers are packaged in transmitters that deliver a modulated optical signal to the transmission fiber. Lasers can be directly modulated by changing the drive current at rates to 1 or 2.5 gigabits (billion bits) per second. At higher data rates, or when the signals

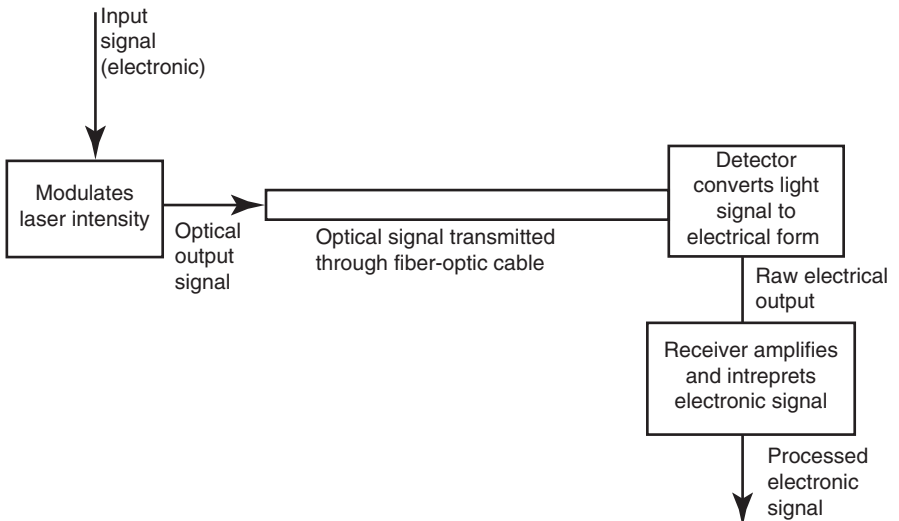


Figure 12-7. Fiber-optic data link.

must travel very long distances, the lasers generate steady beams that are modulated in intensity by a separate modulator before the beam enters the fiber.

Note that Figure 12-7 shows only one *link* that connects two points in the global fiber-optic network. For example, it may carry voice, video and data between your home and a local Internet carrier's operation center. There your data may be converted into electronic form to be merged with other data to and from other homes and converted back to a higher-speed optical signal being sent around the world, like many small streams merging into a large river. And so the process goes, like car traffic on small roads merging onto a main street, which feeds into a superhighway carrying much more traffic. If you are looking at a website at an overseas university, your signals may go halfway around the world on fibers that may carry up to 10 trillion bits per second. Then the signals are split apart at the other end.

12.5.2 Wavelength-Division Multiplexing

As little data streams merge together into bigger ones, the data rate increases. Different parts of the network carry data at different speeds. Your home network may send 100 megabits per second through an optical fiber to a node where your carrier combines it with signals from your neighbors to make 1 gigabit per second, and from there the signals may be merged again into a 10 gigabit per second stream, and ultimately with yet more signals to make 100 gigabits per second.

Different types of laser transmitters are used at each stage, becoming more complex as the data rates get higher. Laser transmitters can send up to about 10 gigabits per second just by turning a single laser beam off and on that many times a second. To get to 100 gigabits, the transmitters have to work harder and modulate not just the laser output power but also shift the phase of the coherent light back and forth, and vary the polarization of the laser light. But that is about as much as a single laser can transmit.

Transmission capacity can be increased by using many different lasers emitting at many different wavelengths to send separate signals at many different wavelengths. This is called *wavelength division multiplexing (WDM)* and it was a tremendous advance in fiber-optic communications because it multiplied the transmission capacity of each fiber a hundred-fold. Each laser is designed to

operate in a very limited band of wavelengths—an *optical channel*. Each optical channel is like a broadcast radio or television channel, a band of wavelengths reserved for one laser to transmit through a single fiber.

A single optical fiber can carry light from about 100 laser transmitters on their own separate optical channels in a band of wavelengths from 1530 to 1565 nm where glass fibers are most transparent. Each laser is modulated separately, and their beams are combined together and transmitted through the fiber and then separated at the other end so that each channel goes to its own receiver. That is enough to carry 10 terabits—10 trillion bits—per second through a single hair-thin strand of glass. That is roughly 20 two-layer Blu-Ray disks—carrying 50 gigabytes or 500 gigabits—every second.

With data traffic continuing to grow rapidly, developers are working on ways to expand fiber capacity further. But that is going to take laying new fiber because 10 terabits per second is as much traffic as the standard single-mode fibers that have been installed since the 1990s can handle. So developers are working on new fibers with even more capacity.

12.5.3 Optical Amplifiers in Telecommunications

Optical fibers are extremely transparent, but they are not perfectly clear, so signals fade with distance. Some form of amplification is needed in long-distance fiber systems after signals have traveled through 50 to 100 km of fiber. (Amplifier spacing generally decreases with the length of the system to avoid the growth of noise, which accumulates more in a series of long spans than when signals undergo smaller amplification.)

The erbium-doped fiber amplifiers described in Section 9.5 are the standard type used for long-distance transmission. Standard erbium-doped fibers amplify light in a band from 1530 to 1565 nm, the wavelengths at which optical fibers are the most transparent and can carry signals the furthest. (That match was sheer good luck because erbium works exceptionally well as an optical amplifier.) Each erbium-fiber amplifier can boost signal strength by a factor of 100 or more so that it can travel through the next length of fiber, as shown in Figure 12-8.

Better yet, an erbium-fiber amplifier amplifies all the signals in its operating range, so if the input signal carries 100 wavelengths in the fiber amplifier range, all of them will be amplified. It takes some




Figure 12-8. Fiber amplifier boosts strength of an optical signal.

optical tricks to make the amplification uniform across that band, but it is possible. That means fiber amplifiers can stretch transmission ranges for WDM as well as for single-channel systems.

12.5.4 Types of Fiber Networks and Lasers

As you learned in Section 12.5.1, fiber-optic links carry data at many different rates. They also operate over widely different distances, even at very high data rates. Fibers need huge capacity to transmit signals across a busy cloud computing data center as well as across the Atlantic. Different types of lasers are used for different types of fiber links:

- *Long-distance telecommunications networks* use InGaAsP lasers emitting precise narrow wavelength ranges in the 1550-nm range and erbium-fiber amplifiers.
- *Campus data networks* use InGaAsP lasers emitting in the 1300 or 1550 nm bands, but do not require such narrow linewidths or optical amplifiers. Some may require lasers emitting within 20-nm bands for coarse WDM.
- *Fiber-to-home networks* like Verizon's FiOS system use InGaAsP lasers emitting at 1490 and 1550 nm, sometimes with a fiber amplifier at 1550 nm, to transmit signals to homes. InGaAsP lasers emitting at 1310 nm transmit signals upstream from homes to the switching office.

- *High-speed data links* use inexpensive 850-nm GaAs VCSELs to span short distances within buildings and 1310-nm InGaAsP VCSELs or edge emitters to span longer distances.
- *Automotive networks* use LEDs rather than lasers in their transmitters because the signals do not have to go far.

12.5.5 Laser Communications through Air and Space

Lasers can send signals through the air and space as well as through fibers, but such systems are rare and have different requirements.

Short atmospheric data links have been tested since the 1960s, but have suffered badly from the weather. Time and again, fog, snow, rain and haze have come in the way of the beam. What has proved practical are short connections carrying high data rates through the air between nearby buildings. Microwave links could serve the same purpose, but they normally require licenses and have limited data rates. Laser links through the air typically use wavelengths longer than 1400 nm to reduce the risk of eye damage from inadvertent exposure to the beam, which in practice means they usually use either 1550 nm erbium-fiber lasers or 1550-nm InGaAsP diode lasers.

Laser communications through space seemed like a good idea back in the 1960s because the highly directional beam could carry more bandwidth than a rapidly diverging radio signal. However, keeping the tightly focused laser beam on target proved a daunting task. Modern technology has eased the problems, and in October 2013 the Lunar Laser Communications Demonstration successfully transmitted 622 megabits per second from the spacecraft to earth, six times faster than the best radio link. The first deep space demonstration of optical communications is now scheduled for launch in 2023 as part of a NASA asteroid mission.

12.6 LASER MEASUREMENT

Lasers have proved invaluable for many types of measurements. Laser instruments have long been used in surveying and construction alignment. Scanning laser instruments can measure the profiles of three-dimensional objects. Pulsed lasers can measure distances. This section briefly describes important measurement techniques. It does not try to cover the many ways in which lasers can be used to perform measurements inside specialized instruments.

12.6.1 Laser Surveying and Construction Alignment

Using a laser beam to draw a straight line may sound trivial, but it can simplify many tasks in building construction, surveying, and even agriculture. Early laser instruments for these applications used the red beams of helium–neon lasers, but modern instruments have shifted to red diode lasers.

An important construction application is defining a plane surface to line up mounts for suspended ceilings or partitions. A laser is mounted on a tripod, with the beam directed up into a prism that redirects the beam by 90 degrees so it emerges pointed horizontally toward a wall. The prism rotates in a full circle, sweeping the beam around the walls at the same height. Construction workers mount the hangers for suspended ceilings at this level, so they are all even around a room. Turn the laser plane generator 90 degrees, and it can mark places for partitions.

Surveyors use laser beams to define straight lines for land measurements.

In agriculture, laser beams define the gradients of irrigated fields. The slope should be large enough that the water does not form puddles, but slow enough that it does not run off too fast. A tripod-mounted laser can draw a straight line at the desired angle (for example, at 1 degree from the horizontal). Then the farmer can mount a sensor on his grading equipment to automatically keep the blade at the right height as it moves around the field.

12.6.2 Laser Radar, Ranging, and Atmospheric Measurement

Radar (radio detection and ranging) has long measured distance with microwaves. The radar transmitter emits a short pulse, and an antenna watches for reflections of the pulse. The time it takes the reflected pulse to return measures the distance to an object. For example, a return time of 1 μs means that a pulse traveling at the speed of light made a round trip of 300 m, so the object is 150 m away.

Laser pulses can be used in the same way as microwave pulses to measure distance. Some police forces use compact laser radars (also called lidars) to time how fast cars are moving (by measuring the change in distance between successive pulses) to foil speeding motorists with microwave radar detectors.

Pulses should be short, or have very sharp rise times, for accurate distance measurements because you need an accurate

measurement of the round-trip time. If your laser fires a 1- μs pulse, the pulse is spread along a distance of 300 m, leaving the object's position uncertain to 300 m. That means it is impossible to measure the range to an object 150 m away with a 1- μs laser pulse, but you could get a more accurate position using a 1-ns pulse, which is only 0.3 m (30 cm) long. In practice, accuracy also depends on pulse timing and measurement electronics.

Eye-safe lasers emitting at wavelengths beyond 1400 nm generally are preferred for all atmospheric measurements where humans might be exposed to the beam. Lidar experts regard 1550 nm as a “sweet spot” which gives the best combination of safety and performance.

Lidar has become a key technology for the development of autonomous cars, which need a suite of multiple sensors to guide the car safely without human intervention. Developers are working on a suite of sensors including microwave radar, cameras, ultrasonic sensors, and lidar to drive the car safely and avoid hitting other cars, wildlife, pedestrians, and bicyclists. One key capability lidar offers is mapping the local environment by firing a million pulses per second to create a point map of nearby objects and their distances, such as shown in Figure 12-1. Another is to measure distances to vehicles and other objects at distances up to 200 to 300 m, the distance needed to stop a car moving at highway speed.

Laser ranging is the measurement of distance to objects using laser radar. During the Apollo era, laser ranging was used to measure the distance to the moon by firing red pulses from a ruby laser through a telescope and timing how long the pulse took to return from a retroreflector that astronauts had placed on the moon. Laser ranging from retroreflectors on satellites can precisely locate points on the earth for geophysical research. A laser instrument called the Mars Orbital Laser Altimeter in the Mars Global Surveyor mapped elevations on the Martian surface by firing 10 pulses per second from a neodymium laser and measuring return times.

Armies use laser rangefinders to measure the distance to potential targets. In modern systems, the rangefinder may provide data directly to a gunnery computer to pinpoint the target. Simple laser rangefinders are used in proximity sensors installed on certain anti-aircraft missiles to detect when they are close enough for their warheads to destroy the target. When the rangefinder measures reflection from a target within that range, it triggers detonation of the warhead.

Lidars and lidar-like systems also can be used to measure the distribution of gases and pollutants in the atmosphere. Return times can tell the distance of the gas, and the wavelengths scattered back to the receiver can indicate which gases are present. These techniques are common for remote monitoring of air pollution.

12.6.3 Laser Surface Profiling and 3D Mapping

Lidar and laser ranging can be combined with other techniques for three-dimensional mapping and profiling ranging from recording the 3D shape of fossils to helping self-driving cars navigate safely through an ever-changing environment. Generally this requires scanning a lidar system that fires a rapid series of low-power, short-duration pulses to measure distances accurately. For example, a one-picosecond pulse extends only 0.3 millimeter, and so with a fast sensor it could in principle record shapes with millimeter-scale accuracy. Measuring the angles at which beams are reflected can determine the shape of the surface. Scanning once can record a stationary object like a dinosaur fossil. Repeated scanning several times a second is needed to help self-driving cars avoid moving objects.

A big advantage of laser surface profiling is its ability to collect data on delicate objects without touching them. For example, paleontologists can scan a delicate fossil to build up a three-dimensional computer model of the object, and then manipulate that model on a computer rather than handle the delicate object. Once the object has been scanned, the digital file can be copied so that others can study the digital model. If the original object was distorted, for example, by the pressure of sediments on a fossil, the digital model can be manipulated to remove the deformation. Three-dimensional digital models also can be used to produce replicas of the original object using lasers, as described in Section 13.5, or other techniques.

Developers of self-driving cars have turned to lidars as a way to help the cars navigate through an ever-changing environment. The first generation of lidars were spinning units mounted on tops of cars that spun, spinning many times a second, often with dozens of lasers and sensors operating to collect a dense point cloud that maps the dynamic environment. Repeated scanning is essential because in order to avoid collisions the car must track the movement of other vehicles, pedestrians, bicycles, and wildlife, and predict where they will go next. This also requires quickly identifying

the difference between a stationary tree and an unpredictable deer. The lidar range for 905 nm lasers is limited to about 30 m by eye safety concerns. Shifting to eye-safe 1550-nm lasers allows higher power that extends range to 200 to 300 m.

12.6.4 Interferometric Measurements—Counting Waves


The most precise optical techniques for distance measurement rely on the interference of light waves, described in Section 2.1.3. The trick behind *interferometry* is to measure distances in units of the wavelength of coherent light.

To visualize the idea of interferometry, consider the arrangement shown in Figure 12-9A, where a beam splitter divides a laser beam into two equal halves, directing each half to a different mirror. The mirrors reflect the laser light back to the beam splitter, and part of the light from each mirror reaches the detector. In this case, the two waves arrive precisely in phase, their amplitudes line up precisely, and they constructively interfere to produce a bright spot.


Now, suppose we move the mirror on the right just one-quarter wavelength farther from the laser, so the wave reflected from that mirror has a path a half-wavelength longer than it previously was, as shown in Figure 12-9B. That shifts the light from the mirror at right half a wavelength out of phase with the light from the mirror at the top, so their amplitudes are now 180 degrees out of phase and add to zero, producing a dark spot at the detector. Move the mirror at right another quarter-wavelength away from the laser, and it will produce another half-wave shift, so the two beams are in phase at the detector, and the dark spot turns light. Counting the number of times the spot turns from light to dark and then back tells us how much the mirror moves in units of the wavelength. Remember that with this arrangement the distance the mirror moves is half the number of wavelengths the round-trip distance changes.

The measurements can be quite precise because the wavelength is small. The detected spot shifts from light to dark when the mirror moves only one quarter of a wavelength. When we use a helium–neon laser emitting at 632.8 nm that means we can detect movement of only 158.2 nm. If the mirror was moved 1 μm , it would go through just over three full light–dark–light cycles, each corresponding to a half wavelength.

Interferometry also can be used to measure distance changes across a broad area. In that case, the interference pattern becomes



(A) Original setup with interference producing a bright spot.



(B) Moving the mirror one-quarter-wavelength produces a dark spot.

Figure 12-9. An interferometer measures distance in units of wavelength. (A) Original setup with interference producing a bright spot. (B) Moving the mirror one-quarter wavelength produces a dark spot.

visible as a set of light and dark fringes, tracing paths over the area being studied. Each fringe indicates a change in distance between the two surfaces of one-quarter wavelength.

12.7 LASER LIGHT SHOWS, POINTERS, AND PROJECTION DISPLAYS

Lasers can entertain as well as measure. Laser light shows date back to the late 1960s and became a part of psychedelic entertainment in the 1970s, when they accompanied some famous rock bands. The first laser pointers were red helium–neon lasers, but it took diode lasers to make them commonplace. Now laser projection displays are used in theaters and other entertainment venues.

12.7.1 Laser Light Shows

Laser light shows have been around for decades but, like fireworks, they remain entertaining. Millions of people around the world have seen brightly colored laser beams light up the night sky, a theater, or a planetarium.

The concept is simple. Just bounce visible laser beams off moving mirrors, which scan patterns in the air, on the clouds, on nearby buildings, or on the screen. At any instant, the beam illuminates only one spot, but the eye sees a line because of the persistence of vision, an illusion which makes the eye seem to light for a fraction of a second after the bright spot has passed. Sparklers and other fireworks can give the same effect.

The beams themselves are invisible in the air unless there is enough haze or smoke to scatter light back to the viewer. The scanning patterns can be controlled by music, by a computer program, or by someone operating the lasers (who becomes the optical analog of a musician). It can be fascinating to watch the patterns unfold.

Light shows emerged after bright argon and krypton lasers became available. Argon provides the green and blue beams and krypton emits in the red. The two gases sometimes were combined in “mixed gas” lasers emitting both argon and krypton lines, with a prism or grating separating the colors. Laser light shows are less common today and now use solid-state lasers which are much more compact, far less expensive, easier to use, and much more durable.

Bright scanning beams can pose safety hazards if they are directed into crowds or reflected by mirror surfaces toward a crowd. This could be a problem with some psychedelic rock shows which sought maximum visual impact as well as earth-shaking sound, but concert audiences probably suffered more damage to their ears than their eyes.

12.7.2 Laser Pointers

Red helium–neon laser pointers were foot-long cylinders an inch or two in diameter and often needed an external power supply. The red diode lasers developed in the 1980s ran on AA batteries and were far easier to handle. Their costs have dropped over the years to the point where they are common trade-show gifts and are sold as cat toys as well as pointers.

A second generation of laser pointers came on the market in the 2000s that were based on the 532-nm green second harmonic of diode-pumped neodymium lasers. The eye is about 10 times more sensitive to green light than to the red laser line, so green laser spots look much brighter than a red spot of the same power, as shown in Figure 12-10. The greater visibility of green spots makes them usable for other purposes, such as amateur astronomers pointing at objects in the sky.

A third generation of “blue” diode laser pointers has become common in recent years. Most use 405-nm violet diode lasers mass-produced for Blu-Ray disk, but a few emit near 450 nm and look bluer and brighter to the eye. The eye’s sensitivity decreases as wavelength decreases below 500 nm in the blue–green laser, so blue lasers look fainter than a green laser of the same power, and violet lasers look fainter than the blue.

As Figure 12-10 shows, violet laser pointers are close to the short-wavelength limit of eye response, so their 405 nm spots look deceptively dim to the eye. However, the violet wavelength stimulates fluorescence from many materials, such as white paper, which glows bluish white when illuminated by a violet laser pointer. What makes the violet-illuminated spots look so bright is that your eyes are much more sensitive to the longer wavelengths that fluoresce than to the 405-nm laser emission.

Laser pointers can make entertaining and educational toys because they can reveal some interesting facets of light and color.




Figure 12-10. The sensitivity of the human eye to light peaks in the green. The eye is much less sensitive to longer or shorter wavelengths, so the spot projected by a green laser pointer will look much brighter than those from red, blue, or violet pointers with equal power. (Color in electronic edition)

However, devices sold as laser pointers can pose some potential hazards.

One big problem is that some lasers sold as “pointers” exceed the maximum safe level for laser pointers established by the US Food and Drug Administration—five milliwatts. One milliwatt is safer. It is unfortunately easy to find lasers described as “pointers” on the Internet that emit 100 mW or more.

Some lasers sold may be mislabeled with power below 5 mW. A study by the National Institute of Standards and Technology found that many of the green lasers they purchased emitted much more power than listed because they failed to block light from the 808-nm pump diodes and the 1064-nm fundamental wavelength of the neodymium laser. It is hard to be sure without proper measurement equipment, and high-precision light meters are not cheap.

Even lasers that do meet safety standards can be dangerous when misused. The most publicized example is people who shine green lasers at airplanes. A bright beam hitting the eyes of a pilot coming in for a landing at night can cause flash blindness at a crucial stage of landing. Many incidents seem to involve lasers with

even higher powers. See Appendix A for more information on laser safety.

12.7.3 Laser Projection Displays

Red, green, and blue lasers can be combined to project bright, full-color images on screens. They are becoming increasingly popular for use in cinemas, where they promise better color, higher contrast, and lower costs than xenon lamps used in the first generation of digital cinema projectors. An alternative is “RB” or red–blue lasers, with the blue laser illuminating a phosphor that produces a green beam as well as delivering a blue beam to go with the separate red beam to provide full color.

12.8 LOW-POWER DEFENSE APPLICATIONS

Relatively low-power lasers have found a number of important military and security applications in which they do not actively destroy anything, but may aid in guidance, deterring attacks, or training. Laser gun sights and laser target designators help soldiers aim weapons. Infrared laser beams can be used to divert attacks by heat-seeking missiles. Lasers also can be used in battle simulation systems and as a nonlethal deterrent to insurgent attacks.

12.8.1 Laser Gun Sights

Lasers can be aligned along gun barrels to aid in aiming. These systems act like laser pointers, with the bright spot pointing out where the gun is aimed. Police may use them as warnings to show criminals that they are literally in police sights. Hunters also use them. Versions are available based on red, green, and blue lasers, but the red and green are preferred; the eye is less sensitive to blue than to green.

12.8.2 Laser Target Designators

A laser target designator is a different type of aiming system intended to guide smart bombs and missiles to their targets. A soldier on the ground or in a plane fires the designator, marking the target with a series of coded infrared pulses recognized by a smart bomb or missile.

The smart bomb or missile contains a sensor that looks for the characteristic series of pulses emitted by a target designator. (The coding prevents the bomb from being misguided by bright lights, fires, reflections of the sun, or other sources of light.) The simplest such sensors focus light from the target onto detectors divided into four quadrants. As long as equal amounts of light fall onto each quadrant, the bomb is on course. If one quadrant starts getting more light than the others, the bomb corrects its course to balance the light reaching all the quadrants.

First used in the Vietnam War, laser-guided bombs have proved much more accurate than conventional unguided bombs. However, the soldier holding the designator must remain in a line of sight to the target in order to keep it “marked” with the laser signal until the bomb hits it. Unfortunately, if the soldier can see the target, the target can also see the soldier and fire bullets back at the source of the laser beam, if they recognize they are being marked.

Early laser designators operated at the 1064-nm neodymium line. However, most countries have now shifted to eye-safe wavelengths longer than 1400 nm to avoid the risk of eye damage posed by shorter wavelengths that can penetrate to the vulnerable retina. That is important because modern soldiers usually spend more time using weapons in training sessions than during actual battle, and training works best if it uses equipment actually used in combat.

12.8.3 Laser Battle Simulation

Lasers also play an important role in training in the Multiple Integrated Laser Engagement System, called “Miles,” which equips soldiers with diode lasers and sensors. The lasers are attached to various weapons and each fires a characteristic sequence of pulses. One code indicates that the pulses come from a rifle, another indicates a bazooka, and a third denotes heavy artillery. Sensors are strapped on trucks and tanks as well as soldiers.

When the war games start, the soldiers fire laser pulses at each other and the sensors keep score. A laser-simulated rifle shot can “kill” a soldier. Tanks, however, can only be knocked out by certain types of weapons. (To keep things honest, when the sensors on a tank detect a “kill,” they turn off the controls and fire a plume of purple smoke to tell everyone on the battlefield that the tank is dead.)

Similar laser battle simulation systems are used today by armies around the world. Perhaps someone who saw one of them invented “laser tag” games, which do not actually use lasers.

12.8.4 Laser Countermeasures and Aircraft Defense

Infrared lasers can also be used to foil attacking weapons, particularly missiles aimed at aircraft. Inexpensive heat-seeking missiles that home on aircraft engine exhaust are readily available to insurgents around the world. They are guided by infrared sensors, which look for particular infrared wavelengths emitted by hot gases emerging from engines.

The countermeasure is an aircraft-mounted weapon containing an infrared laser that can fire at attacking missiles. The laser beam looks brighter than the plane’s engine to the missile’s infrared sensors and is modulated in a way that makes it seem to be moving in a different direction than the plane. Once the attacking missile shifts its lock from the engine to the laser beam, the modulated light diverts it from its target. The military has deployed such systems. Similar systems have been proposed to protect civilian airliners from terrorists with heat-seeking missiles, but the costs have been considered too high.

Effective countermeasures require lasers in the 3- to 5- μm mid-infrared range, which are few and far between. Early systems used optical parametric oscillators pumped by neodymium lasers, but quantum cascade lasers can now generate watts at 4 μm , making them an attractive alternative.

12.8.5 Laser Dazzlers

Frequency-doubled green laser *dazzlers* typically emitting 100 to 300 milliwatts have been developed as nonlethal weapons to deter insurgent attacks and have been deployed at checkpoints in the Middle East. The dazzlers are pulsed to get people’s attention and are intended to serve two purposes. One is to warn noncombatants away from dangerous military areas or alert them to the presence of a checkpoint ahead, reducing civilian casualties. The other is to deter attacks by focusing a light on insurgents that is so bright that they cannot effectively attack their targets.

The dazzler’s beam spreads at a broader angle than an ordinary laser. At 100 m, it spreads as wide as a person’s head and shoulders.

At 500 m, the beam is as wide as the front of a car or truck. Similar lasers are being developed for marine defense to block attacks by insurgents or pirates in small boats. Military versions often are mounted on firearms.

Dazzler beams are so intense that an inadvertent exposure at close range can cause eye damage. Models designed for police and civilian security include lidars that measure the range to the person or vehicle and reduce output power to a safe level if they detect the target is too close for safe exposure. Military versions lack that feature, but rely on the soldier's skills and judgment to control the weapon's use.

12.9 SENSING AND SPECTROSCOPY

A broad range of low-power laser applications have emerged from research on the interaction between light and matter. The study of those interactions and how they depend on wavelength is called *spectroscopy*, and it has been a particularly fruitful field for lasers. Once we understand how materials react to light at various wavelengths, we can devise ways to sense the presence of various materials and measure the quantities present.

From an applications standpoint, both detection and measurement are important. Specificity is critical. Whether you are looking for a specific biomolecule to diagnose the presence of disease or warn of chemical attack, you want to make sure you are not getting false alarms from misidentifying other molecules in a complex environment.

Lasers are only one of many tools used in spectroscopy and sensing, so this section will concentrate on laser-based techniques.

12.9.1 Fluorescence Spectroscopy

Intense light can excite atoms or molecules to short-lived, high-energy states that quickly release much of the energy they absorbed at a longer wavelength. This process is called fluorescence and occurs in many materials. You can test for yourself what familiar materials fluoresce with a violet laser pointer. Fluorescence spectroscopy can be used for remote sensing of pollution from balloons or aircraft for in-factory inspection of newly manufactured products. Fluorescence measurements can identify materials by the laser

wavelengths that excite fluorescence, as well as by the fluorescence wavelengths that the lasers excite.

For example, fluorescence spectroscopy could be used to inspect clothes that might have been contaminated by oil from factory machinery that would make them unsellable. If you knew 350-nm ultraviolet light would make cloth fluoresce at 450 nm only if it was tainted with the oil, you could build an inspection station on the production line with a sensor that would reject any clothing that fluoresced at 450 nm, and so only clean clothes would be shipped.

Fluorescence is a powerful tool that can measure how much material is present as well as detect the material's presence. However, care must be taken because many materials fluoresce at similar wavelengths. For example, if you were testing clothing for oil contamination, fluorescence from dyes used to color the fabric might overlap the wavelengths of oil fluorescence. So you might need detailed spectra to check a number of wavelengths where only oil fluoresces.

12.9.2 Absorption Spectroscopy

Absorption spectroscopy identifies materials by the wavelengths they absorb rather than those they emit. It essentially checks what wavelengths are subtracted from the illuminating light when it is reflected.

You do not need a laser for absorption spectroscopy. You can start with an ordinary light source that emits the full spectrum of visible light, and then spread out the transmitted spectrum to see what wavelengths have been absorbed. However, the laser gives much finer resolution, so you can identify the precise wavelengths that have been absorbed. With an ordinary light source, you might know only that light was absorbed near 750 nm. A laser could tell you that light was absorbed at two wavelengths in that range: 749.87 and 750.13 nm. Because the laser can concentrate light at a single wavelength, it also can spot absorption lines too weak to see using other techniques.

Tunable lasers are the most common types used in absorption spectroscopy. Often, they are “swept” in wavelength, by tuning throughout their emission range. For example, a titanium–sapphire laser could be swept from 740 to 770 nm once every 30 seconds,

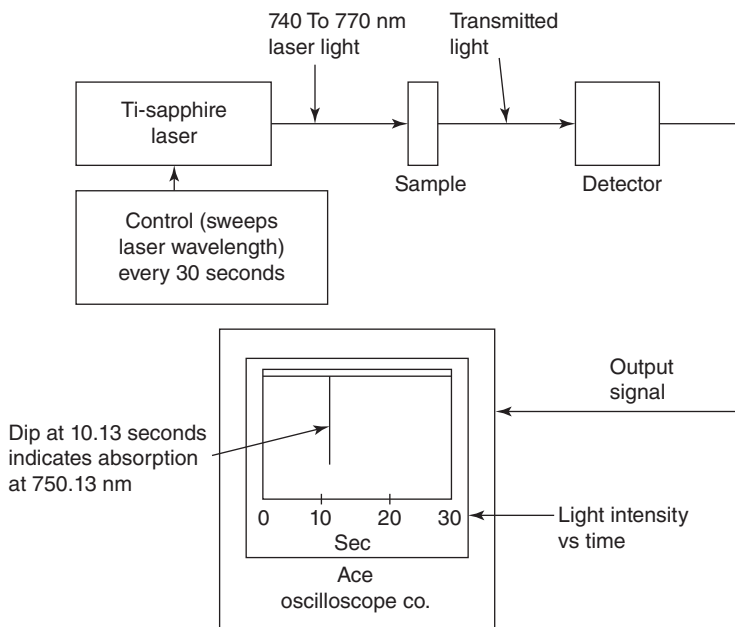


Figure 12-11. Absorption is measured by sweeping Ti-sapphire laser wavelength from 740 to 770 nm every 30 seconds.

as shown in Figure 12-11. If the wavelength varied uniformly with time, you could measure the wavelength at which light was absorbed by timing when the sensor measured absorption. For a 30-second scan at a uniform rate, detecting absorption 10.13 seconds after the start of the sweep would mean the wavelength was 750.13 nm.

Infrared absorption spectroscopy at wavelengths of a few micrometers is particularly useful for organic compounds because absorption bands arising from molecular vibration and rotation—the most useful for identification—lie in the infrared. Laser sources are available in that region, but non-laser sources are easier to use when the extremely high resolution of laser sources is not essential.

12.9.3 Laser Remote Sensing

Many important applications of laser spectroscopy are found in remote sensing, which monitors distant objects or senses conditions remotely. For example, laser remote sensors might probe for

ozone in the upper atmosphere by transmitting a wavelength known to be affected by ozone. The laser would fire a pulse through the atmosphere and then look for returns reflected by ozone molecules. The return time would reveal the altitude of the ozone; the intensity of the return would indicate the amount of ozone. Likewise, a laser could be aimed at the plume of gas emerging from a distant smokestack to look for the emission of particular pollutants, or a laser tuned to particular wavelengths could monitor for chemical weapons or biological agents on a battlefield.

Remote sensing is a broad field that does not require lasers. Spy satellites and earth-resource satellites both perform remote sensing by photographing the Earth's surface looking for particular features. Such sensing is called *passive* because the sensors merely watch. Laser remote sensing is called *active sensing* because it probes objects with the laser beam and monitors their response.

Laser remote sensing can work in many ways. Typically, developers find a wavelength that is either absorbed by the material being sought or causes fluorescence and look for a suitable signal showing the material's presence. To verify their measurements, they may probe with a second wavelength that does not cause absorption or cause fluorescence from the material.

The design of sensing systems depends on both the material being sought and on the application. To measure the emission of a known pollutant from a smokestack, all you need is a laser tuned to a particular absorption line unique to that pollutant and suitable detectors that can measure the absorption. Developing a military or security system to warn of the presence of dangerous chemical agents is more difficult because you cannot be sure what agents might be used, but you need an immediate warning if any might be present. Then the system should identify the agent to indicate what protective action should be taken. Other considerations include such issues as materials that might falsely trigger the sensor, the consequences of false positives, and balancing the problems of false positives against the dangers of missing dangerous materials.

12.9.4 Close-Up Laser Sensing

Lasers also can be used for a variety of close-up sensing, detection, and measurement applications, and often are packaged into measurement instruments. For example, fluorescence induced by blue or ultraviolet light can reveal important details in living

cells viewed under microscopes, so some high-performance microscopes have optional laser accessories to help visualize internal cell structures.

Lasers can be used for some simple spectroscopic measurements of important medical parameters. For example, oxygenated blood absorbs light at different wavelengths than blood lacking oxygen, so it is possible to measure oxygen levels in blood by shining red light through a fingertip.

Some biological tests may use lasers to read out results. For example, the best way to identify specific microbes is by testing with an antibody known to react with that microbe. A reaction indicates the microbe is present, but the reaction may not be easy to detect. Laser illumination might detect fluorescence if the antibody has reacted with the microbe, but not if there is no reaction. Similarly, medical patients could be given fluorescent dyes that are collected by cancer cells and then illuminated with lasers that could excite the fluorescence.

Measuring the scattering of light by particles in a gas or transparent liquid can indicate both particle velocity and size, which are important for applications from medicine to aeronautics. The scattering of femtosecond laser pulses that illuminate the skin can reveal structures in the outer layer of the skin and potentially identify pathologies; the technique is called *optical coherence tomography*, and, like X-ray tomography, the data it generates can be used for three-dimensional imaging of small organs, such as the retina of the eye. Figure 12-12 shows an OCT scan of a healthy retina.

Simple tasks such as counting blood cells can be automated with laser-based sensors. The blood sample can be passed through a transparent tube so thin that only a single blood cell can fit through it at one time. If the tube were illuminated by a laser, the passage of a cell would block some of the light, triggering a sensor to count the cell. Similar laser measurements could be used on assembly lines to verify the presence of parts and to show that they are the correct size and orientated properly.

12.9.5 Spectroscopy in Scientific Research

The cutting edge of spectroscopy is in scientific research, where lasers offer extremely narrow linewidth and the ability to measure small differences in wavelength or frequency, such as the difference




Figure 12-12. OCT scan shows a cross-section through the retina of a human eye, which is only about $500\ \mu\text{m}$ thick. The patient looks into the instrument and the light scans repeatedly across their eye, looking through the layers of the retina and underneath for any signs of retinal disease. Note the scale bar at bottom left, showing how much details is revealed. The patient sees the light scanning the eye, but does not feel a thing. (Courtesy of iStock.)

between the same transitions of two different isotopes of the same element. You will learn more about lasers in research in Chapter 14.

12.10 HOLOGRAPHY

Holography is a method for producing truly three-dimensional images. Today holograms typically are made with lasers. However, Hungarian-born engineer Dennis Gabor invented holography in 1948 as a way to improve the resolution of electron microscopes, a dozen years before the first laser. His idea was to record the phase of waves, information lost in photography because film or digital detectors directly record only the light amplitude. Gabor realized that recording both phase and amplitude would allow him to reconstruct the whole original light wave and thought that would help him better record two-dimensional images.

Gabor originally recorded holograms only of flat two-dimensional transparencies, and by the early 1960s interest in his

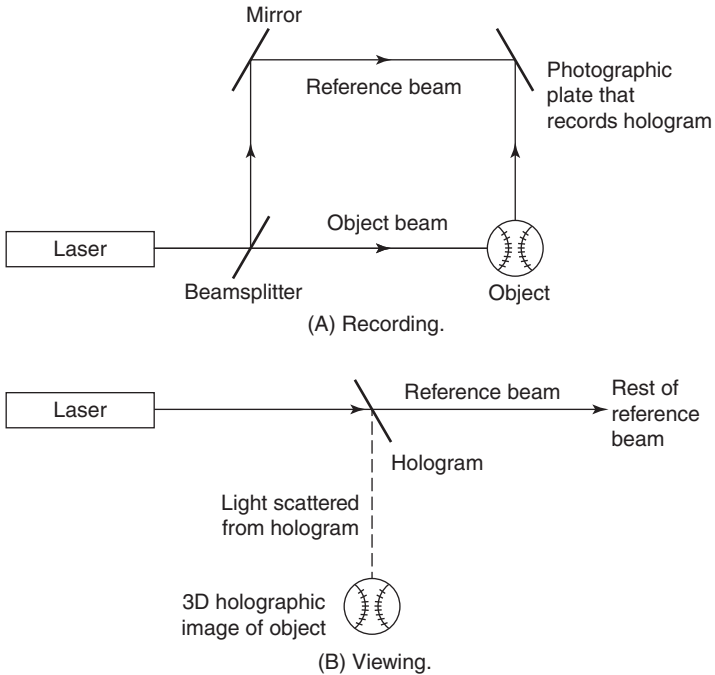


Figure 12-13. Recording and viewing a hologram in coherent laser light.

idea had faded. Then in 1963 Emmett Leith and Juris Upatnieks invented a new three-dimensional form of holography using the coherent light from a red helium–neon laser. The top part of Figure 12-13 shows their process. A beam splitter splits the input laser beam into two beams going in different directions. One beam, called the *object beam*, illuminates an object that reflects the light. The other beam, the *reference beam*, follows a separate path. Separate optics focus the two beams onto a photographic plate or other imager, where the two coherent beams interfere to form a pattern of light and dark areas by constructive and destructive interference.

That interference pattern in the image plane is recorded as an intensity pattern called a *hologram* that bears no obvious relation to the illuminated object. However, when a beam of light following the same path as the reference beam illuminates the hologram, scattering of that beam by the interference pattern reconstructs a wave front that creates a three-dimensional image of the original object, floating in space at the object's position, as shown at the bottom of Figure 12-13.

Holography depends on the coherence length of the illuminating laser beam, which has to be long enough to illuminate the whole object. Most diode lasers are not coherent enough to make holograms because they have small cavities and short coherence lengths, but helium–neon, argon and some other gas lasers have the long coherence required to record large 3D holograms. Some solid-state lasers also have the required long cavities and high coherence.

The first laser holograms had to be viewed in laser light to look three-dimensional. However, later advances made the images viewable in three dimensions with ordinary white light. The most widely used white-light holograms today are called *rainbow* holograms because their color ranges across the spectrum depending on the viewing angle. Rainbow holograms can be mass-produced by pressing them into metallized plastic and have been widely used as eye-catching images on stickers and on the covers of books and magazines.

The unusual nature of holographic images also has led to the use of holograms on credit cards and security labels. In this case the need of sophisticated equipment to make holograms is a good thing, because it makes it hard for crooks to forge holograms that verify the integrity of a credit card or shows that no one has tampered with a parcel. The idea is that because sophisticated equipment is needed to produce holograms, crooks will have a hard time forging them.

These mass-produced holograms can show some three-dimensional effects, but they only have to show brief 3D flashes. To see the full glory of a 3D holographic image floating in space, you need just the right illumination and the right viewing angle, as well as a well-produced hologram. Move your head, and you can see the real three-dimensional depth of the image.

12.10.1 Fake Holograms

The word “hologram” is often misused to describe other types of displays that look lifelike. Futurist Ray Kurzweil has sometimes given talks “via hologram,” and the late rapper Tupac Shakur likewise has performed “live” on stage as a “hologram.” The images are full-color and look lifelike and 3D to the audience, but they are not laser-produced holograms. They are a very effective optical illusion called Pepper’s Ghost that dates back to the 16th century and was demonstrated most famously by John Henry Pepper in a 1862 stage play. However, they are not laser holograms.

12.10.2 Holographic Interferometry

Holography can be combined with interferometry to make sensitive measurements of how objects deform under stress. The trick is to record two holograms with a pulsed laser, one with no stress on the object, the other with a stress applied. This tests how well the object can withstand the applied stress.

The combined hologram shows a single image of the object crossed with bright and dark lines—an interference pattern that measures how much the object's surface moved between exposures with the extra force applied. Figure 12-14 shows vibrations in a guitar. Such vibration patterns can reveal the internal stress within objects, such as aircraft tires, tested at two different pressures. If




Figure 12-14. Holographic interferometry reveals the vibrations of a guitar soundboard; color in electronic edition. (Courtesy of Bernard Richardson, Cardiff University)

the lines change slightly and uniformly, they indicate that the tire expanded slightly when pressure increased, as a sound tire should. If the pattern shows other features, such as a series of concentric rings surrounding one region that indicates a flaw in that area, which might cause the tire to fail. In this way, holographic interferometry can be used to assess the quality of components expected to suffer stress and to measure the effects of changes in gas flow.

12.10.3 Holograms as Lenses

A hologram reconstructs an image of an object by diffracting the light in the reference beam in a way that replicates the wave front of light reflected by the original object. In changing the direction of the input light, the hologram acts as an optical element that diffracts light in a predictable and reproducible way.

This diffraction effect also can be used to make a *holographic optical element* that acts like a lens. In this case, a computer designs a diffraction pattern needed to direct light passing through it in the desired way.

Holographic optical elements do not redirect light as efficiently as ordinary lenses or mirrors, but they can do things that are difficult or impossible with other types of optics. For example, they can be made large and flat to project the image of an aircraft control panel on the windshield in front of the pilot as a “heads up” display, so the pilot does not have to look down at the instruments at a bad time. This application requires large optics with a short focal length that would otherwise be impractical.

12.11 OTHER LOW-POWER APPLICATIONS

This chapter presents only a sampling of low-power laser applications, concentrating on the ones you are most likely to have encountered, the ones that give good examples of how low-power lasers can be used, and some applications that seemed particularly interesting. A few others are listed below, but there are many more.

- DNA and protein sequencing machines
- Positioning patients in medical imaging equipment
- Laser acupuncture

- Precision measurements of the interior of the eye for ophthalmology, including imaging the retina and measuring the interior shape of the eye
- Illuminating a layer of sodium atoms in the upper atmosphere so they glow as laser “guide stars,” bright points in the sky which large telescopes on the ground can use to sense atmospheric turbulence and correct for it to improve how sharply they can see objects in the sky.

12.12 WHAT HAVE WE LEARNED?

- Low-power lasers generally deliver well under a watt and have little effect on materials they illuminate unless the materials are designed to be light-sensitive, such as photographic film.
- Lasers produce well-controlled light that can be focused precisely.
- Laser beams mark a straight line.
- Lasers are monochromatic and generally coherent.
- Diode lasers efficiently turn electrical energy into light.
- Low-power laser beams can read printed symbols, particularly bar codes printed on food packages, magazine covers, and many other products.
- Reading and recording optical disks is a major consumer use of low-power lasers.
- 780 nm lasers are used for CDs, 650 nm for DVDs, and 405 nm for Blu-Ray.
- Laser printers write text by raster scanning a laser beam on a photocopier-like drum and transferring the image to paper.
- The backbone of the global telecommunications network relies on diode lasers transmitting high-speed digital signals through optical fibers.
- Many lasers transmitting at different wavelengths can send signals simultaneously through the same optical fiber.
- Fiber amplifiers boost the strength of optical signals sent through long lengths of fiber across continents or below oceans.
- Laser beams can transmit signals in space.
- Lidar can produce point cloud maps of the local environment for use by self-driving cars.
- Laser ranging measures distances very accurately using radar techniques.

- Lasers can profile the shape and curvature of fragile three-dimensional objects without touching them.
- Lasers can measure distances to within the wavelength of light by counting interference fringes.
- The human eye is more sensitive to green laser pointers than to red laser pointers.
- Violet laser pointers look faint to the eye, but their light can produce bright fluorescence.
- Laser target designators fire a series of coded pulses that smart bombs can decode and use to home in on targets.
- Soldiers use a laser battle simulation system in war games; the laser beams are coded to represent weapons that can kill different targets.
- An infrared laser mounted in an aircraft can blind a heat-seeking missile trying to hit the plane.
- Fluorescence spectroscopy illuminates a material with one wavelength and looks for emission of characteristic longer wavelengths in response to identify materials.
- Absorption spectroscopy monitors the presence of certain materials by observing changes in the intensity of wavelengths of light the materials absorb. An example is monitoring smokestack emissions.
- The coherence of laser light is needed to record three-dimensional holograms, but most holograms now made can be seen in white light.

WHAT'S NEXT?

Chapter 13 covers applications of high-power lasers, which produce major changes in the materials they illuminate.

QUIZ FOR CHAPTER 12

1. You are designing an optical system that uses light to check if a hard-to-reach vent is blocked by shining light through the vent to a sensor on the outside. Your main concern is that birds might build a nest in the vent, but that has only happened once in the past 5 years. Why might you want to use a laser rather than an incandescent light bulb as the light source?

- a. The laser beam could set the nest on fire and unclog the vent.
 - b. Inexpensive diode lasers last much longer than light bulbs.
 - c. You could use laser interferometry to measure the exact size of the bird's nest.
 - d. The laser light would go in a straight line through the vent to the sensor.
 - e. The laser has no advantage.
2. A supermarket laser scanner looks up from underneath a check-out counter. Why is not it blinded by overhead fluorescent lighting?
- a. The laser only scans during the flickers of the fluorescent lights.
 - b. The clerk leans over the scanner to block the light.
 - c. A narrow-line optical filter in the scanner blocks all wavelengths but the laser line.
 - d. The scanner only reads the bar codes when they are flat on the glass surface, so the fluorescent light does not illuminate off the bar code.
 - e. The scanner uses a sensor that only detects red light, which is not present in fluorescent light.
3. Why do Blu-Ray disk players use violet lasers to play video programs?
- a. The short wavelength can be focused to a small spot to help squeeze more data onto a small disk.
 - b. To avoid patents on using other lasers.
 - c. The violet lasers erase information stored on the disks, so they can only be played once.
 - d. Violet lasers are needed only for recording data; Blu-Ray disks can be played in ordinary CD or DVD drives.
 - e. The violet beam is scattered inside the player and makes the front glow.
4. A CD is played with a 780 nm laser and stores 700 megabytes. If the only difference between a DVD player and a CD player was that the DVD was played with a 650 nm laser, how much data could a DVD hold?
- a. 840 megabytes
 - b. 1.008 gigabytes
 - c. 2.1 gigabytes
 - d. 4.7 gigabytes
 - e. 10 gigabytes

5. What other factors allow more data to be stored on DVDs?
 - a. None
 - b. Magnetic recording material
 - c. More sensitive recording material
 - d. Better optics and electronic data compression
 - e. Better optics only
6. If an automotive lidar which fires pulses lasting 1 nanosecond can detect objects up to 300 m away and has to wait long enough for a signal to return before it can fire another pulse, how many points per second can it measure?
 - a. 500,000 points
 - b. 1,000,000 points
 - c. 2,000,000 points
 - d. 5,000,000 points
 - e. 300,000,000 points
7. How accurately can a 1500-nm automotive lidar measure the distance of an object if its output pulses last 1 ns?
 - a. 3 μm
 - b. 3 mm
 - c. 3 cm
 - d. 30 cm
 - e. 3 m
8. How does a color laser printer print full color with only a laser to write on only one electrostatic drum?
 - a. The laser writes on different spots which respond to different colors.
 - b. The laser writes on different spots, which convert the toner to different colors before it is transferred to the paper.
 - c. The printer contains three colors of toner which the laser transfers to the drum as the paper passes through three times.
 - d. An inkjet stage sprays ink onto the paper, and then the laser vaporizes the unneeded color.
 - e. The printer uses a special paper that responds directly to laser light by changing color.
9. A standard digitized version of a telephone voice signal has a data rate of 64 kilobits per second. How many of these voice signals can be combined to make a single data stream at 10 Gbit/s?
 - a. 178
 - b. 1780
 - c. 17,800

- d. 178,000
 - e. 1,780,000
10. The average distance of the moon is 384,000 km from the Earth. If you are building a laser radar to precisely measure the distance from the Earth, how much time do you have to wait for the reflected light to return?
- a. 1 ms
 - b. 1 s
 - c. 1.28 s
 - d. 1.82 s
 - e. 2.56 s

HIGH-POWER LASER APPLICATIONS

ABOUT THIS CHAPTER

Chapter 12 described major uses of low-power lasers, loosely defined as those that do not change objects not intended to be light sensitive. This chapter covers major applications of high-power lasers, which deliver enough energy to significantly alter materials they illuminate. A wide variety of applications are now practical in manufacturing, including cutting, welding, and drilling of metals and other materials, as well as three-dimensional printing. Careful selection of beam parameters allows lasers to machine difficult materials such as glass. High-power lasers and laser-like sources are also used to manufacture electronic components and to perform surgery. Other applications remain in development, notably military weapons, controlling chemical reactions, and thermonuclear fusion.

13.1 HIGH- VERSUS LOW-POWER LASER APPLICATIONS

As you learned in Section 12.1.1, a laser's output power is crucial in determining what it can do. This chapter is about applications that require beams powerful enough to alter materials that are not inherently sensitive to light. A laser that cuts or welds a material clearly is making a major change, whether the material is sheet of metal in

a factory or living tissue during surgery. Writing serial numbers or product codes is a small change, but we consider it materials working if the laser beam burns off paint or surface material to mark the object. Laser scanning of a bar code on a food package is a low-power application because it does not change the package.

The last chapter showed the diversity of low-power laser applications. Less diversity is evident in the use of high-power lasers. Many applications are variations on the basic theme of laser materials working—heating an object to temperatures sufficient to vaporize and remove its exposed surface. You could even say that laser surgery is merely materials working on people, and that laser weapons perform materials working on unfriendly objects.

Recent improvements in solid-state and pulsed lasers have enabled new applications. Ultrashort pulses can cut brittle materials such as glass by ablating very thin layers. High-power solid-state lasers can build up three-dimensional structures from powders, a process called *three-dimensional printing* or *additive manufacturing*. Dramatic improvements in efficiency of fiber and solid-state lasers have made higher powers more practical, opening the way for applications from welding thin steel sheets to mobile weapons able to shoot down rockets and drones.

13.2 ATTRACTIONS OF HIGH-POWER LASERS

In Chapter 12, we listed attractions of low-power lasers. Let us take a similar look at the attractions of lasers for high-power applications:

- Lasers produce well-controlled light that can be focused precisely onto small spots to achieve the high energy density needed to alter materials.
- Lasers can generate light in short pulses, concentrating energy to produce very high peak power for short intervals, which can change materials more than steady beams.
- High-power pulses can ablate very thin layers from a strongly absorbing material.
- Power densities can reach tremendous levels if high-power pulses are focused onto small areas.
- Laser beams do not apply mechanical pressure to objects, allowing noncontact processing that does not distort objects.
- Robotic systems can control lasers and can manipulate optical fibers delivering high-power laser beams.

- Some choice of wavelength is available to match laser output to material absorption.
- Laser light travels a well-defined straight line.
- Glass optical fibers can transmit reasonably high laser powers in the near infrared and visible.

This list is not identical with the one in Chapter 12, but some points are quite similar. On the other hand, the list of disadvantages of high-power lasers is rather different:

- Stringent requirements for safe use of lasers require isolation of the beam in an area separate from workers, which can complicate operations.
- Limited range of lasers that can produce high power.
- High-power lasers are expensive.
- Metals strongly reflect many wavelengths.
- Some materials reflect strongly at the wavelengths of the most powerful lasers.
- Ablated material can get in the way of incoming laser beam and may be deposited on optics, machinery, and other objects.
- Interaction with the air limits transmission of high-power beams over long distances.
- Laser light lacks the momentum of a physical drill or projectile, which may limit its impact.

These advantages and limitations shape the high-power laser applications covered in the rest of this chapter.

13.3 IMPORTANT CONSIDERATIONS AND TRENDS

High-power laser applications involve a number of considerations that may not be immediately obvious but are essential for successful use of lasers. This section briefly describes these factors.

13.3.1 Special Safety Requirements

Laser safety requirements for low-power applications focus primarily on eye safety because the eye is the only part of the body sensitive enough to light to be imperiled by low-power lasers. Laser powers above 500 milliwatts fall under Class 4 of laser safety codes and can burn exposed skin depending on details such as the wavelength, how tightly the beam is focused, and how long the exposure

lasts. They may also ignite other things exposed accidentally to a beam that is misdirected or reflected from a reflective surface.

For workplace safety, Class 4 lasers often are operated in a *laser enclosure*, which completely contains the laser beam. During normal operation, no people may be inside the laser enclosure when the laser is operating. Robotic equipment moves the beam delivery system and/or the workpiece as needed. The details depend on the type of laser and its application. If you work in such an environment, your management should inform you of proper procedures.

13.3.2 Wavelength Effects and Energy Deposition

The success of high-power laser applications depends on how well the laser beam can transfer light energy to the object. That requires selecting a laser wavelength that will be absorbed by the object. The details are complex, but the simple message is that it is important to match laser wavelength to material absorption for high-power applications.

Material absorption can vary widely. Table 13-1 lists surface absorption of some common metals at important laser wavelengths. In general, metals tend to be strongly reflective, but the reflectivity varies across the spectrum. Copper looks darker than aluminum because it absorbs much more green light, but the two metals absorb roughly the same fraction of infrared light at 1 μm . For comparison, the table lists the absorption of dark and light human skin and a typical white paint, but the absorption of those materials and other nonmetals is a more complex matter.

Table 13-1. Surface absorption (percent) at selected laser wavelengths

Material	Laser and wavelength			
	Doubled Nd or Yb (near 500 nm)	Ruby (694 nm)	Nd or Yb (near 1 μm)	CO ₂ (10.6 μm)
Aluminum	9%	11%	8%	1.9%
Copper	56%	17%	10%	1.5%
Human skin (dark)	88%	65%	60%	95%
Human skin (light)	57%	35%	50%	95%
Iron	68%	64%	~35%	3.5%
Nickel	40%	32%	26%	3%
Titanium	48%	45%	42%	8%
White paint	30%	20%	10%	90%

The complexity of nonmetals comes from the fact that they transmit light as well as reflect and absorb it, so the light can go deeper into them than metals. How deep can it go depends on the wavelength. All you need to see an example is a bright lamp like a white LED flashlight and your hand. Close your fingers together, turn on the light, and you will see red light seeping through between your fingers, because tissue absorbs less red light than other colors. You can check that if you have red and green laser pointers. Put the tip of your thumb over the end of the red laser, and your thumb will glow red, but if you use a green laser, the green light will not come through.

What this means for laser applications is that the beam deposits heat below the surface as well as at the top. Whether that is good or bad depends on the application and the material.

How lasers deposit energy also depends on the intensity and length of the laser heating. As you will learn later in this chapter, a very bright and very short laser pulse at a wavelength that is strongly absorbed by a material can blast a thin layer of atoms off the surface without disturbing the underlying material. This can be used to perform delicate surgery or to machine strong but brittle materials such as glass, because the process happens so fast that very little heat gets through to the underlying material.

Other complications arise if the laser emits a steady beam. As metals absorb laser energy, they eventually melt and can even vaporize. The metals also absorb more light as they get hotter, so once the surface starts to melt, the process goes faster. However, vaporization can form a layer of vapor or plasma between the laser and the object, which can block the light from the underlying material, as shown in Figure 13-1.

You will learn more about specific laser processes later in this chapter.

13.3.3 Beam Delivery and Beam Quality

Beam delivery is an important part of any laser process. The beam can be delivered to the target by aiming the laser's output port at it, by moving the target in front of the laser, or by coupling light from the laser into a flexible optical fiber or an articulated mechanical arm and moving the end of the fiber or arm. Robots can be programmed to perform repeated operations on a production line. The




Figure 13-1. Formation of a plasma blocks the laser beam from reaching the object.

optics closest to the target must be cleaned or replaced regularly if dust, vapor, or other contaminants accumulate on them.

The beam delivery system must focus the laser light into a well-controlled pattern on the target material. That may be a small spot for cutting or drilling or a large area of uniform intensity for heat treating a surface. The beam quality, described in Section 5.8.4, is important for laser machining and other applications requiring concentrated deposition of laser power.

13.3.4 Industrial Laser Technology Trends

Some important new trends in industrial lasers have emerged in recent years, notably major increases in laser output power, machining new materials with ultrashort laser pulses, and development of diode lasers for industrial applications. Each needs a brief explanation.

13.3.4.1 Multikilowatt Lasers

Fiber laser powers available commercially have increased markedly over the past decade, helping open new applications and improving laser performance. Multikilowatt lasers used to be rarities limited to the most demanding industrial applications, but now fiber lasers with several kilowatts of continuous output are readily available, and powers in the tens of kilowatts are standard products. Thin-disk

lasers can generate over 10 kW. Prices have also come down, making a broader range of applications practical because high-power lasers can work faster and on thicker materials.

13.3.4.2 *Ultrafast Machining*

The industrialization of ultrafast lasers makes it possible to machine sensitive materials such as glass one layer at a time. A picosecond or femtosecond laser can vaporize the surface layer without damage to underlying material because the pulse is so fast that heat does not have time to penetrate into the bulk. Ultrafast machining has spread into such high profile as the iPhone, where it can cut glass without cracking it.

13.3.4.3 *Direct Diode Lasers*

The high wall-plug efficiency of diode lasers has been alluring for industrial applications, but poor beam quality has been a serious drawback. *Direct diode systems* are being developed to improve diode beam quality to the level needed for industrial applications. Heat treating was a relatively easy target because it involves heating relatively large areas. New improvements in beam combining and beam management have improved power density, so diodes can be used in soldering and brazing. The challenges in improving beam quality and intensity are serious, but the rewards of success would be high and progress is continuing.

13.4 MATERIALS WORKING

Materials working involves cutting, welding, drilling, and otherwise modifying industrial materials, including both metals and nonmetals. The 1960s engineers who played at measuring laser power in Gillettes were doing materials working by drilling holes in razor blades. Materials working has become the single largest market for lasers (in total dollars), nearly \$3.2 billion in 2015 according to the *Industrial Laser Solutions* magazine. This included about \$1.89 billion in machining, \$720 million in “micromaterials” (largely electronics), and \$570 million in marking. To understand how lasers work on materials, we will start with two of the earliest applications, which used lasers to drill very different materials and then move on to more general applications.

13.4.1 Drilling Diamond Dies

One way to make thin metal wires is by pulling the metal through tiny holes in diamond “dies.” Diamond is an excellent die material because it is the hardest known substance. It also conducts heat well, so it can dissipate the frictional heat generated by pulling metal through the hole.

The problem is that diamond is so hard that conventional drills will not penetrate it. Only diamond bits can drill into diamond, and those diamond bits get dull as they drill. Laser light, however, does not get dull. Focusing a laser beam onto a tiny spot in the diamond heats carbon atoms, so they evaporate or combine with oxygen in the air. If you deposit enough laser energy, you can drill through the diamond. Laser diamond drilling was invented in the 1960s, remains in use and is a nice example of how a laser can solve an unusual problem.

13.4.2 Drilling Baby-Bottle Nipples

The rubber used to make baby-bottle nipples is on the other end of the hardness scale from diamond, but that does not make it easy to drill holes through it. The holes in a nipple must be small for milk to flow properly, so tiny wire pins used to be pushed through the molded nipples. However, that is like trying to punch holes in Jell-O with a thin wire; the flexible rubber can catch, bend, and break the fine pins.

A laser beam is an excellent alternative because nothing solid contacts the nipple to push it out of shape. The laser pulse simply burns a tiny hole through the rubber, without touching anything to the nipple. Other techniques can also be used; it is a nice example because it shows that lasers can machine the very soft as well as the very hard.

13.4.3 Noncontact Processing

An important factor in drilling both diamond dies and baby-bottle nipples is that laser drilling is a noncontact process. There is no physical laser “bit” that can get dull while drilling diamond. Nor can the laser beam distort the soft rubber nipple or be caught while drilling the hole.

The noncontact nature of laser processing is important for many types of laser drilling, cutting, and welding. Interestingly,

many materials machined by lasers fall into categories similar to diamond and rubber—either very soft or very hard. Titanium, a light metal used in military aircraft, is too hard to cut easily with conventional metal saws or drills but can be cut readily with a carbon dioxide laser. Other advanced alloys and hard materials are also often cut with lasers. On the other extreme are some rubbers, plastics, and other soft materials that are usually cut with lasers.

13.4.4 How Laser Drilling Works

Drilling makes holes through objects. Laser drilling is done with short laser pulses, each of which removes some material, making the hole deeper. The process requires high peak power and works well for short repetitive pulses. (Long pulses may not reach the high intensity needed to vaporize and remove material from the hole.) The number of pulses needed depends on factors including wavelength, peak power, the nature of the material, and repetition rate.

Penetration depth depends on the optics that focus the laser beam. Figure 13-2 shows that a short-focus lens focuses light onto a tight spot, but the beam spreads out quickly beyond the focal point, reducing the light intensity deeper in the hole. The longer-focus lens does not produce as tight a focal spot, but the beam intensity




Figure 13-2. Effect of focal length on depth of focus; the longer the depth of focus, the deeper the hole that the beam can drill.

remains high deeper into the hole, so it can drill deeper. Drilling deep holes may require optics that change their focus as the hole penetrates deeper.

In the simplest case, a series of laser pulses drills one hole and then the laser is moved to drill another hole with another series of pulses. However, multiple holes may be drilled simultaneously by passing a single high-power pulse through beam splitters to divide it into several lower-power pulses that can be focused onto many spots simultaneously. Usually, this is done with soft, thin materials. One example is perforating ventilation holes in cigarette papers with pulses from a carbon dioxide laser.

13.4.5 Marking and Scribing

Laser marking and scribing are similar to drilling in that they evaporate material from individual spots, but they differ in important ways. In marking, the goal is to write an indelible message, serial number, or trademark on a component by making a series of shallow pits deep enough to be visible as dots that form the desired pattern. The shallower the pits, the faster the job can be done. Alternatively, the laser pulses can be focused through a stencil-like mask to remove material from a larger area. The laser beam can mark a bare surface, but marking efficiency can be enhanced by coating the surface with a layer of light-absorbing paint that evaporates easily. The laser-produced dots are less likely to rub off than ink dots.

Scribing involves drilling a series of deeper holes that form what is essentially a perforated line in a brittle ceramic material. The holes may go through the material but do not have to penetrate completely. Bending the ceramic breaks it along the weakened line scribed by the laser pulses, a technique used to split semiconductor wafers into chips in electronics manufacture.

13.4.6 Laser Cutting

Laser cutting is somewhat like drilling a series of holes that overlap. In cutting, the beam and/or the object move continuously, and the beam itself is normally continuous rather than pulsed. You can think of laser cutting as a process something like slicing with a knife, but in this case, the knife is a beam of light.

Cutting is typically done with the assistance of a jet of air, oxygen, or dry nitrogen. For nonmetals, the role of the jet is to blow

debris away from the cutting zone and improve the quality of the cut. Lasers can cut readily through sheets of wood, paper, and plastic, but thick materials are more difficult, and foods tend to char when cut, leaving an unappetizing black layer. Efforts to slice bread with lasers have produced burnt toast.

Metal cutting works differently and uses a different type of jet. The laser beam heats the metal to a temperature hot enough that it burns as oxygen in the jet passes over it. This process is properly called “laser-assisted cutting,” because the oxygen in the jet actually does the cutting.

13.4.7 Laser Welding

Welding may seem similar to cutting, but the two processes differ in fundamental ways. Cutting separates one object into two (or more) pieces. Welding joins two (or more) pieces into a single unit. Both require heating the entire thickness of the material, but the conditions differ.

Figure 13-3 shows how two pieces of metal are welded together. To form a solid bond, they must be fitted together precisely, leaving




Figure 13-3. Two metal plates being welded with a laser.

very little room between them. The laser beam heats the edges of the two plates to their melting points, causing them to fuse together where they are in contact. If the pieces are not in contact along the entire junction, the weld will be flawed. To make a proper joint, the laser beam must penetrate the entire depth of the weld (the thickness of the material) and melt it fully. The heat of welding should fuse the metal through the entire depth of the weld and in a zone extending on both sides of the junction. The width of the welded zone depends on the nature of the materials being welded, on the power delivered by the laser, and on how fast the laser beam moves along the joint. The further from the mid-point of the weld, the less the heat affects the metal.

The compositions of the pieces being joined together are very important in welding. The materials do not have to be identical, but metallurgists have learned that some compositions do not bond together well, no matter how thoroughly they are heated or what technique is used for welding.

Like cutting and drilling, laser welding is limited in the thickness of material it can handle. However, it has gained wide acceptance in manufacturing high-technology products. The cases of heart pacemakers and the blades of Gillette sensor razors are among products welded by lasers.

13.4.8 Laser Soldering and Brazing

Soldering and brazing are processes in which a filler metal is melted into a joint to seal together other metals that melt at higher temperatures. Brazing is done at higher temperatures, closer to melting point of the metals being sealed, and requires pieces that are fitted more closely together. Soldering is done with metals that melt at lower temperatures and are not fitted together closely, such as wires soldered to circuit boards. Both can be done with lasers, but brazing requires higher and more tightly focused laser power because the materials melt at higher temperatures.

13.4.9 Heat Treating

High-power lasers are also used in heat treating, which raises the temperature of the surface of a metal to convert it to another crystalline state that is harder and more resistant to wear than the rest

of the material. That hard crystalline form of the metal may be too brittle to use in bulk form, but strong enough if it merely forms a layer on another type of metal crystal.

The laser beams used in heat treating—usually high-power continuous-beam carbon dioxide or diode lasers—typically cover a wide area and are scanned across the metal. The surface may be covered with a coating to help it absorb the laser energy more efficiently. Metal surfaces can be heat treated in other ways, but a laser beam can reach into areas that other processes cannot, such as the insides of engine cylinders.

A related application is the annealing of silicon wafers used in electronics or solar cells to change their crystalline state.

13.4.10 Laser Paint Removal and Cleaning

Lasers have found a peculiar niche in removing paint from the skins of aircraft, and cleaning molds used in casting tires and other difficult surfaces. Typically, the materials being removed absorb much more laser energy than the surface being treated, the laser spot burns off the old paint or dirt, exposing a reflective surface that absorbs little laser energy.

Paint removal is important for aircraft because layers of unneeded paint add to the weight of the plane, increasing fuel costs in the highly competitive airline industry and affecting performance of cutting edge military jets. Other processes can be used for much of the aircraft, but lasers are needed for difficult areas.

Laser cleaning is particularly attractive for molds on factory production lines, which normally operate at high temperatures and are often assemblies of parts needed to mold complex structures like tire treads. The laser can be used right on the production line, without need to disassemble the mold or for time-consuming cooling of the hot mold to lower temperatures for cleaning.

13.4.11 Types of Materials Machined

The vast bulk of materials machined with lasers are metals, with cutting and welding of metals accounting for \$1.81 of the \$1.89 billion spent on large-scale machining, according to *Industrial Laser Solutions*. But, plastics can also be welded and cut, generally at much lower powers and different wavelengths than used for metals.

Special processes have been developed to cut glass, sapphire, and other hard, brittle materials, for applications including protective screens of tempered glass for iPhones and other portable electronics.

13.4.12 Types of Lasers Used in Materials Working

Industrial Laser Solutions estimated total sales of lasers for materials working and electronics production in 2015 at \$3.18 billion. They divided the total among four major types of lasers:

- Fiber lasers, virtually all ytterbium doped, used at powers to tens of kilowatts for a wide range of applications working metals and other materials and additive manufacturing (also known as 3-D printing)—\$1.71 billion.
- Carbon dioxide gas lasers, used in a wide range of applications—largely for nonmetals—\$657 million.
- Solid-state lasers based on rods or other bulk materials, largely doped with neodymium, for a range of applications—\$463 million.
- Excimer lasers, largely used in semiconductor photolithography, and diode lasers, used for some low-intensity applications, such as heat treating and soldering—together coming to \$346 million.

Advances in technology have led to rapid growth of fiber laser sales for industrial applications over the past decade. Fiber lasers are now preferred for many applications because of their high power, high efficiency, excellent beam quality, compact size, and compatibility with robotic systems. CO₂ lasers continue to dominate other applications because their long infrared wavelength is better suited for some materials and will be used in new-generation semiconductor photolithography. Thin-disk solid-state lasers have become an important part of the bulk solid-state market because of their high power and efficiency. The high-energy pulses and low initial cost of flashlamp-pumped solid-state remain important in some applications.

Behind the scenes, diode lasers are vital components in pumping all fiber lasers and many bulk solid-state lasers. An emerging trend is the growing use of high-power diode lasers for industrial applications that do not require high beam quality or intensity, notably heat treating and soldering. The high efficiency of diode

lasers is pushing improvements in beam quality and power to make *direct diode* systems more competitive in other applications.

Special processes based on extremely short pulses have led to development of new laser systems to cut glass, sapphire, and other hard, brittle materials for applications including protective screens of tempered glass for iPhones and other portable electronics.

13.5 ADDITIVE MANUFACTURING AND THREE-DIMENSIONAL PRINTING

Laser machining normally produces things either by removing pieces from a block of material or by welding, soldering, or brazing pieces together to make a larger whole. Lasers have also become important with a different approach to production that builds up objects of varied shapes from raw materials. As it has developed, it has been called *desktop fabrication* and *three-dimensional printing* by the public and *rapid prototyping* and *additive manufacturing* by the industry. Now that it comes into wider use, additive manufacturing is becoming the more common term.

The basic idea is to build up a structure from raw material. In industry, it started as *rapid prototyping*, used to produce three-dimensional versions of computer-designed objects for visualization or testing. In the public eye, it originated as *desktop fabricators* for hobbyists who wanted to produce their own three-dimensional objects. The first versions were homemade, but commercial versions followed soon. Futurists soon envisioned a day when home fabricators would be commonplace and used to produce a wide variety of everyday objects. However, early versions were limited to melting a plastic material and using an ink-jet-like printing device to build up structures layer by layer, which was limited to homogeneous objects with limited structural strength.

More advanced techniques have now emerged, some of which use lasers. An approach called *stereolithography* uses ultraviolet light to cause chemical reactions that make molecules of an organic liquid combine to form a solid. In such a system, an ultraviolet laser beam is focused onto a platform sitting just below the surface in a suitable liquid called a *photopolymer*. Controlled by a computer, the beam scans across the surface of the liquid, forming a thin layer of solid on the platform. Then, the platform is lowered a bit deeper into the liquid, and the laser beam scans the surface again, following

Additive manufacturing

How burner tips are made: Using additive manufacturing, spare parts for smaller gas turbines are already being produced today.




Figure 13-4. Complex mechanical parts are built by using a laser to melt patterns in a thin layer of powder, then adding more powder and repeating the process layer by layer. Color in electronic edition. (Courtesy of Siemens)

another pattern to build the next layer of the model. The cycle repeats to deposit many more layers when building three-dimensional objects. The process can be quick and easy with suitable materials, but the special materials can be expensive.

Another approach, called *selective laser sintering*, uses a high-power laser beam to melt a powdered material to form a solid. Figure 13-4 shows how this process can make replacements for a complex burner used in a gas turbine. After a digital blueprint for the part is made, a thin layer of metal powder is spread on a platform, and a laser beam scans across the powder, melting the desired shape. Then, the platform is lowered and a new layer of the powder is applied and scanned, producing another thin solid layer. The

complex part is built up layer by layer and the powder removed to leave the complex three-dimensional object. The aerospace industry is becoming a leader in using this sort of additive manufacturing to make spare parts from a digital library.

A major advantage of this process is that it can fabricate materials from metals, ceramics, and glass as well as plastic. This flexibility has attracted companies/industries, such as aerospace, because it can fabricate custom components in shapes hard to make with other processes and reproduce those components when needed to create a library of spare parts that could be produced quickly when and where needed. Such systems typically use fiber lasers.

13.6 SEMICONDUCTOR ELECTRONICS FABRICATION

The electronics industry uses some types of laser machining described earlier, but its most important laser applications are in the fabrication of semiconductor integrated circuits. Lasers play a role in making the master patterns used in producing integrated circuits, called *photomasks*, but their most important role is in the photolithographic process used to mass-produce chips—a multibillion-dollar business that has become the cornerstone of modern electronics. As of 2017, standard commercial chips were produced using an ultraviolet laser process, but 2018 saw the introduction of a new generation of systems that use shorter-wavelength *extreme ultraviolet* light produced by a different laser process.

13.6.1 Semiconductor Photolithography and Wavelength

The operation of semiconductor integrated circuits depends on a series of extremely thin layers of varied composition deposited on a substrate of silicon. The dimensions of these layers are measured in nanometers. The features on each layer are defined by computer-generated patterns called *photomasks*, which are copied layer by layer by a process called *photolithography*. First, a thin layer of light-sensitive material called *photoresist* is spread across the surface, and then a light is focused through a photomask onto the surface to write the desired pattern. Depending on the type of photoresist, exposed or unexposed areas are etched away chemically, leaving the new layer on the surface. Then, the process is repeated for each layer needed.

The performance of the chip depends on how finely features can be fabricated. The smaller the feature size, the more transistors can be squeezed onto a chip, and the more transistors, the faster the chip can process data and the more data it can store. This is the power of Moore's law, which has driven decades of advances in semiconductor electronics. A key factor in the steady shrinking of features on chips has been the optical resolution of the photolithographic system, which depends on the wavelength of the light used to expose the photoresist.

Optical lithography started with mercury lamps emitting in the blue, then the violet and 365 nm in the near ultraviolet. In the 1980s, developers turned to the 248-nm krypton fluoride excimer laser, which allowed features to be shrunk to somewhat smaller than the wavelength. Argon fluoride excimer lasers emitting at 193 nm took over when the nominal feature size shrunk to 130 nm.

Developers had planned to shift to molecular-fluorine lasers emitting at 157 nm for smaller geometries, but that technology did not live up to expectations. Instead, they turned to optical tricks to enhance the resolution possible at 193 nm. Immersing the process in water effectively reduced the wavelength by increasing the refractive index, allowing 45-nm resolution. Multiple patterning added another couple of steps. But, that was as far as resolution could shrink without a shorter-wavelength light source.

13.6.2 Extreme-Ultraviolet Photolithography

Conventional lasers cannot generate shorter wavelengths, in a region near 10 nm called the *extreme ultraviolet*. Free-electron lasers can produce wavelengths well into the X-ray realm, as you learned in Section 11.6. But, they require a large electron accelerator and huge amounts of electrical power. High harmonic sources are too inefficient to generate the 200 watts needed for an extreme ultraviolet source.

Instead, developers have turned to laser-produced plasma sources described in Section 11.5.4. Powerful pulses from a CO₂ laser are fired at 30- μ m droplets of molten tin, blasting electrons from tin atoms to produce a hot plasma emitting brightly at 13.5 nm. The process is relatively inefficient, so the CO₂ laser must generate over 20 kilowatts in the infrared to produce 200 watts at 13.5 nm, and the pulse must be specially shaped with a "pre-pulse" that heats the tin for optimum excitation. It has taken a lot of efforts. But, at

this writing in late 2018, it looks like this source will be used to make the next generation of chips in the development pipeline, in which feature size will be a mere 7 nm.

13.6.3 Scribing, Mask Repair, and Circuit Changes

Lasers are also used in other electronics fabrication operations. One example is the scribing semiconductor wafers, so they can be split easily into chips, as described in Section 13.4.5.

Resistor trimmers are laser systems designed to modify thin-film resistors deposited on chips or hybrid circuits. Their resistance depends on their surface area, which can be reduced by vaporizing some material with laser pulses, fine tuning the resistor to the desired value.

Other laser systems repair the photomasks used to expose integrated circuits. Masks are costly devices, and a laser pulse can vaporize dirt or other excess material, repairing the mask so it can remain in manufacturing use.

Lasers can also anneal solar cells and other semiconductor devices, changing their crystalline form and/or reducing internal stress to improve their performance. The process is similar to heat treating in metals, described in Section 13.4.9.

13.7 LASER MEDICAL TREATMENT

Physicians started experimenting with lasers soon after the first lasers were demonstrated. The first medical specialists to use lasers were those who were already familiar with light, notably eye specialists (ophthalmologists) and skin specialists (dermatologists). Other specialties followed, and laser medicine is now well developed.

13.7.1 Tissue Interactions and Surgery

Laser medicine is based on understanding how laser light interacts with tissue. Tissue is largely water, so it is a reasonable first approximation to consider water as behaving like tissue. Water absorbs about 80% of incident 10.6- μm light from a CO_2 laser within about 20 μm , roughly the thickness of the skin surface. Water has even higher absorption peaks at 3 and 6 μm , but its absorption varies




Figure 13-5. A CO₂ laser removes the top layer of cells with minimal damage to those underneath.

widely in the infrared. The skin-color variation caused by melanin is evident only at visible wavelengths.

That high absorption and the ready availability of CO₂ lasers long have made the 10.6- μm line a favorite for laser surgery. The absorption is so strong that a tightly focused CO₂ laser beam can vaporize surface cells. In fact, so much of the laser energy is absorbed in the top layer that cells below the surface survive with little damage, as shown in Figure 13-5. The penetration into tissue is just deep enough to seal small blood vessels and stop bleeding, an effect called cauterization that can be helpful in surgery. This is especially valuable for surgery in regions rich in blood vessels, such as the gums and the female reproductive tract, where it allows surgeons to remove thin layers of tissue without causing heavy bleeding.

Carbon dioxide lasers are also used in a type of heart surgery that creates new paths for blood vessels in the heart, called transmyocardial revascularization.

Bone contains less water than other tissue, so CO₂ lasers cannot cut it. However, water absorption peaks so sharply near 3 μm that the 2.94- μm Er-YAG laser can be used for bone surgery. In practice, erbium lasers are used only in certain types of bone, such as oral surgery or implants.

Lasers at other wavelengths in the ultraviolet, visible, and near infrared have other advantages and are valuable for other types of treatment. For example, skin blemishes and tattoos can be bleached by illuminating them with laser beams that match the peak absorption wavelength of the blemish. Lasers that emit ultrashort pulses lasting picoseconds or femtoseconds can perform delicate microsurgery that removes small amounts of tissue without damaging underlying cells. Many recent advances have been made possible by matching the laser wavelength, power level, and pulse duration to specific treatment needs.

In the decades since lasers were first used in medical treatment, interest has shifted away from general surgery to more specialized treatments. Let us look at some important procedures to understand how they work.

13.7.2 LASIK and Other Refractive Surgery

The best known and most intensely promoted laser surgery today uses an argon fluoride excimer laser to reshape the lens of the eye to correct refractive defects that otherwise would force people to wear eyeglasses or contact lenses. The idea of refractive surgery emerged in the early 1980s after the discovery that 193-nm pulses from ArF lasers could efficiently ablate tissue from the cornea, the transparent front layer of the eye, shown in Figure 13-6. That opened the possibility of surgically correcting eye's refractive defects by reshaping the cornea.




Figure 13-6. The anatomy of the eye. Refractive surgery removes tissue to change the shape of the cornea, the outer layer that provides much of the eye's refractive power.

Refractive surgery has improved considerably since that discovery. The most widely used technique today is LASIK (laser-assisted in situ keratomileusis), which peels back the top layer of the cornea, known as the epithelium, before removing tissue from the cornea, then replaces the epithelium, which has been found to reduce complications. Current systems measure the refractive profile of the eye before surgery and compute the precise corrections needed during surgery. However, some complications remain, and accuracy of the correction depends on the healing response of the eye. Some patients require a second surgery to complete correction.

13.7.3 Laser Surgery inside the Eye

The first major successes of medical lasers were in other types of eye surgery. These treatments focused the laser beam through the eye to treat damaged or diseased retinas, instead of performing risky open-eye surgery, and remain important. A newer treatment also focuses light inside the eye, but to help clear up a common complication of cataract surgery.

The simplest treatment to explain is one that can repair a retinal detachment, in which the light-sensitive layer at the back of the eye comes loose from the eyeball. The damage can spread, and without treatment, the entire retina can come loose, causing blindness. The treatment focuses laser pulses to cause small burns that form scar tissue which heals to “weld” the retina down to the back of the eyeball, so it cannot break free. That was a huge advance in treating detached retinas, which otherwise could require dangerous surgery or cause blindness.

A more common laser procedure is treatment of *diabetic retinopathy*, the spread of abnormal blood vessels across the surface of the retina in many people with diabetes. These blood vessels are very fragile, so they can leak blood into the normally clear liquid of the eye, reducing vision and leading to blindness. Ophthalmologists found that shining less than a watt of visible light onto the retina can close the abnormal blood vessels and slow or stop the progress of the disease. It is among the most common types of eye surgery.

Cataracts occur when the natural lens of the eye becomes cloudy and are treated by replacing the natural lens with a plastic one. The surgery leaves behind a membrane that holds the natural

lens in place, and in about one-third of all cases, the natural membrane becomes cloudy after the surgery, again obstructing vision. That can be treated by directing pulses from a neodymium laser through the implanted lens to move the cloudy membrane. A short focal length lens focuses the energy onto the thin membrane, but it spreads across a large area inside the eye, so it does not harm the retina or other parts of the eye. The operation is done in a physician's office. Like other treatments that focus light into the eye, it avoids the time and expense of hospitalization and the risk of surgically opening the eye.

13.7.4 Laser Dermatology

Dermatologists were early laser enthusiasts, but development of standard laser treatments for skin conditions took a while. Today, lasers are used to bleach tattoos and skin blemishes called port wine stains, although both are stubborn targets. Lasers are also used for cosmetic skin treatments to remove wrinkles, improve skin tone, and remove unwanted hair.

The first major success of laser dermatology came in treating dark-red birthmarks called "port wine stains," which often appear on the face or neck. The discoloration comes from abnormal blood vessels just under the surface of the skin. Because port wine stains are near the skin surface, they cannot be treated by conventional surgery, but lasers emitting at wavelengths they absorb strongly can bleach them.

Laser bleaching of tattoos has received considerable attention but remains a mixed success. The idea is to illuminate tattoos with laser light at a wavelength absorbed by the tattoo pigment, breaking down the pigment molecules to bleach the tattoo. How well it works depends on the ink; some almost disappear, but others fade only slightly. If you are lucky, laser treatment will leave a slightly discolored spot where a traditional tattoo used to be; if you are not, you will still wear your pledge of undying love for Annie to your wedding to Zelda.

A new family of encapsulated tattoo inks designed for people who might have second thoughts could make laser tattoo removal much easier. Tiny transparent plastic spheres contain the dyes and protect them from degrading in the body. A single laser treatment splits open the spheres, dumping the dye into cells that break down the molecules to bleach the color.

Other types of laser skin treatment are largely cosmetic. Laser skin resurfacing scans a carbon dioxide laser across aged areas, particularly the face, neck, and hands, to remove the surface layer. This removes surface wrinkles and blemishes, exposing a fresh layer of skin, which requires a few weeks to heal and return to normal color.

Laser hair removal aims a pulsed near-infrared laser emitting in the near infrared at hair follicles. Melanin in the hair and follicle absorbs the light, heating and killing the follicle, and thus stopping hair growth in areas where hair is undesired. The effectiveness of the treatment depends on the contrast between skin and hair color; the darker the hair and lighter the skin, the more effective the treatment. Different lasers may be used for different skin and hair combinations. The 755-nm output of a pulsed alexandrite laser works well, but only on light skin. Arrays of diode lasers emitting pulses at 810 nm are used for light to medium skin, and neodymium lasers emitting at 1064 nm are used for darker skin. Removing light-colored hair is difficult.

13.7.5 Laser Dentistry

Some dentists heavily promote “laser dentistry,” but lasers are not plug-in replacements for mechanical dental drills. However, lasers are valuable for a number of dental procedures and are alternatives to mechanical drills for treating some types of cavities. There are three main applications.

The obvious one is “drilling” teeth, using solid-state erbium lasers emitting at 3 μm . These lasers cannot remove solid enamel well, but they can remove decayed areas and prepare cavities for repair. More important to both dentist and patient, laser pulses are less threatening than the traditional whirring mechanical drill and can be less painful as well. However, lasers cannot drill teeth with old fillings or crowns and cannot be used to drill cavities between teeth, which limits their use. Laser dentistry also tends to be more expensive because dental lasers cost around \$40,000, many times the price of a standard dental drill.

As described in Section 13.3.2, CO_2 lasers can seal blood vessels while removing tissue, so they are used to treat gum disease by removing swollen tissue. Diode and neodymium lasers can also be used to treat soft tissue and can stimulate regrowth and kill bacteria.

Several types of lasers can activate solutions that whiten teeth, a popular cosmetic treatment. The U.S. Food and Drug Administration has approved the use of argon, CO₂, and diode lasers, but blue diode, alexandrite, erbium, and neodymium lasers can also work with some bleaching agents. The choice of laser depends on the costs of the laser and the type of whitening agent used.

13.7.6 Fiber-Optic Laser Surgery

Conventional laser surgery works well for external parts of the body, but delivering the beam inside of the body can be a problem. It can be solved for some types of surgery by delivering the laser energy through an optical fiber that can be threaded into the body without major surgery.

An important example is treatment of kidney stones by threading a fiber-optic instrument through the urethra into the kidney or bladder. The fiber is pointed directly at the stone, and then laser pulses are fired through the fiber to shatter the stone into pieces small enough to pass through the urethra without pain.

A similar treatment can be used to treat benign enlargement of the prostate gland by threading an optical fiber into the urethra. In this case, green pulses from a frequency-doubled neodymium laser are fired through the fiber to remove tissue from the enlarged prostate, allowing urine to flow more freely through the urethra.

Lasers can also be used in minimally invasive or laparoscopic surgery performed through a small incision in the body. A fiber-optic probe carrying the laser beam is threaded through the incision to reach the zone being treated.

13.7.7 Micro- and Nanosurgery with Ultrafast Lasers

Short-pulse lasers can deliver very high peak power to small areas for nanoseconds, picoseconds, or femtoseconds for use in surgery on very small scales. The high peak power vaporizes tissue but does not penetrate deeply. Q-switched nanosecond pulses remove more material at a time and are useful for microsurgery. Picosecond and femtosecond pulses ablate thinner layers for smaller-scale nanosurgery, including selective ablation of parts of cells for biological research. The pulses may be transmitted through optical fibers for some applications.

13.7.8 Laser Cancer Treatment—Photodynamic Therapy

Early attempts to use lasers in cancer surgery failed because laser ablation splattered cancer cells around the surgery site, making the procedure counterproductive. However, lasers now treat some cancers in a different way, called *photodynamic therapy (PDT)*.

The treatment requires two steps. First, the patient is injected with a compound that strongly absorbs red light and is strongly absorbed by cells that are growing rapidly, particularly cancer cells. After a waiting period, which may last a couple of days, a red laser illuminates the affected zone, where the light-absorbing compound absorbs the light and releases reactive compounds that destroy the cancer cells. The laser may directly illuminate the skin to treat skin cancer, but more often, a fiber-optic endoscope delivers the beam to irradiate areas inside the body, such as the lungs and esophagus.

The laser light can penetrate only about a centimeter into tissue, so PDT is limited to treating small tumors close to the surface. The FDA so far has approved the chemicals used for PDT for treatment of esophageal and non-small-cell lung cancer, and some skin cancers, which may also be treated with LEDs. One significant drawback of PDT is that it leaves the patient hypersensitive to light for about four weeks, until the compound has dissipated from the patient's body.

13.7.9 Low-Level Laser Therapy

All the laser treatments described so far produce observable changes in some tissues. However, low-power lasers emitting only a few milliwatts are used to treat pain and assorted other conditions that often do not produce readily observable changes. Most of these treatments are considered alternative medicine, and the orthodox medical establishment is generally skeptical about their effectiveness. Examples include:

- Laser “acupuncture,” in which a low-power laser beam illuminates the acupuncture points originally defined for classical Chinese acupuncture with needles. The effects are said to be similar to needle acupuncture.
- Alleviation of chronic pain by illuminating affected areas with a low-power laser.
- Speeding of wound healing by illumination with a low-power laser beam.

13.8 PHOTOCHEMISTRY AND ISOTOPE SEPARATION

The narrow range of wavelengths in laser light allows it to selectively trigger chemical reactions by illuminating certain materials with specific wavelengths of light. The light may remove an electron from the atom to ionize it, excite an atom or molecule to a highly reactive state, or split apart a molecule. These are called *photochemical* reactions, and their selectivity makes them attractive for separating materials that otherwise behave very similarly, such as different isotopes of the same element, which absorb light at slightly different wavelengths because they have different atomic weights.

Laser photochemistry has long been considered a promising technology, but that promise has been hard to realize. A big problem has been that the equipment required for selective photochemistry is so expensive that it makes sense only if the products are very expensive and are hard to produce in other ways. That led to a focus on isotope separation, particularly in a major U.S. government research effort to separate uranium and plutonium isotopes for use in nuclear reactors and nuclear weapons, respectively. The laser uranium program began in the 1970s was turned over to private industry in the 1990s and was largely shelved in 1999, but efforts have been made to revive a privately developed laser process called Silex. Plutonium enrichment is no longer pursued actively in the United States.

13.8.1 The Isotope Problem

Isotopes are atoms of the same element that contain different numbers of neutrons in their nucleus but the same number of protons. The number of protons in the nucleus defines the atomic number and, hence, the identity of an element. For example, all atoms with 92 protons in their nucleus are uranium atoms, but uranium atoms may contain different numbers of neutrons. The two most common uranium isotopes are uranium-238 (with 146 neutrons) and uranium-235 (with 143 neutrons). Uranium-238 is the most common in nature, but it cannot sustain the fission chain reaction needed to drive a nuclear reactor. To sustain a chain reaction in a conventional reactor, uranium needs to contain at least a few percent of U-235, more than the 0.7% in natural uranium. That requires enriching the U-235 concentration above the natural level.

The chemical properties of the two isotopes differ too little for chemistry to do the job, but two conventional processes, called gaseous diffusion and centrifuges, can. In theory, laser-based processes can separate isotopes more efficiently.

13.8.2 Atomic Vapor Laser Isotope Enrichment

One approach to enrich isotopes is to select individual uranium atoms in the vapor form. The Lawrence Livermore National Laboratory in California developed a process in which copper-vapor lasers excited dye lasers tuned to a narrow range of wavelengths that excited U-235 atoms but not U-238. The process shown in Figure 13-7 was supposed to collect ionized U-235 atoms on an electrode whilst U-238 atoms remained in vapor. However, the process proved to be much more complex, and handling uranium vapor posed additional problems. Those complications and the low demand for enriched uranium brought the program to a halt.

Livermore also developed a similar process for plutonium, designed to reduce levels of the heavier, nonfissionable plutonium-240 and -241 isotopes that could slow down the nuclear chain reaction of plutonium-239 in nuclear bombs. That program was abandoned after the end of the Cold War left the United States with a surplus of weapon-grade plutonium.




Figure 13-7. To enrich concentration of the fissionable U-235 isotope, a dye laser is tuned to a wavelength that ionizes U-235 but not the more abundant U-238.

13.8.3 Molecular Laser Isotope Enrichment Processes

Uranium isotopes can also be enriched in uranium hexafluoride (UF_6), which is a gas at room temperature. The Los Alamos National Laboratory in New Mexico developed a laser process that used light near $16\ \mu\text{m}$ to selectively excite UF_6 containing U-235, but that project was abandoned. An Australian company called Silex Systems Ltd. later developed a similar molecular process, which it licensed for use in producing reactor fuel. That program still exists, but development has slowed, and its long-term fate is unclear.

13.8.4 Laser Separation of Other Isotopes

Laser techniques can separate isotopes of other elements, for applications in medicine, research, or industry. The need for these isotopes is much smaller than the need for uranium or plutonium isotopes, but laser techniques are promising for production of relatively small quantities of isotopes including calcium, gadolinium, and erbium.

13.9 LASER-DRIVEN NUCLEAR FUSION

Government research on laser-driven nuclear fusion began in the 1960s and has continued since then with the dual goals of developing civilian fusion reactors and military simulation of the explosion of hydrogen (thermonuclear) bombs. The civilian side of the project was long stressed in public, but the military program has long paid for most of the research.

The idea behind laser fusion is to implode a mixture of hydrogen's deuterium and tritium isotopes, so they reach temperatures and pressures high enough to fuse their nuclei together in the same way fusion occurs in stars and hydrogen bombs. The laser's role is to fire intense pulses that illuminate the outside of a pellet containing fusion fuel uniformly, so it implodes and produces the required conditions for billionths of a second. Inertial forces cause the implosion, so this approach is called *inertial confinement fusion*, in contrast to the older concept of *magnetic confinement fusion*, in which fusion fuel is heated to high temperatures while confined for longer in a magnetic field. Today, magnetic confinement fusion funded by civilian agencies in the United States and overseas, but the largest

source of inertial fusion funding is the U.S. nuclear weapon program.

One key challenge in inertial fusion is producing a symmetrical explosion that compresses the target uniformly and heats it to the conditions needed for fusion of the hydrogen isotopes. Another is using that energy to generate electric power.

Government scientists have built a series of increasingly larger and more fusion powerful lasers, each of which showed that a bigger and more powerful laser was needed to achieve fusion. Today's biggest fusion laser is the National Ignition Facility (NIF) at the Lawrence Livermore National Laboratory in California. (The name comes from the goal of igniting a self-sustaining fusion reaction.) Construction began in 1997, but experiments did not start until 2009.

NIF is the biggest laser ever built, occupying three connected buildings that together are 704 feet long, 403 feet wide, and 85 feet tall, shown in Figure 13-8. NIF is a massive laser oscillator amplifier. A fiber-laser oscillator emits a one-nanojoule seed pulse, which is amplified and divided among parallel chains of amplifiers. A total of 192 beam lines converge on a single target chamber. At the




Figure 13-8. View of the National Ignition Facility, with the surrounding building cut away to expose the laser equipment. The beam lines feed into the target chamber at lower right. (Lawrence Livermore National Laboratory illustration; government work not subject to copyright.)

fundamental wavelength of 1053 nm, the neodymium-glass amplifier delivers 4.2 megajoules per pulse. The final optics convert the infrared pulse to the 351-nm third harmonic at the cost of about half the energy, leaving 1.8 megajoules to be delivered to the target in a pulse lasting 3 to 20 nanoseconds.

NIF failed to reach ignite a fusion plasma by the end of 2015, and experts now believe that extensive changes or a new laser would be needed to reach the ignition threshold. NIF is now being used to study weapon physics as well as assess prospects for future fusion energy. France built a similar giant fusion laser called Laser Megajoule, which began operation in October 2014 and is still undergoing testing.

13.10 HIGH-ENERGY LASER WEAPONS

The idea of laser weapons goes back to the dawn of the laser age in 1959, when the Pentagon's Advanced Projects Research Agency issued a \$1 million contract to try to build a laser based on ideas proposed by Gordon Gould. Since then, the Pentagon has spent billions of dollars trying to develop high-energy lasers of various types for use on the ground, at sea, in the air, and in space.

That long history reflects tough technical problems and changes in military goals. In the 1960s, the goal was defense against Soviet nuclear missiles, but in the 1970s, that shifted to putting lasers on the battlefield in tanks, planes, and ships. The 1980s brought Ronald Reagan's ambitious quest for an effective way to block Soviet nuclear attacks. That shifted again after the end of the Cold War to defense against rogue states. Major advances in solid-state laser technology since 2000 have stimulated a new wave of development of laser weapons to defend against rockets, artillery, mortars, drones, and small boats, particularly in the hands of insurgents.

To understand why laser weapon development has taken so long, we will look at the missions for laser weapons, the technology required, the problems that have been encountered, and the types of lasers that have been investigated.

13.10.1 Missions for Laser Weapons

The prime missions for laser weapons are defensive. That is, their goal is to destroy or disable offensive enemy weapons. During the

Cold War, the prime concern was nuclear-armed ballistic missiles. Later, concern about nuclear missiles shifted to those that might be launched by a few “rogue states.” A growing concern since the 1990s has been the relatively short-range and low-cost weapons wielded by insurgents: rockets, artillery, mortars, drones, and small boats that may be aimed at civilians as well as military forces.

Disabling a weapon can be as simple as blinding the sensor that guides it to its target. Some projectiles can be destroyed by frying their electronic guidance systems. Rockets and other weapons carrying explosives can be destroyed by heating the deadly payload to the point where it detonates. Lasers can also heat outboard motors on small boats or engines on drones to ignite the fuel or disable the engine. Heating can soften the fuselage of a drone or rocket to the melting point, or ignite flammable components.

Generally, achieving those goals requires a battle management system that can identify a target, track its motion, aim the laser beam at the target, and keep the beam on target long enough for it to deliver a lethal dose of energy. Ideally, the battle management system should be able to verify the kill by seeing the target catch fire, explode, or fall out of control. Judging from videos displaying kills, the laser typically dwells on a target for seconds to disable or destroy it, so the battle management system must keep the beam on target for at least that long. It is also important for the system to verify the kill, so it can move onto other targets without having to worry about the first one.

13.10.2 Requirements for Laser Weapons

A high-power laser alone does not make a laser weapon. Practical laser weapons require four key subsystems, generally with a human operator overseeing them:

- A battle management system, which identifies the target, tracks it, and gathers data on its location and track to aim the laser at the target.
- The laser weapon itself, including the laser, the power supply for the laser, and other components.
- Optics for beam control and pointing to deliver the laser beam to the target.
- The “platform” that houses and carries the laser and associated subsystems, which may be a plane, ship, ground vehicle, satellite, or ground station.

These systems have to work in the battlefield environment. Ship-board lasers have to withstand the moist marine environment. Mobile lasers must withstand vibrations and together with their power source must fit into a vehicle. Optics must be sealed away from dirt or have interfaces that withstand it. Logistics managers want fuels that are readily available in the field, so the armed forces have chosen electrically powered solid-state lasers that can be powered by diesel generators.

Lasers and optical systems must be able to deliver lethal energy through the air to their targets. The air is not as calm and clear as it looks on a clear sunny day, especially when tens or hundreds of kilowatts of laser light are passing through it on their way to a target. The further the beam has to travel, the harder it is to deliver lethal energy. Today's most promising applications for laser weapons are against targets that are within several kilometers. Existing lasers and optical systems cannot reliably destroy targets hundreds or thousands of kilometers away, a capability sought for nuclear missile defense.

Several prototype battlefield laser weapons are now being tested, with power levels from 10 to about 150 kilowatts. Tests of these systems will help identify the power levels needed to kill various types of targets.

13.10.3 Laser Weapon Physics

The physics of laser weapons differs greatly from those of guns and rockets because light is massless. Projectile weapons damage targets because they have a momentum that depends on their mass and velocity, and it is that momentum that damages the target. Light carries energy, and lasers have to deposit that energy on a target to cause damage. A bright laser can damage electronic sensors or the human eye. Lasers damage other targets by transferring enough heat to them to cause damage, such as by detonating an explosive warhead in a rocket, igniting gasoline on an outboard motor, or melting a wing or propeller on a drone.

Target vulnerability to laser weapons varies widely. A rocket covered with black plastic that strongly absorbs laser light would be far more vulnerable than one covered with shiny aluminum. Some parts of targets may be far more vulnerable than others, so lasers could destroy them faster by focusing on their known weak points.

Humans also have uneven vulnerability. Our eyes are extremely sensitive to light, so an antipersonnel laser could blind the eye at power well below what would burn the skin. Intentional blinding is considered inhumane and banned under international rules of war.

13.10.4 Laser Types Tested for Weapon Use

Several types of lasers have been scaled to continuous beams generating tens of kilowatts and up to test their potential for laser weapons. From the 1960s through the 1990s, all known weapons tests used gas lasers, which burned chemical fuels to generate energy and flowed the resulting hot gas through a laser cavity. They worked something like rocket engines, but energy from the hot gas excited laser action instead of propelling flight. They eventually were judged unsuitable for battlefield use because they would have required continuous supplies of exotic and dangerous fuels. The main types were:

- *Hydrogen fluoride chemical lasers* emitting a megawatt at 2.6 to 3.0 μm tested on the ground for potential use in space-based missile-defense lasers because those wavelengths are absorbed strongly in the atmosphere. No programs are active.
- *Deuterium fluoride chemical lasers* emitting a megawatt at 3.6 to 4.0 μm using hydrogen's heavy isotope deuterium to produce a wavelength readily transmitted by the atmosphere. They were tested on the ground and in the air, but no program is active.
- *Chemical oxygen iodine lasers* emitting at 1.3 μm . One was installed in a Boeing 747 for the Air Force's Airborne Laser in an effort to develop a defense against nuclear missiles launched by rogue states. It approached megawatt power but could not attain the range needed for the mission, so it was grounded in 2012 and later scrapped. No further development is under way.
- *Diode-pumped solid-state and fiber lasers* are being used in all current demonstrations of defense against rockets, artillery, drones, small boats, and other relatively short-range targets.

13.10.5 Current Laser Weapon Developments

Current development of laser weapons is focused on defense against rockets, artillery, mortars, drones, and small boats used by

insurgents. Lasers would be used against targets up to several kilometers away, a distance much shorter than the hundreds to thousands of kilometers needed for nuclear missile defense.

Military research agencies considered developing chemical lasers for the task, but logistics experts insisted that any lasers used on the battlefield would have to be run on electricity from diesel generators, which are already in the field. That led to a successful effort to demonstrate a 100-kilowatt diode-pumped solid-state laser.

The Air Force, Army, Navy, and Marines all are working on field demonstrations of diode-pumped solid-state or fiber lasers, which can be powered by diesel generators. So far, tests have ranged from kilowatt-class industrial fiber lasers mounted on HUMVEEs and used to detonate unexploded ordnance to custom-built solid-state lasers generating several tens of kilowatts, with plans for tests above 100 kilowatts. The armed services are developing laser weapons for use on the ground, in the air, or at sea. Key goals include demonstrating reliable operation with wall-plug efficiency of 30% to 40%, and reducing size and mass for field use. All currently being studied emit near 1 μm . Figure 13-9 shows the Navy's



Figure 13-9. The Navy LaWS experimental laser weapon, installed on the USS Ponce in the Persian Gulf in November 2014. This 30-kW system focused light from six industrial fiber laser to shoot down drones and detonate explosives on small boats. Color in electronic edition. (U.S. Navy photo by John F. Williams)

experimental Laser Weapon System (LaWS) on the USS Ponce. The telescope on top collected light from six 5.5-kW industrial fiber lasers and focused it on targets in the air or on the water.

13.11 WHAT HAVE WE LEARNED?

- High-power lasers can significantly change materials not intended to be sensitive to light.
- Short laser pulses concentrate energy to produce very high peak power for short intervals, affecting materials more than exposure to the same power over longer time.
- Lasers do not move or distort objects, allowing noncontact processing.
- Materials working includes cutting, welding, drilling, and heat treating both metals and nonmetals.
- Laser pulses can drill holes in extremely hard materials and punch holes in very soft materials without distorting them.
- Laser pulses do not get dull or bend.
- Lasers emitting more than 500 milliwatts fall under special safety rules intended to reduce human exposure.
- Energy deposition depends on the wavelengths materials absorb.
- Plasma released from an illuminated object's surface can block the laser beam.
- Beam quality is vital in laser machining.
- Drilling requires high peak power from pulsed lasers.
- Lasers can mark or scribe materials without drilling holes all the way through.
- Laser cutting of nonmetals is aided by jets of air or dry nitrogen to blow away debris.
- Laser cutting of metals is aided by oxygen, which reacts with the hot metal.
- Heat treating transforms the surface of a metal to a harder form.
- Lasers can remove paint and surface contaminants from metals.
- Materials-working applications are the largest market for lasers in terms of dollar sales.
- Additive manufacturing builds up three-dimensional components that are hard to make by conventional machining techniques.

- The manufacture of electronic integrated circuits uses ultraviolet light from excimer lasers to form patterns on semiconductor wafers.
- The shorter the wavelength used for photolithography, the finer the details that can be produced, and the more electronic components that can fit on a chip.
- The next step in photolithography is the use of extreme ultraviolet light produced by vaporizing tin droplets with laser pulses.
- Laser medicine is based on understanding the interactions of laser light with tissue.
- Most surgical lasers emit at wavelengths strongly absorbed by water, which is the most important component of tissue.
- The 10.6- μm CO_2 laser can vaporize tissue and prevent bleeding.
- LASIK and other refractive surgeries uses 193-nm ArF lasers to change the refractive power of the cornea.
- Lasers can treat detached retinas and diabetes-related blindness.
- Pulsed dye lasers can bleach dark birthmarks called port wine stains.
- Laser success in erasing tattoos depends on the ink used. Many traditional inks are hard to treat.
- Laser skin resurfacing removes surface wrinkles and blemishes, exposing a fresh layer of skin.
- Lasers can stop growth of unwanted hair by killing hair follicles.
- Three-micrometer erbium lasers can remove decayed regions of teeth and prepare cavities for filling but do not remove solid enamel well.
- CO_2 lasers can treat gum disease.
- Laser energy delivered through optical fibers can shatter kidney stones and treat benign enlargement of the prostate.
- Cold laser treatment includes acupuncture, pain treatment, and wound healing.
- Lasers can selectively excite certain isotopes to purify them.
- Inertial confinement fusion uses high-energy laser pulses to heat and compress fusion fuel to the high densities and temperatures needed for nuclear fusion.
- Laser light carries very little momentum, so laser weapons must damage targets by transferring energy to them.
- Solid-state electrically powered laser weapons can destroy rockets, artillery, mortars, drones, and small boats at distances to several kilometers. They are being developed for use on ground vehicles, ships, drones, and fighter jets.

QUIZ FOR CHAPTER 13

1. Which of the following is not an attraction of high-power lasers?
 - a. Ability to generate extremely high powers in very short pulses
 - b. Amenable to robotic control
 - c. Do not apply physical force to objects
 - d. Can burn chemical fuels
 - e. Can focus light energy tightly onto a small spot to generate very high powers
2. Which of the following is most important for laser drilling of baby-bottle nipples?
 - a. Ability to generate extremely high powers in very short pulses
 - b. Amenable to robotic control
 - c. Do not apply physical force to objects
 - d. High efficiency
 - e. Can focus light energy tightly onto a small spot to generate very high powers
3. A 1000-watt continuous-wave carbon dioxide laser illuminates a titanium sheet for one second. How much energy does the titanium absorb (assuming that absorption does not change with heating)?
 - a. 80 watts
 - b. 80 joules
 - c. 120 joules
 - d. 480 joules
 - e. 1000 joules
4. What requires the most concentrated peak power?
 - a. Drilling diamond
 - b. Heat treating steel
 - c. Soldering electronic components
 - d. Tattoo removal
 - e. Paint removal
5. What type of laser machining is assisted by a jet of air or oxygen?
 - a. Cutting
 - b. Drilling
 - c. Heat treatment
 - d. Scribing
 - e. Resistor trimming

6. How could additive manufacturing be used?
 - a. By hobbyists making plastic toys.
 - b. To make 3-D prototypes to evaluate prospects for industrial products.
 - c. To make complex 3-D components impossible to make by conventional molding.
 - d. To make small runs of spare parts for aircraft.
 - e. All of the above.
7. What types of lasers are used in semiconductor photolithography?
 - a. Fiber lasers
 - b. Excimer lasers
 - c. Diode lasers
 - d. Argon-ion lasers
 - e. Bulk solid-state lasers
8. Laser hair removal works best for people with
 - a. Light hair and dark skin
 - b. Light hair and light skin
 - c. Dark hair and dark skin
 - d. Dark hair and light skin
 - e. Patients whose skin absorbs a dye before treatment
9. What types of lasers are used in LASIK refractive surgery?
 - a. Fiber lasers
 - b. Excimer lasers
 - c. CO₂ lasers
 - d. Argon-ion lasers
 - e. Bulk solid-state lasers
10. What types of lasers are now being developed for use as weapons?
 - a. Deuterium fluoride chemical lasers
 - b. Hydrogen fluoride chemical lasers
 - c. Chemical oxygen iodine lasers
 - d. Diode-pumped solid-state and fiber lasers
 - e. All of the above
11. What types of targets are envisioned for laser weapons now in development?
 - a. Rockets
 - b. Artillery shells
 - c. Drones
 - d. Small boats
 - e. All of the above

12. What target is most vulnerable to damage from a high-energy laser?
 - a. Ballistic missiles
 - b. Mortar rounds
 - c. Helicopters
 - d. Tanks
 - e. The retina of the human eye

LASERS IN RESEARCH

ABOUT THIS CHAPTER

When Theodore Maiman introduced the laser to the world at a 1960 press conference, he predicted it would become a powerful tool for probing the nature of matter. Basic research was one of the laser's first applications and remains important. One recent headline-grabbing triumph was the detection of gravitational waves still shaking the universe after the two black holes merged more than a billion light years away. A laser-based instrument detected space–time vibrating by only one thousandth of the diameter of a proton.

What makes lasers such important research tools is their extraordinary ability to manipulate and measure light, time, space, atoms, and energy. That pays big dividends many areas of science and technology. We cannot cover all the ways lasers are used in research, so we will just sample the high points and fun stuff.

14.1 LASERS OPEN NEW OPPORTUNITIES

The coherent, highly directional, monochromatic, and intense beam from the laser excited some scientists, engineers, and physicians so much that they lined up to borrow, buy, or build their own lasers soon after Maiman's announcement. They soon began getting interesting results.

One important example was frequency doubling. Physicists already knew that nonlinear interactions could generate harmonics of electromagnetic waves. However, the effect was very small at low intensities, so it could not be observed at the high frequencies of light waves, until the laser came along. Months after Maiman's announcement, Peter Franken focused a ruby laser into quartz and photographed weak emission at 347 nm in the ultraviolet, the second harmonic of the 694-nm ruby wavelength. The second-harmonic spot was so faint that an editor mistook it for a flaw in the photo and removed it in the published paper. But, Franken's breakthrough opened the way to the whole new field of nonlinear optics. You can see it at work today in green laser pointers, which emit the 532-nm second harmonic of neodymium.

Many other research advances followed, some important enough to be rewarded with Nobel Prizes. Charles Townes, Nikolai Basov, and Alexander Prokhorov were the first to share one, the 1964 physics prize for developing the principles behind the maser and the laser. (Basov and Prokhorov were recognized for pioneering theoretical work on microwave masers.) Table 14-1 lists all 33 laser-related Nobel Prizes through 2018, all in physics or chemistry.

Looking at the list of Nobels shows a clear pattern—elegant measurements and experiments using lasers have been richly rewarded, showing their importance to both physics and chemistry. The exceptions were for groundbreaking inventions which have changed our society and our lives—the laser itself, semiconductor lasers, fiber-optic communications, and the blue light-emitting diode.

14.2 LASER SPECTROSCOPY

Lasers have brought a revolution in spectroscopy, the study of the interaction between light and matter. The crucial breakthrough was the development of tunable laser sources. The first was the tunable dye laser in the late 1960s, the first laser that could be tuned across a broad range of wavelengths. White-light sources could produce a few photons within a narrow range of wavelengths if their light was spread into a spectrum, but tunable dye lasers could deliver something available from no other light source—a large number of photons concentrated in a very narrow range of wavelengths. That opened the doors to more spectroscopic experiments. Later,

Table 14-1. Laser-related Nobel Prizes through 2018

Year and prize	Recipients	Research
1964, Physics	Charles Townes, Nikolai Basov, Alexander Prokhorov	Fundamental research leading to the maser and laser
1971, Physics	Dennis Gabor	Holography (made practical by the laser)
1981, Physics	Nicolaas Bloembergen, Arthur Schawlow	Development of laser spectroscopy
1997, Physics	Steven Chu, Claude Cohen-Tannoudji, William Phillips	Laser trapping and cooling of atoms
1999, Chemistry	Ahmed Zewail	Studies of chemical reaction dynamics on femtosecond time scales
2000, Physics	Zhores Alferov, Herbert Kroemer	Invention of heterostructures, essential for high-speed optoelectronics and fiber-optic communications
2001, Physics	Eric Cornell, Carl Wieman, Wolfgang Ketterle	Producing Bose–Einstein condensates, sometimes called “atom lasers”
2005, Physics (separate citations)	Roy Glauber	Quantum theory of optical coherence
2005, Physics (separate citations)	John Hall, Theodor Hänsch	Ultraprecise laser spectroscopy and frequency-comb generation
2009, Physics	Charles Kuen Kao	Invention of fiber-optic communications
2012, Physics	Serge Haroche, David J. Wineland	Methods to measure and manipulate quantum systems
2014, Chemistry	Eric Betzig, Stefan W. Hell, William E. Moerner	Development of super-resolved fluorescence microscopy
2014, Physics	Isamu Akasaki, Hiroshi Amano, Shuji Nakamura	Invention of efficient blue light-emitting diodes
2017, Physics	Rainer Weiss, Barry Barish, and Kip Thorne	Observation of gravitational waves
2018, Physics	Arthur Ashkin	Optical tweezers
2018, Physics	Donna Strickland, Gerard Mourou	Chirped Pulse Amplification

the titanium–sapphire laser offered even broader tunability from a solid-state source, and now quantum cascade lasers, optical parametric oscillators, and optical parametric amplifiers also offer broadly tunable output.

Spectroscopy has achieved fascinating things, but there is only room here to describe a few notable achievements.

14.2.1 Doppler-Free Spectroscopy

Atoms and molecules in gases normally move randomly at velocities that depend on the temperature. Those random motions change the wavelengths at which the atoms and molecules absorb and emit, an effect called *Doppler spreading* because the wavelength shifts caused by motion are known as Doppler shifts. This effect had limited how precisely atomic and molecular spectra could be measured.

The high intensity of pulsed laser beams led to an ingenious scheme called Doppler-free saturation spectroscopy, devised by Theodor Hänsch and Arthur Schawlow at Stanford University, both of whom later received Nobel Prizes. They recognized that a dye laser had a line width much narrower than the Doppler width of a gas transition, so they used the laser pulse to saturate absorption in a narrow slice of the spectrum. The ingenious trick was to split the laser beam into two unequal parts, passing in opposite directions through a tube containing the gas, as shown in Figure 14-1. The brighter beam excited all the atoms or molecules to a higher state, saturating absorption so they could not absorb more light at that




Figure 14-1. Doppler-free saturation spectroscopy.

wavelength. A chopper was placed in the path of the high-power beam, and when it passed the high-power beam, the weaker probe beam could pass through the saturated gas in the other direction and be detected. That could effectively cancel the natural Doppler spreading to get Doppler-free gas spectra, a key achievement in spectroscopic resolution.

14.2.2 Femtosecond Spectroscopy and Chemical Reactions

Instead of studying very narrow slices of the spectrum of light, Ahmed Zewail studied narrow slices of time in his Caltech chemistry lab. His goal was to probe how chemical reactions occur.

Chemical reactions are fast; chemical bonds are made and broken in picoseconds to femtoseconds (10^{-12} to 10^{-15} second). To freeze those changing bonds in time, Zewail flashed femtosecond pulses at the reacting atoms and molecules. Making precise measurements required carefully timing the pulses and the course of the chemical reactions in order to determine the exact sequence of events and gave tremendous new insight into chemical reactions.

14.2.3 Frequency Combs and Atomic Clocks

We tend to think of femtosecond pulses as isolated extremely short events, and that view works well when studying the dynamics of chemical reactions. However, mode-locked lasers naturally generate a series of picosecond or femtosecond pulses separated by regular intervals, which offers additional possibilities if you look at the pulses on a frequency scale rather than on a time scale.

Understanding what that means requires a brief look at the ways to describe the variation of laser power, or any signal, over time. Normally, we think of the variation as a function of time, such as a sine wave that varies regularly, or a signal turning off and on regularly to form a series of square waves. However, that variation can also be described as the sum of a number of sine waves having different frequencies and amplitudes, using a technique called *Fourier analysis*. The mathematical details are complex, but the point is that any repetitive wave can be described as the sum of sine and/or cosine waves at a number of different frequencies. For a sine wave, that description is just a single wave at the wave's own frequency. For a square wave, it is the sum of sine waves at odd harmonics of

the square wave frequency, with the intensity decreasing at higher harmonics.

Performing a Fourier transform on a series of short repetitive pulses in time converts the signal into a series of continuous signals at a series of frequencies, as shown earlier in Figure 11-6. The separation between the frequencies equals the repetition rate of the laser pulses. The result is called an optical frequency comb, and the shorter the pulses, the wider the range of frequencies they span. Each frequency component is a continuous sine wave at precisely that frequency, with an amplitude that depends on the original pulse pattern.

Hänsch, then working at the Max Planck Institute for Quantum Physics in Germany, and John Hall at the National Institute of Standards and Technology transformed this interesting bit of esoteric physics into a powerful measurement technique by using another elegant optical trick. When intense light passes through a highly nonlinear medium, the wavelengths are shifted, spreading the light across a wider range of wavelengths. The result is called a *white-light continuum* or *supercontinuum* (described in Section 11.3) to stress how much the wavelengths are being spread. Stretching the range of wavelengths makes it possible to squeeze the pulse duration to shorter intervals. This pulse squeezing is possible because of a variation on the uncertainty principle—the wider the range of frequencies in a pulse, the better it can be defined in time and, correspondingly, the shorter the pulses, the wider the frequency range. With short enough pulses, the frequency comb can span a whole octave—a factor of two—in frequency.

Once a frequency comb spreads across an octave of frequencies, it is possible to double one of the lowest frequencies in the comb and find the doubled frequency on the high-frequency end. Knowing the pulse repetition rate makes it possible to calibrate the whole spectrum of millions of precisely spaced optical frequencies and use them to make extremely precise measurements. This marked a major advance, because precise calibration of optical frequencies had been extremely difficult, and Hänsch and Hall shared a well-deserved Nobel Prize for developing the elegant technique. The technique is now being used in extremely accurate tests of fundamental theories of physics.

It promises to open the door to a new generation of atomic clocks based on light rather than microwave oscillation of a cesium atom at 9,192,631,770 hertz that has long been the primary time

standard. Optical oscillations at much higher frequencies promise much better time resolution.

14.3 MANIPULATING TINY OBJECTS

Another fascinating research use of lasers is in manipulating tiny objects with light. Two important variations are optical trapping, which physically manipulates a tiny object in space, and laser cooling, which slows the thermal motion of an atom or molecule so it reaches extremely low temperatures.

14.3.1 Optical Trapping and Tweezers

The radiation pressure of light—the net force that light applies to an object it encounters—can be used to trap and manipulate tiny objects, an effect called *optical tweezers* or *optical trapping*.

First demonstrated in 1970, optical tweezers use a laser beam to manipulate nanometer- or micrometer-size objects made of nonconductive materials, such as tiny glass spheres. A lens or microscope objective focuses the laser beam on a small spot, creating a narrow waist in the beam; beyond that point, the focused beam expands to larger size and is usually focused through another lens. Radiation pressure from the light in the beam pushes an object toward the most intense zone of the beam—the sharply focused waist—and holds it close to that point. The force of photons passing through that zone tends to push the object slightly away from the point of peak laser intensity.

This trapping technique holds the tiny object at a point where its properties can be measured. The technique has also been extended to molecules and single biological cells, allowing them to be studied in great detail.

14.3.2 Laser Cooling

Laser cooling or *laser refrigeration* slows down atoms or molecules in a gas by removing the thermal energy that keeps them moving. That cools the gas to temperatures that can approach absolute zero.

The trick to laser cooling is to get the atoms or molecules to absorb one photon and then emit another photon that includes a bit more energy. To do that, a laser is tuned to emit photons with a little less energy than an electronic transition in the gas. Atoms

or molecules at rest or moving away from the laser will not absorb the photons because their wavelength is too long. But those moving toward the laser will see the wavelength blue-shifted to a shorter wavelength, matching the absorption line so they will be absorbed.

Eventually, the atoms and molecules that absorb these blue-shifted photons will re-emit the energy. But when they do, they re-emit the energy at the shorter wavelength at the center of the transition. Because they did not pick up quite as much energy from the absorbed photon as needed for the transition, they will lose a little bit of its thermal energy and slow down. The more times the atom absorbs and emits a photon in that way, the more thermal energy it loses, and the cooler it gets.

Atomic motion is in three dimensions, so a single laser beam is not enough to cool atoms to seriously low temperatures. That requires multiple beams, which serve to slow down atoms going in any direction. It also requires adjusting the laser wavelength, because as the atoms cool down, they move more slowly, and their motion does not blue-shift the laser light as much as when they were warmer. Thus, the laser wavelength must be tuned closer to the transition wavelength as the atoms are cooled.

Eventually, laser cooling can slow atoms down enough to produce *optical molasses*, gas so cold that a set of intersecting laser beams can hold them in place. At this point, their temperatures are in the microkelvin range (millionths of a degree above absolute zero).

14.4 ATOM LASERS AND BOSE-EINSTEIN CONDENSATES

Laser cooling did more than set records by creeping ever closer to absolute zero. It also pointed the way to a new state of matter, called an *atom laser* or a *Bose-Einstein condensate*.

Albert Einstein and Indian physicist Satyendra Nath Bose predicted in the 1920s that cooling certain particles, now called bosons, below an extremely low temperature should cause them to collectively drop into the lowest possible energy state, a process called condensation. Laser cooling experiments started pushing toward that goal around 1980, but it was not reached until 1995, when Carl Wieman and Eric Cornell of the University of Colorado cooled a few thousand rubidium atoms below about 170 billionths




Figure 14-2. Bose–Einstein condensation shows atoms condensing from thermal state at left to condensate at right. Color in electronic edition. (Courtesy of Immanuel Bloch, Max Planck Institute of Quantum Optics.)

of a degree Kelvin, and the atoms became a virtually motionless cluster for up to 15 seconds at a time. Like photons in a laser beam, all the atoms occupied the same quantum state.

Condensation comes gradually, as shown in Figure 14-2. The image at left shows a cluster of cold atoms above the condensation threshold. The center image shows the atoms at a temperature just above the threshold. The right image shows a nearly fully condensed cloud of atoms, with almost all of them occupying the same state, and thus producing a strong peak at center.

It was an elegant research achievement that took only six years to earn Wieman, Cornell, and Wolfgang Ketterle of the Massachusetts Institute of Technology a Nobel Prize. “This is a completely new state of matter,” Wieman said when he described the experiment. “It never existed before because the universe was way too hot.” Nor could it ever occur naturally as the universe cools in the far distant future. It is a state distinct from a solid, liquid, gas, or plasma. Laser cooling was not enough to produce a Bose–Einstein condensate; a magnetic trap and further cooling were needed. But the laser cooling was a vital step in the process.

Later experiments showed that a Bose–Einstein condensate could be made to release individual atoms that briefly remained in the same state. This effect was called an “atom laser” because the condensate was releasing atoms that were coherent with the condensate. This coherence could also be seen in the condensate,

where the atoms belonged as a collective whole. The coherence is most evident when the atoms are considered as “matter waves,” which are all coherent with each other.

Although the matter waves in condensates are similar to light waves in many ways, differences are inevitable because atoms differ from photons. Condensates are superfluids that move without friction. Atoms interact directly with each other; direct interactions between photons are rare except in nonlinear materials. Further experiments have shown that coherent atoms interact in ways that are similar to familiar optical effects, although there are important differences.

14.5 DETECTION OF GRAVITATIONAL WAVES

After completing a major upgrade after years of preliminary research, the Laser Interferometer Gravitational-Wave Observatory (LIGO) in 2015 detected gravitational waves for the first time. The discovery made headlines when reported in early 2016, and LIGO continues to discover gravitational waves from far-off events rippling through the universe. Figure 14-3 shows the first five discoveries.

A century ago, Albert Einstein’s theory of relativity predicted the existence of gravitational waves radiating energy through the




Figure 14-3. Gravitational waves reveal the physics of colliding black holes and neutron stars. (Courtesy Caltech/MIT/LIGO Laboratory.)

universe. In theory, powerful events such as the merger of two black holes should shake the fabric of space–time, sending ripples through the universe. Efforts to detect gravitational waves began in the 1960s, physicists began trying to detect gravitational waves, but early detectors were not sensitive enough.

Laser interferometers offered a new way to detect those faint gravitational ripples. The idea was to suspend massive mirrors at the ends of the two perpendicular arms of a long interferometer and look for the tiny movements expected from passing gravitational waves. Detection of the minute movements from gravitational waves required a giant interferometer with arms 4 km long, and extraordinarily precise lasers and optics. LIGO includes two interferometers, one in Washington state and the other in Louisiana, to verify that signals are global.

The first run from 2002 to 2010 failed to detect any gravitational waves, so LIGO was upgraded with new lasers and optics to increase its sensitivity. The first detection came soon after the advanced version of LIGO began operation. More detections have followed. Further upgrades are planned, the Virgo interferometer is now operating in Italy, and a fourth is planned in India. These improvements should both enhance sensitivity and help astronomers pinpoint the sources of gravitational waves. It is the start of a new type of astronomy, made possible by lasers.

14.6 LASER GUIDE STARS FOR ASTRONOMY

Another laser technique has helped ground-based optical astronomy approach the resolution of the *Hubble Space Telescope*. It uses a laser beam to illuminate sodium atoms in the upper atmosphere, creating a bright spot called a *guide star* that a ground-based telescope can use to measure and correct for the atmospheric turbulence that limits how clearly ground-based telescopes can see astronomical objects.

The *Hubble* made breakthroughs in astronomy because it was the first astronomical telescope to get above the Earth's atmosphere, which blurs seeing on the ground. Astronomers had long dreamed of ways to clear up the air, and in the 1970s, military researchers began working on *adaptive optics* that could measure atmospheric effects on light and correct for them by adjusting the shape of a deformable mirror. The military goal was to improve their own

surveillance, including their ability to observe other countries' satellites.

Improvements in sensing systems and electronics made it seem feasible for adaptive optics to improve seeing. But, sensing the state of the atmosphere at night required a bright star in the field of view, and bright enough stars were not available in most of the sky. Military researchers realized that they could create an artificial star by shining a laser beam up into the sky. Laser light scattered from the atmosphere comes right back to the telescope, creating a false star bright enough to help the adaptive optics compensate for turbulence. Those laser guide stars are opening more of the sky to adaptive optics for today's giant 10-m class ground telescopes, and for the coming generation of 20-m class telescopes.

14.7 SLOW LIGHT

Another intriguing area of laser research is slowing light down to a veritable snail's pace. As you learned earlier, the speed of light in a material equals the speed of light in vacuum divided by the refractive index, so increasing the refractive index can slow light several times below its normal 300,000 km/s. The refractive index can reach extremely high values, but those peak values are possible only at wavelengths where the materials have resonances that make them absorb light very strongly.

Steve Harris at Stanford University found a way around this problem by illuminating a material with so much laser light that he saturated the transition. In 1995, he used this approach to slow down the velocity of light at a resonant wavelength in lead vapor to 1800 km/s, a factor of 165. The idea did not capture much public attention until 1999, when Lene Hau at Harvard University slowed light to 17 m/s by passing it through sodium atoms cooled below the threshold for forming a Bose–Einstein condensate and illuminated by a low-power laser tuned to the same wavelength. That was an impressive achievement, but it required extremely low temperatures and extremely narrow laser line widths.

Since then, the technique has been extended to other materials and warmer temperatures, using variations on making materials transparent at wavelengths at which their refractive index is extremely high. Slow light has also stimulated practical interest,

because slowing down the speed of light could allow construction of optical buffers able to store optical signals for use in fiber-optic communication systems.

14.8 NANOSCALE LASERS

Nanoscale lasers are hard to build almost by definition because optical photons have wavelengths of several hundred nanometers. Electrons can be confined to a much smaller scale because they are considerably smaller. As described in Section 10.5.2.2, heterostructures can confine electrons in all three dimensions in quantum dots, which range in size from 2 to 50 nanometers. However, quantum-dot lasers are not single quantum dots; they are larger semiconductor lasers with active layers which contain one or more quantum dots in the junction plane. Nanolasers are more complex, and their operation generally involves the nanoscale properties of electrons.

14.8.1 Photonic Crystal Lasers

Laser action has been demonstrated in *photonic crystals*, which are periodic structures consisting of many very small elements that act together to confine light. Typically, the elements are stacks of very thin layers with different refractive indexes that collectively act to block transmission of light at certain wavelengths in ways similar to thin-film optical coatings. In some cases, the layers are two different compounds; in others, one is solid glass and the other layer is a material containing holes smaller than the wavelength of light, which reduce the effective refractive index of that layer.

Photonic crystals can create cavities with modal volumes smaller than the cube of the emitted wavelength (in the solid) and with high cavity confinement. This allows photonic crystal lasers to have extremely low thresholds, although their output power is low, and single-photon emission is difficult. A number of approaches have been demonstrated in various materials.

14.8.2 Single-Atom Lasers

The ultimate in miniaturization is a single-atom laser, which has been demonstrated by placing a single atom inside a resonant

cavity. Light can excite a single atom or ion that is either trapped in a resonant optical cavity or one of a series falling through the cavity. Jeff Kimble's group at Caltech in 2003 trapped a single cesium atom at a time in an optical cavity for about 100 ns and optically excited it to emit a series of photons resonant with the cavity. Other experiments have used a similar scheme but observed light emitted by a series of atoms falling through a resonant cavity, where more than one atom may be in the cavity at a time.

The fact that a resonant cavity contains only a single atom does not mean that the volume of the laser cavity is the size of that one atom. Typically, it is larger, defined by a pair of cavity mirrors on either end of the volume containing the atom. That volume may be considerably larger than a single wavelength if, as in Kimble's experiment, the light-emitting atoms are falling through the space between the two mirrors.

14.8.3 Spasers and Surface Plasmon Lasers

True laser action requires oscillation, which is hard to achieve on a nanoscale with photons because their wavelengths are hundreds of nanometers. However, electrons are much smaller, and researchers have found compact clusters of electrons oscillating at optical frequencies, called *surface plasmons*, can oscillate on a nanoscale, in volumes less than a cubic wavelength.

Surface plasmons form at the interface between a highly conductive metal and a semiconductor or dielectric material. Photons can excite surface plasmons, and the two can couple together to form a *surface plasmon polariton*, which propagates along the interface. Although they oscillate at optical frequencies, that oscillation is over an area smaller than a wavelength.

A number of groups have demonstrated devices that they call either *nanolasers* for their scale or *spasers* for Surface Plasmon Amplification by the Stimulated Emission of Radiation, in analogy with the acronym that gave us the word laser. Typically, these are long, thin rod-like devices, longer than the emitted wavelength, but much thinner, some with volume smaller than a cubic wavelength. This makes it possible for a surface plasmon "laser" to operate in a volume smaller than a cubic wavelength.

So far, nanolasers remain in the research stage, part of a broader field of nanoscale photonics that includes solar cells and sensors as well as light emitters. Individual devices are so small that they may

find more applications as parts of larger integrated photonic devices than as individual components.

14.9 STRANGE LASERS

Researchers have come up with countless variations on the basic principle of the laser. The basic building blocks are simple: a resonant cavity that contains a light-emitting material that can be excited to yield stimulated emission. These ingenious twists on the laser principle range from the amusing to the potentially useful. A few deserve mention to show how versatile the laser principle is.

14.9.1 The Jell-O Laser

Early laser researchers tried making lasers from a wide variety of materials and had so much success that the playful physicist Arthur Schawlow suggested “anything will lase if you hit it hard enough.” To test the theory, he and his Stanford research associate Theodor Hänsch focused pulses from a nitrogen laser into prepared samples of colored Jell-O dessert. Schawlow had bought 12 different flavors, which they tried one by one, but none of them worked.

Not willing to give up, Schawlow decided that the dye sodium fluorescein was “almost” nontoxic, so he mixed some with a clear batch of gelatin. Much to his delight, it lased, showing an edible laser was possible in theory. But Schawlow did not venture to eat the successful sample, as he had the unsuccessful Jell-O flavors. The two published a brief report on their success in 1971, which did not stop the two from separately receiving Nobel Prizes.

14.9.2 Living Lasers

Forty years later, two scientists at the Harvard Medical School went further and demonstrated laser action in a living cell. They genetically engineered cultured human kidney cells to produce the green fluorescent protein, originally found in jellyfish. Then, they placed one cell between a pair of mirrors 20 μm apart that formed a laser resonator. When they illuminated the cell with weak pulses of blue light, the genetically engineered cell emitted green light that produced a visible laser beam. The power levels were so low that the cell survived without problem. Their technique could be used to

study individual cells outside the body or processes within a living animal.

14.9.3 Anti-Lasers

In theory, a laser could be made to run backward in time by absorbing light rather than emitting it. In 2011, Yale University physicists Hui Cao and Douglas Stone demonstrated such an anti-laser. Normally, a laser amplifies light photons that bounce back and forth between a pair of mirrors at a resonant frequency, so an integral number of light waves fit into the cavity. They had predicted that laser light directed into an optical cavity in which it could resonate would be absorbed if the cavity was filled with an unexcited material that would emit at the resonant wavelength. The Yale experiment showed that light at the resonant frequency was absorbed very strongly, but light at other wavelengths were not—so the anti-laser works. It might even be useful for optical sensing light or modulating it.

14.10 EXTREME POWER ULTRASHORT PULSE LASERS

Shrinking the duration of a laser pulse concentrates its energy in time and thus can generate extremely high peak powers. A 10-joule pulse squeezed into 10 femtoseconds has a peak power of one petawatt (10^{15} watts), more than the peak power from the 1.8-megajoule National Ignition Facility (NIF). This has led to development of a new class of lasers that can explore both extremely short intervals of time and extremely high concentrations of laser energy.

Producing these extremes of time and power density requires pushing technological frontiers. As you learned in Section 4.6.2, producing extremely short pulses requires extending the range of frequencies in a laser pulse and poses major challenges in pulse amplification as well as duration. This section will explain that technology by tracing its development over two decades, from the pioneering Petawatt Laser at the Lawrence Livermore National Laboratory to European Union's ambitious multisite Extreme Light Infrastructure project.

14.10.1 Petawatt Lasers

The technology for petawatt lasers emerged from development of high-power fusion lasers at Livermore in the mid-1990s. The first

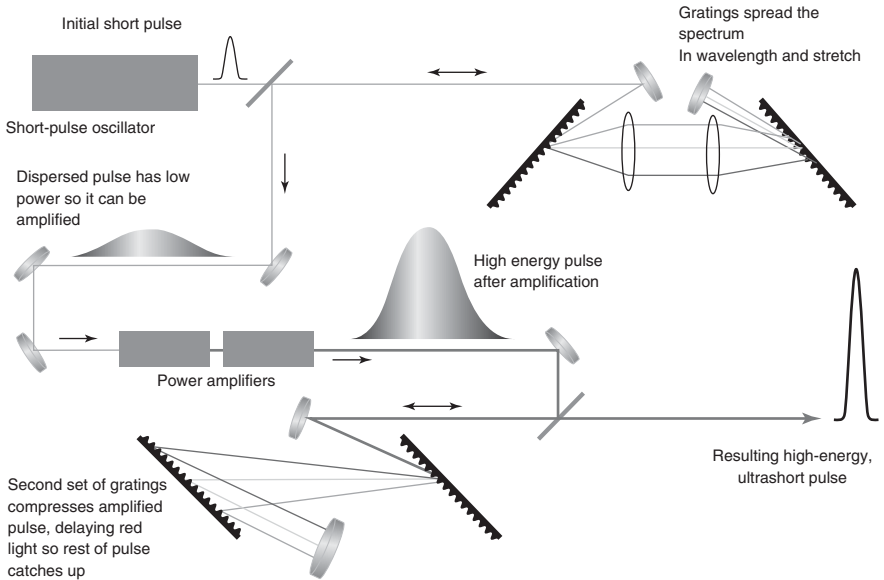


Figure 14-4. Chirped-pulse amplification permits pulses to be amplified to very high peak powers without damaging amplifier optics. (Adapted from Government work, not subject to copyright, Lawrence Livermore National Laboratory.)

petawatt laser incorporated a chain of neodymium-glass amplifiers from the Nova fusion laser, the predecessor of NIF. Up to that time, the amplification of high-power pulses had been limited by nonlinear effects and by optical damage to the final amplifier stages. Livermore overcame that problem by using a technique called *chirped-pulse amplification*, shown in Figure 14-4.

Chirped-pulse amplification was developed to avoid damage by stretching pulses in time before and during amplification in a solid material that could be damaged by extremely high peak power. Once the amplified pulse exits the solid, it can be compressed in time to produce extremely high power. In Figure 4-5, the input pulse first passes through a pair of diffraction gratings arranged so the shorter wavelengths travel a longer distance than the longer wavelengths. (The shorter wavelengths are called “blue” and the long ones called “red” no matter what the actual wavelengths.) This “chirping” process spreads the pulse over a longer time interval, with the “red” wavelengths at the front and the “blue” wavelengths at the back because they traveled a longer distance. This reduces the peak power to a level low enough that when amplified it will not damage the optics. Once the pulse is amplified, it passes into

the air and goes through a second pair of gratings, arranged to delay the red light so the blue light could catch up, recreating the original short pulse in the air, which can tolerate the extremely high peak power that would have damaged the amplifier.

The low-gain bandwidth of the neodymium-glass laser in Livermore's original petawatt laser limited pulse duration to hundreds of femtoseconds. Amplifiers with broader gain bandwidth can generate shorter pulses. Titanium-sapphire and optical parametric amplifiers can generate chirped pulses approaching one femtosecond.

14.10.2 Attosecond Pulses

The oscillation period of light waves limits the minimum pulse duration. A single wave cycle at 800 nm, in the center of the tuning range of the Ti-sapphire tuning range, lasts only about 2.7 femtoseconds, so a Ti-sapphire laser cannot produce a pulse much shorter. That is a fundamental limit, but it is eased at shorter wavelengths. At 400 nm, the short end of the visible range, a single wave cycle lasts only about 1.3 femtoseconds. So, it is not surprising that the shortest pulses on record from Ti-sapphire lasers are in the femtosecond range.

This fundamental limit arises from the uncertainty principle. As described in Section 4.6.2, the uncertainty in the frequency (the bandwidth) of a pulse times the duration of a pulse must be at least 0.441. This means that generating a 5-femtosecond pulse requires a bandwidth of at least 88 terahertz, or roughly 710 to 910 nm from a Ti-sapphire laser. Using the same formula, you need a bandwidth of 441 THz to produce a 1-femtosecond pulse. That is a little bigger than the visible spectrum, which is about 420 THz wide. Producing even shorter pulses, with duration measured in attoseconds (10^{-18} second), requires higher frequencies in the ultraviolet or X-ray bands.

A leading way of generating attosecond pulses is *high harmonic generation*, described in Section 11.5.2, which focuses intense femtosecond pulses of laser light into a jet containing atoms of a rare gas such as argon or xenon. Each femtosecond pulse contains a series of peaks in the electric field marking oscillations at the frequency of the laser light. At these peaks, the electric field is powerful enough to pull valence electrons out of the atoms in the gas, but as the field drops back to a lower level, the atomic nucleus pulls the electron

back into place. The recaptured electron releases the extra energy it collected from the field of the laser pulse in an attosecond burst of extreme ultraviolet waves and X-rays.

Attosecond pulses offer new opportunities for basic research. One use is to probe what happens on the extremely short attosecond time scale. Another is to study how matter interacts with the extremely high intensities that can be produced in attosecond pulses.

14.10.3 Extreme Light Infrastructure

The European Union is nearing completion of the first three parts of a major international facility for ultrashort pulse research called the Extreme Light Infrastructure, which began commissioning in late 2017. These are:

- The ELI-Attosecond Light Pulse Source in Szeged, Hungary, will use broadband attosecond pulses to take snapshots of electron dynamics in atoms, molecules, plasmas, and solids and to study effects of ultrahigh intensity lasers. <http://www.eli-alps.hu>
- The ELI-Beamlines Facility, near Prague, Czech Republic, aims to develop new sources of short pulses of light and particles and to perform experiments with electromagnetic field intensities to about 10^{23} W/cm². <http://www.eli-beams.eu>
- The ELI-Nuclear Physics Facility in Magurele, Romania, will focus on laser-based nuclear physics and intense gamma-ray beams, as well as study effects of fields on the order of 10^{23} – 10^{24} W/cm². <http://www.eli-np.ro>

The center of the ELI-Beamlines Facility is a diode-pumped laser built by Livermore that is able to fire pulses of up to 16 joules lasting 28 femtoseconds. The laser is designed to deliver pulse intensities to 10^{23} watts/cm², probing a new realm of light–matter interactions. Figure 14-5 shows the layout of the laser, which covers an area of 4.6×17 m plus a 4-m extension for beam compression. That is only a small fraction of the size of the NIF, shown in Figure 13-8, despite the fact that the ELI laser will generate higher peak intensities. NIF requires much more space because it can deliver pulses to 1.8 megajoules and uses flashlamp pumping.




Figure 14-5. The High-Repetition-Rate Advanced Petawatt Laser System built for the ELI Beamlines facility in the Czech Republic. Designed to produce extremely high intensities in 28-femtosecond pulses, the laser stretches 17 m long. (Courtesy of Lawrence Livermore National Laboratory.)

14.10.4 Laser Acceleration

The extremely high electric and magnetic fields generated by ultrashort pulses can accelerate charged particles to high energies over very short distances. Theoreticians Toshi Tajima and John M. Dawson proposed the idea in 1979, but lasers then available were not powerful enough for the job. The situation changed with the development of chirped-pulse amplification and petawatt lasers, and recent years have seen a series of experiments with electrons, protons, and heavier ions.

Traditional charged particle accelerators use arrays of massive electromagnets in long tubes, with power through the electromagnets switched to keep the fields pulling the particles along. Small accelerators for laboratory use have meter-long tubes, but the big ones used in major research programs are miles long. At best, they can increase electron energy by tens of millions of volts per meter. In contrast, pulses from ultrashort lasers fired into a plasma can reach electric fields of teravolts (trillions of volts) per meter at their peak and accelerate electrons to 100 million electron volts in only 1 mm.

Laser accelerators will not replace the giant accelerators used in particle-physics research. However, they can fit in ordinary laboratories to allow smaller-scale research. And laser accelerators can

also produce more powerful beams, making them better-suited for research in fields such as medical treatment with particle beams.

14.11 X-RAY FREE-ELECTRON LASERS

The free-electron lasers described in Section 11.6 have emerged as important X-ray “light sources” for research. The big advantage of X-ray lasers is that their peak power per unit wavelength is far brighter than other X-ray sources. The first of the big hard X-ray free-electron lasers, the Linac Coherent Light Source (LCLS) at the SLAC National Accelerator Center, was a factor of 10^8 brighter than conventional sources.

Research X-ray sources are large and very expensive instruments and are generally operated as large “user facilities” open to a broad range of researchers. X-ray free-electron lasers are similarly large and expensive because the electron beams that provide their energy come from large accelerators. SLAC renovated the last kilometer of the two-mile Stanford Linear Accelerator, built in the 1960s for particle physics, to power LCLS. Its output can be tuned between 0.15 and 1.5 nm, and a second-harmonic accessory can produce wavelengths down to 0.7 nm, although at greatly reduced power.

Highly reflective resonator mirrors are not available in the X-ray band, so LCLS and other X-ray free-electron lasers are amplifiers rather than oscillators. As in conventional lasers, the initial spontaneous emission is amplified by stimulated emission, but X-rays make only a single pass through the amplifier, a process called *self-amplified spontaneous emission* (SASE). The process is chaotic so the output has sharp peaks lasting 0.3 to 2 femtoseconds, depending on wavelength. Grazing-incidence mirrors deflect LCLS output to one of six stations with measurement instruments.

X-ray lasers have proved invaluable in studying materials and crystalline structure. Much of the work has been on proteins, but they also have been used to simulate a rain of tiny diamonds thought to form in the interior of the ice giants Uranus and Neptune.

Two other large X-ray free-electron lasers are now operating. SACLA at RIKEN in Japan began operation in 2012. The European X-Ray Free-Electron Laser began operation in Hamburg, Germany, in 2017 and is now the world’s most powerful laser. SLAC is planning to add a second section of accelerator to build LCLS II, which

will operate in parallel with the original, but fire a million pulses a second instead of the present 120. The \$1 billion new machine will begin operation in the early 2020s.

14.12 OTHER EMERGING RESEARCH

There is no shortage of ideas for laser research, and not enough room to describe them all. In many cases, lasers play only a small part in the research effort; in others, the laser is central. Here are a few emerging technologies and ideas.

14.12.1 Quantum Computing and Quantum Cryptography

Quantum optics also promise weird and wonderful capabilities in cryptography and computing. The laws of quantum mechanics tell us that a quantum state, such as the polarization of a photon, does not have a fixed value until we observe it. If that photon shares a common origin with another, and if the two have been produced in ways that connect or “entangle” their properties, observing the properties of one photon can simultaneously determine the value of that property for the second photon. One emerging application is using that quantum entanglement as a way to encrypt data and develop inherently secure data transmission techniques.

Quantum computing involves performing operations on the quantum states of objects. These quantum states have peculiar properties, including that they are nominally superpositions of a number of possible states, which are not resolved until someone observes them. In principle, quantum computing could be very powerful for a few particular applications, notably the factoring of large numbers, which is now used to ensure the security of computer transactions. Researchers are still working on ways to make the concept practical. One requirement is a source of single photons.

14.12.2 Photonic Integration

Chapters 10 and 11 mentioned the possibility of combining lasers with semiconductor electronics for photonic integration on a single substrate. The idea dates back decades, but the reality has long been challenging. Problems remain in combining active light emitters made of III–V compound semiconductors with electronics, largely silicon. But, integrated versions of many optical components have

been integrated, including lasers. This book does not explore the details because they do not directly relate to laser physics, but the emerging technology is important.

14.12.3 Deep-Space Communications and Optical SETI

Communicating through space with light has two big attractions. Laser beams concentrate light in a much narrower angular range than microwave antennas, so less power is wasted transmitting signals out of reach of optical receivers. Light also has much higher frequencies than microwaves, so it can carry more information from space probes. However, realizing those advantages has been difficult because the tight focus of a laser beam also had a big disadvantage—it was hard to keep it focused on a very small spot far away in space. New technology has greatly improved beam direction. NASA has demonstrated laser communications from lunar orbit to the Earth and hopes to use high-speed laser beaming to transmit data from future Mars probes at much higher data rates than now possible.

Lasers might even permit interstellar communications. In 1962, Charles Townes calculated that a 10-kW laser beam focused by a 5-m mirror in space would look brighter than the sun at the laser line when viewed from deep space and visible to the naked eye up to 0.1 light year away. A 5-m telescope could detect the signal up to 100 light years away. A handful of astronomers are scanning the sky with optical telescopes, searching for extraterrestrial intelligence transmitting signals to us on laser lines. They have not spotted ET yet, but neither have astronomers searching at microwave frequencies that many had thought were better candidates for SETI.

14.12.4 Power Beaming and Laser Propulsion

Laser beams can also beam energy through the air to power remote equipment without wires. The beam can deliver energy to solar cells customized to convert the laser light efficiently into electricity, charging batteries or powering equipment directly. PowerLight Technologies of Seattle, formerly LaserMotive, has shown that they can transmit hundreds of watts of power across hundreds of meters, either through the air or through a fiber-optic cable. The laser beam can track a slow-flying drone through the air as well as deliver power to stationary equipment.

Power beaming on a much larger scale, from orbiting solar power satellites to the ground, is an idea that dates back to the energy crisis of the 1970s. Lasers were proposed because they could deliver power to much smaller receivers than microwave downlinks. Advocates of solar power satellites turned to microwaves because they feared intense laser beams might be deliberately misdirected as weapons. However, the solar power satellite idea never caught on because of the difficulty of launching the giant solar arrays that would have been needed.

Powerful laser beams pointed upward might propel spacecraft by evaporating fuel kept on the satellite. One idea is to nudge satellites into different orbits by evaporating ice they carry, so the force from the evaporating gas would push the satellite in the right direction. The Breakthrough Starshot project has proposed a very different sort of power beaming—using ground-based lasers to accelerate ultra-light miniature space probes attached to light sails to speeds up to 100 million miles per hours so they could reach Alpha Centauri in 20 years. Then, they could use other lasers to beam back data on the star system. It is a far-out idea, but an intriguing one.

14.12.5 Lasers to Save the World—Asteroid Defense

Someday, lasers might even save the world in a very different way, by deflecting a wayward asteroid away from the Earth. The radiation pressure delivered by a laser beam is a weak force, but if we knew decades in advance that an asteroid might be a threat, we could point a powerful laser beam toward it. The laser light would push just a tiny bit on the asteroid, but if the asteroid were small and the lead time long, over many years that tiny push could nudge the asteroid's orbit far enough to one side that it missed the Earth. It may sound like science fiction, but it is a much easier, if less dramatic, way to save the world from an asteroid than sending Bruce Willis and a crew of astronauts armed with nuclear weapons at the last minute.

14.13 WHAT HAVE WE LEARNED?

- Research has been a major laser application since the laser was invented.

- Thirty-three people have received laser-related Nobel Prizes.
- Laser-related inventions honored with Nobel prizes are the laser, holography, semiconductor heterostructures (which enabled diode lasers), fiber-optic communications, and blue light-emitting diodes.
- Laser spectroscopy can be used to eliminate the Doppler spreading caused by the random motion of atoms or molecules, giving very sharp spectra as if the particles were not moving.
- Femtosecond spectroscopy can show the progress of a chemical reaction.
- Trains of femtosecond pulses contain a “comb” of uniformly spaced laser frequencies, which can measure frequencies precisely across the optical spectrum.
- Lasers can trap and manipulate tiny objects, a technique called optical tweezers.
- Laser cooling can slow atoms to a virtual halt by draining away their thermal energy.
- Atoms cooled to extremely low temperatures can all occupy the same low-energy state in a Bose–Einstein condensate, making them coherent with each other.
- LIGO, the Laser Interferometer Gravitational-Wave Observatory, detected gravitational waves for the first time and launched gravitational-wave astronomy.
- Laser guide stars and adaptive optics allow giant ground-based telescopes to see sharply through the atmosphere.
- Light can be slowed to a virtual crawl by passing through materials with very high refractive index.
- Laser action has been demonstrated in nano cavities, photonic crystals, surface plasmons, and from single atoms.
- Laser action has been demonstrated in living cells and gelatin.
- Chirped-pulse amplification made it possible to amplify very short pulses to peak powers in the petawatt region.
- Pulses have to have extremely wide bandwidth to be made extremely short. Attosecond pulses are possible only at the high frequencies of the extreme ultraviolet.
- The Extreme Light Infrastructure will allow a new generation of experiments with extremely intense laser pulses lasting femtoseconds.
- The strong electromagnetic fields of short laser pulses can accelerate charged particles over short distances.

- X-ray free-electron lasers reveal new information on materials and crystal structure, valuable in research ranging from protein structure to conditions inside the planets Uranus and Neptune.
- Quantum computing and quantum cryptography are important new fields.
- Photonic integrated circuits combine semiconductor electronics with lasers and optical detectors.
- Laser beams could transmit signals between planets or perhaps between stars.
- High-power lasers might someday be used for propulsion in space and to deflect asteroids away from the Earth.

QUIZ FOR CHAPTER 14

1. What does Doppler-free spectroscopy do?
 - a. Stops gas atoms from moving
 - b. Measures the motion of atoms in an absolute rest frame
 - c. Uses two laser beams pointed in opposite directions to cancel the Doppler spreading caused by the random motion of gas atoms
 - d. Uses a powerful beam to trap atoms so they cannot move and a weak probe beam to measure their spectra
 - e. Produces a Bose–Einstein condensate
2. What do you need to generate a frequency comb?
 - a. A series of very short pulses repeated at regular intervals
 - b. A single pulse stretched across a wide range of frequencies
 - c. A set of narrowband optical filters, each selecting one frequency
 - d. A set of harmonic generators that multiply frequency
 - e. An atomic clock
3. How does laser cooling reduce the temperature of atoms?
 - a. Atoms absorb and reemit photons repeatedly, each time losing a little bit of energy.
 - b. Atoms absorb and reemit photons repeatedly, each time gaining a little bit of energy.
 - c. Two laser beams pointed in opposite directions trap the atoms so they cannot move.
 - d. A single powerful beam traps the atoms so they cannot move.
 - e. Photon pressure pushes the atoms into a refrigerator.

4. What is the essential property that defines a Bose–Einstein condensate?
 - a. Atoms are trapped by a laser beam.
 - b. Atoms are stimulated to emit other atoms.
 - c. Atoms can never escape from it.
 - d. Atoms occupy the same quantum state.
 - e. Atoms freeze into a perfect crystal.
5. How do lasers detect gravitational waves?
 - a. Using Doppler-free spectroscopy
 - b. By measuring tiny displacements of massive mirrors at the ends of long interferometers.
 - c. By observing Bose–Einstein condensates form near absolute zero.
 - d. By timing fluctuations of light with atomic clocks.
 - e. By observing vibrations in the upper atmosphere with laser guide stars.
6. How do you make a nanoscale laser with volume smaller than a cubic wavelength of visible light?
 - a. By trapping a single photon.
 - b. By forcing a single electron to oscillate in a quantum dot.
 - c. By creating compact clusters of electrons called surface plasmons.
 - d. By creating Bose–Einstein condensates.
 - e. It is impossible.
7. Suppose you could compress the duration of a pulse to 1 attosecond. How much energy would the pulse have to contain to reach peak power of a petawatt during that attosecond?
 - a. 1 mJ
 - b. 100 mJ
 - c. 1 J
 - d. 2 J
 - e. 10 J
8. How small a spot do you have to focus a petawatt pulse onto to achieve a peak power of 10^{23} watts per square centimeter?
 - a. 0.3 cm
 - b. 1 mm
 - c. 0.3 mm
 - d. 0.1 mm
 - e. 0.03 mm
 - f. 0.01 mm

9. A laser in orbit around Mars must direct signals to Earth so they can be detected from any point on the surface of the planet facing Mars. Assume the distance to Mars is 75 million kilometers and the Earth's diameter is 12,800 km. What should the beam divergence be?
- 0.00107 milliradian
 - 0.0017 milliradian
 - 0.017 milliradian
 - 0.107 milliradian
 - 0.170 milliradian
10. The laser orbiting Mars has a diffraction-limited beam at 1- μm wavelength. How big must its output mirror be to emit a beam with the divergence you calculated in Question 9? (Approximate divergence as λ/D .)
- 6 mm
 - 6 cm
 - 60 cm
 - 1.7 m
 - 6 m

ANSWERS TO QUIZ QUESTIONS

CHAPTER 1

1. c
2. b
3. a
4. c
5. d
6. b
7. e
8. b
9. c
10. No single correct answer

CHAPTER 2

1. b
2. c
3. b
4. a
5. c
6. d
7. d
8. c
9. d
10. c

CHAPTER 3

1. b
2. a
3. c
4. e
5. c
6. e
7. a
8. b
9. c
10. a
11. b
12. b

CHAPTER 4

1. b
2. e
3. a
4. e
5. c
6. e
7. b

- 8. d
- 9. c
- 10. b
- 11. d
- 12. a

CHAPTER 5

- 1. a
- 2. d
- 3. d
- 4. c
- 5. b
- 6. b
- 7. e
- 8. e
- 9. c
- 10. a
- 11. a
- 12. e

CHAPTER 6

- 1. d
- 2. c
- 3. d
- 4. e
- 5. d
- 6. b
- 7. e
- 8. a
- 9. d
- 10. b
- 11. e
- 12. a

CHAPTER 7

- 1. b
- 2. b

- 3. d
- 4. a
- 5. c
- 6. b
- 7. d
- 8. e
- 9. e
- 10. a
- 11. a
- 12. b

CHAPTER 8

- 1. c
- 2. b
- 3. a
- 4. c
- 5. e
- 6. d
- 7. a
- 8. b
- 9. c
- 10. a
- 11. d
- 12. e
- 13. d
- 14. b

CHAPTER 9

- 1. d
- 2. a
- 3. b
- 4. e
- 5. b
- 6. b
- 7. a
- 8. d
- 9. c

10. a
11. b
12. e

CHAPTER 10

1. b
2. a
3. c
4. d
5. b
6. e
7. d
8. a
9. c
10. d
11. e
12. b
13. e
14. d
15. e

CHAPTER 11

1. a
2. e
3. e
4. e
5. b
6. c
7. d
8. a
9. d
10. a

CHAPTER 12

1. b
2. c

3. a
4. b
5. d
6. a
7. d
8. c
9. d
10. e

CHAPTER 13

1. d
2. c
3. b
4. a
5. a
6. e
7. b
8. d
9. b
10. d
11. e
12. e

CHAPTER 14

1. c
2. a
3. a
4. d
5. b
6. b
7. a
8. e
9. e
10. a

LASER SAFETY

If you are going to work regularly with lasers, you will need to know much more about laser safety than this brief summary can tell you. But if you are going to do anything with lasers, you should understand the key points of laser safety described here. Do not be scared, but do be cautious, and be sure you know what kind of laser you have and its potential dangers before you turn it on.

Like most tools, lasers can be dangerous if they are misused or poorly made. The big issue for lasers is eye safety. Like looking directly at the sun, looking directly into a laser beam can damage the retina, the light-sensitive cells at the back of your eye that you see with. *Never* stare into a laser beam of any kind, even one that seems too feeble to do any harm. Light that your eyes cannot sense can still damage the retina very seriously.

A second rule is to keep your fingers out of laser power supplies unless you know what you are getting into. You cannot hurt yourself changing the batteries in a pocket laser pointer, but power supplies for other lasers can electrocute you, like household line current.

A.1 POWER SUPPLY HAZARDS

Electrically, the only lasers you can consider harmless are battery-powered laser pointers. Anything that plugs into household line current or high-voltage lines is potentially dangerous. Line current can kill you. Gas lasers that plug into line current often include

transformers that generate very high voltages (thousands of volts or more) to power a discharge in the laser gas.

Pulsed lasers often include capacitors that store a powerful electric charge to excite the laser medium or to power a flash lamp. Capacitors retain a high-voltage charge for some time after the power is switched off or disconnected. Learn how to handle electricity safely before you poke around inside the laser or the power supply.

Modern commercial lasers are built to comply with electrical safety codes, so you should not encounter exposed high voltages unless you open the case. However, you cannot assume dangerous voltages are safely concealed in a laser laboratory or workshop.

A.2 LASER BEAM HAZARDS

The human eye has evolved to be very sensitive to visible light, and this sensitivity makes it extremely vulnerable to excess light levels. The brightest light in nature is the sun, and we have developed an aversion response that makes our eyes instinctively look away if we happen to glance at the sun. There is good reason for this, because staring directly at the sun can permanently damage the retina at the back of the eye, impairing vision for the rest of your life.

Lasers that emit significant beam power into the open air are required to carry warning labels. Optical disk players do not need warning labels because the beam does not go outside the case. Supermarket scanners do not need warning levels because the beam going into the air is weak and moves continually so you could not stare into it. Laser pointers and other lasers that emit more than a milliwatt directly into the air *should* bear warning labels.

The hazards from a milliwatt-level visible laser beam are similar to those from looking directly at the sun. Although the laser delivers only a tiny fraction of the power delivered by the sun, the light rays in a laser beam are all parallel to each other, and the beam is small enough to pass through the pupil of your eye, so the lens focuses all that light into a tiny spot on the retina. A momentary accidental glance into a 1-mW laser beam will dazzle your eyes, but it will not cause permanent damage unless you stare into the beam.

The most vulnerable part of your eye is the retina, the layer in the back of your eye that senses light. If you have normal color

vision, your eye responds weakly to infrared light at wavelengths beyond 700 nm, with the response declining as the wavelength increases, so you can see a little light at the 780 nm wavelength emitted by the diode lasers in CD players. However, the rest of the eye transmits wavelengths as long as about 1400 nm to the retina, so those invisible wavelengths are as much of a hazard to your retina as the laser light you can see. Those wavelengths are particularly dangerous because you cannot see them, so your eye does not instinctively turn away from the light, as it does if you accidentally turn toward the sun.

Ultraviolet wavelengths shorter than 400 nm and infrared wavelengths longer than about 1400 nm do not penetrate deeply into the eye. They can damage the cornea, the surface layer of the eye, and the lens just underneath the cornea, but the power thresholds for damage to the cornea and lens are higher than for damage to the retina.

You do not have to be looking at the laser to suffer an eye injury. One common cause of accidents is reflection of the beam off a shiny object in an unexpected direction, hitting someone in the eye who is looking away from the laser.

The best way to prevent such accidents is to wear special laser safety goggles made of glass or plastic that block selected laser wavelengths. The goggles must be matched to the lasers being used. It does no good to wear goggles that block a wavelength your laser does not emit.

Let us look briefly at the types of lasers and the recommended precautions.

A.2.1 Lasers Packaged inside Equipment

Lasers are packaged inside a wide range of equipment, and in most of them the beam is always contained safely inside the case. You will never see the beam from a laser printer or a CD, DVD, or Blu-Ray disk player unless you take them apart, so this equipment requires no special warnings or precautions.

Laser scanners used in supermarkets are an exception because the beam must emerge from the scanner in order to read the bar codes on packages. However, the scanners are designed so the beam power is very low and the beam scans so fast that it does not emit enough light into the air that it requires a warning label. Engineers worked very hard to achieve this goal, because they did not think

customers would be eager to shop at stores in which the checkout counter bore a “DANGER: LASER RADIATION” sign.

A.2.2 Battery-Powered Laser Pointers

The type of lasers you are most likely to encounter as separate devices are pointers. United States laser safety regulations limit laser pointers to power levels less than 5 mW. Even at that level, they are required to bear warning labels that say “Danger—Laser Radiation, Avoid Direct Eye Exposure,” identify the maximum power level, and specify the class of laser product according to rules specified by the Food and Drug Administration (FDA).

Red laser pointers use diode lasers and are unlikely to exceed their 5-mW rating, so they are safe if you use common sense and do not point them at anybody’s eyes.

Green laser pointers pose somewhat different issues. They look much brighter to the eye because the human eye is much more sensitive to green light than to red light, although the eye hazard from a green pointer at a given power level is similar to that from a red pointer. However, green pointers may exceed their rated output level, and some with powers in the 100-mW range are sold over the Internet in pointer-like packages, without standard safety labels or equipment. Because green pointers are based on diode-pumped, frequency-doubled neodymium lasers, near-infrared light at 808 and 1064 nm is present inside the laser. Better pointers include filters that block the infrared light from exiting in the beam, but tests at the National Institute of Standards and Technology have shown that many green pointers lack filters, so their output includes those invisible infrared wavelengths, which pose additional hazards as mentioned above.

Thus, green pointers in the 5-mW range deserve somewhat more care than red pointers; those with higher output power deserve much more care. When a green pointer is used to identify objects in the sky, special care should be taken to avoid shining it at aircraft because the bright light could flash-blind a pilot and get you in serious trouble.

Blue and violet laser pointers are also widely available, and these pose different hazards. True blue diode lasers emit at about 450 nm, where the eye is fairly sensitive, and which the eye perceives as sky blue. However, 405-nm violet diode lasers are sometimes sold as “blue” although the color is better described as violet.

The human eye cannot see 405 nm light well, so what looks like a dim violet beam when the laser spot is on some surfaces actually is quite bright (and hence more dangerous). You can see that by aiming the blue/violet laser at paper, which generally fluoresces brightly, producing a white spot that looks much brighter than it does on a non-fluorescent object like aluminum. (Be sure if you point it at metal to aim it so the reflection will not hit you in the eye.) That bright fluorescent spot is the real brightness of the beam.

A.2.3 Scientific, Industrial, and Medical Lasers

If you are working with scientific, industrial, or medical lasers, they are likely to pose more serious hazards. These lasers are classified according to potential hazards by safety regulating agencies. In the United States, the agency responsible for equipment safety is the Center for Devices and Radiological Health in the FDA. Table A-1 shows the FDA classes, the maximum power level, and the warning

Table A-1. Laser safety classification and warning requirements in the United States for continuous-wave lasers at 400 to 1400 nm, somewhat simplified

Safety class	Max CW power accessible	Warning class	Label text	Special requirements
1	Up to 0.39 mW	None required	None required	None
IIa	0.39–1.0 mW	None required if exposure under 1000 seconds	None required	None
II	Up to 1 mW	CAUTION, if potential exposure over 1000 seconds	Laser radiation—do not stare into beam	None
IIIa	Up to 5 mW Irradiance over 2.5 mW/cm ²	CAUTION	Laser radiation—avoid direct eye exposure	
IIIb	Up to 500 mW	DANGER	Laser radiation—avoid direct exposure to beam	Key interlock, emission indicator
IV	Over 500 mW	DANGER	Laser radiation—avoid eye or skin exposure to beam	Key interlock, emission indicator; cannot operate with cover off

label requirements for lasers emitting continuous beams at 400 to 1400 nm. The U.S. Occupational Safety and Health Administration is responsible for safe operation of lasers in the workplace.

As you might expect, the hazards become more serious as power levels increase, and operating requirements become more stringent. The need for an emission indicator might seem odd, but remember that many laser beams are not visible, or are not emitted continuously, so people working nearby might otherwise not know the laser was in operation.

You are not likely to encounter class IIIb or IV lasers other than in an academic, medical, or industrial laser laboratory. These are the places where laser goggles are a must. Goggles are designed to protect eyes from accidental exposure both from the front and from the sides. They block only specific laser wavelengths or bands; for example, the blue–green band that includes the 488- and 514.5-nm argon-ion lines and the 532-nm line of doubled neodymium emitted by green pointers. The goggles should transmit as much light as possible at other wavelengths so the wearer can see well enough to work safely in the laboratory.

Anyone using lasers should be familiar with the basics of laser safety, but if you are going to work with class IIIb or IV lasers, be sure to spend some time getting to know the details, particularly for the types of lasers you will be working with. The web links in Table A-2 are starting points.

Table A-2. Laser safety web resources

British laser safety advice https://www.gov.uk/government/publications/laser-radiation-safety-advice/laser-radiation-safety-advice
International Laser Safety Conference https://www.lia.org/conferences/ilsc
Laser Product Safety LLC, private company with training programs and other information http://www.laserproductsafety.com/lpssr.htm
Rockwell Laser Industries, private company with collection of laser safety articles online https://www.rli.com/Resources/articles.aspx
Sam's Laser FAQ—Laser Safety http://www.repairfaq.org/sam/lasersaf.htm
U.S. Food and Drug Administration Laser Products and Instruments https://www.fda.gov/Radiation-EmittingProducts/RadiationEmittingProductsandProcedures/HomeBusinessandEntertainment/LaserProductsandInstruments/default.htm
U.S. Occupational Safety and Health Administration on Laser Hazards https://www.osha.gov/SLTC/laserhazards/index.html

HANDY NUMBERS AND FORMULAS

PHYSICAL CONSTANTS

Boltzmann constant (k): 1.380658×10^{-23} J/K (joule/Kelvin)

Boltzmann constant (k): 8.617385×10^{-5} eV/K

Electron mass: $9.1093897 \times 10^{-31}$ kilogram

Planck's constant (h): 6.626070×10^{-34} joule-second

Planck's constant (h): 4.135667×10^{-15} electronvolt-second

Proton mass: $1.6726231 \times 10^{-27}$ kilogram

Rydberg constant (R): 1.0973731534×10^7 per meter

Rydberg constant (R): 13.6056981 eV

Speed of light in vacuum (c): 2.99792458×10^8 meters/sec

CONVERSIONS

Frequency in terms of wavelength and speed of light: $\nu = c/\lambda$.

Wavelength in terms of frequency and speed of light: $\lambda = c/\nu$.

Electronvolts to joules: $1 \text{ eV} = 1.60217733 \times 10^{-19}$ joule

Photon energy and frequency: $E = h\nu$

Photon energy in joules: $E = h\nu = 6.63 \times 10^{-34} \nu$

A 1-eV photon has frequency: $2.41798836 \times 10^{14}$ hertz

Photon energy and wavelength: $E = \frac{hc}{\lambda}$

A 1-eV photon has wavelength: 1.2399 μm

SYMBOLS TO REMEMBER λ (Greek lambda): wavelength n : refractive index ν (Greek nu): frequency ω (Greek omega): angular frequency = $2\pi\nu$ c : speed of light h : Planck's constant Δ (Greek delta): change or increment θ (Greek theta): angle**IMPORTANT FORMULAS**Angle of refraction from Snell's law: $\theta_2 = \arcsin\left(\frac{n_1 \sin \theta_1}{n_2}\right)$ Approximate diffraction limit in radians for aperture D : λ/D Decibel ratio of two powers: $dB = 10 \times \log_{10}\left(\frac{P_{\text{out}}}{P_{\text{in}}}\right)$ Far-field beam diameter at distance D for beam with divergence θ : $2D \times \tan \theta$ Frequency as a function of speed of light and wavelength: $\nu = c/\lambda$ Fresnel reflection at surface: $R = \left(\frac{n_1 - n_2}{n_1 + n_2}\right)^2$ Ratio of populations in states at energies E_1 and E_2 :

$$\frac{N_2}{N_1} = \exp\left[\frac{-(E_2 - E_1)}{kT}\right]$$

Resonance condition in cavity of length L and refractive index n : $2nL = N\lambda$ Snell's law of refraction: $n_1 \sin \theta_1 = n_2 \sin \theta_2$ Wavelength as a function of speed of light and frequency: $\lambda = c/\nu$

Metric unit prefixes and their meanings

Prefix	Symbol	Multiple
exa	E	10^{18} (quintillion)
peta	P	10^{15} (quadrillion)
tera	T	10^{12} (trillion)
giga	G	10^9 (billion)
mega	M	10^6 (million)
kilo	k	10^3 (thousand)
hecto	h	10^2 (hundred)
deca	da	10^1 (ten)
deci	d	10^{-1} (tenth)
centi	c	10^{-2} (hundredth)
milli	m	10^{-3} (thousandth)
micro	μ	10^{-6} (millionth)
nano	n	10^{-9} (billionth)
pico	P	10^{-12} (trillionth)
femto	f	10^{-15} (quadrillionth)
atto	a	10^{-18} (quintillionth)

RESOURCES AND SUGGESTED READINGS

FURTHER READING AND RESOURCES

Recommended Online References and News Sources

RP Photonics Online Encyclopedia <https://www.rp-photonics.com/encyclopedia.html> (A comprehensive online encyclopedia of photonics, it's my top laser reference source on the web)

Sam's Laser FAQ, <http://www.repairfaq.org/sam/lasersam.htm> (A vast "practical guide to lasers" written for experimenters and hobbyists that is a great help to newcomers)

Sam's Safety Guidelines for high-voltage and/or line-powered equipment, <http://www.repairfaq.org/sam/safety.htm>

Sam's Laser FAQ–Laser Safety, <http://www.repairfaq.org/sam/lasersaf.htm>

optics.org <http://optics.org/> (news site and archive run by SPIE)

Laser Focus World magazine <http://www.laserfocusworld.com>

Photonics Spectra magazine <https://www.photonics.com/>

Updates to this book <http://www.understandinglasers.com>

Introductory Books on Lasers and Optics

J. Warren Blaker and Peter Schaeffer, *Optics: An Introduction for Technicians and Technologists* (Prentice-Hall, Upper Saddle River, NJ, 2000)

Jean-Claude Diels and Ladan Arissian, *The Power and Precision of Light* (Wiley, Weinheim, 2012)

Galen C. Duree, *Optics for Dummies* (Amazon Digital, 2011)

Grant R. Fowles, *Introduction to Modern Optics* (Dover, New York, 1989)

C. Breck Hitz, J. J. Ewing, and Jeff Hecht, *Introduction to Laser Technology*, 4th ed. (IEEE Press/Wiley, Hoboken, NJ, 2012)

Orazio Svelto, *Principles of Lasers*, 5th ed. (Springer, New York, 2009)

Advanced Books on Lasers and Optics

Liang Dong and Bryce Samson, *Fiber Lasers: Basics, Technology and Applications* (CRC Press, 2016)

Understanding Lasers: An Entry-Level Guide, Fourth Edition. Jeff Hecht.

© 2019 by The Institute of Electrical and Electronics Engineers, Inc. Published 2019 by John Wiley & Sons, Inc.

- Eugene Hecht, *Optics*, 5th ed. (Pearson, 2016)
- Walter Koechner, *Solid-State Laser Engineering*, 6th ed. (Springer-Verlag, Berlin, 2011)
- Rudiger Paschotta, *Encyclopedia of Laser Physics and Technology* (Wiley, Weinheim, 2008); also available online
- Frank L. Pedrotti, *Introduction to Optics*, 3rd ed. (TBS, 2006)
- Richard S. Quimby, *Photonics and Lasers: An Introduction* (Wiley, Worcester, MA, 2006)
- Bahaa E. A. Saleh and Malvin Carl Teich, *Fundamentals of Photonics*, 2nd ed. (Wiley-Interscience, Hoboken, NJ, 2007)
- Anthony Siegman, *Lasers* (University Science Books, Mill Valley, CA, 1986)
- Warren J. Smith, *Modern Optical Engineering*, 4th ed. (McGraw-Hill, New York, 2007)

History of Laser Development

- Joan Lisa Bromberg, *The Laser in America 1950–1970* (MIT Press, Cambridge, MA, 1991)
- Jeff Hecht, *Beam: The Race to Make the Laser* (Oxford University Press, New York, 2005)
- Theodore Maiman, *The Laser Odyssey* (Laser Press, Blaine, WA, 2000)
- Nick Taylor, *Laser: The Inventor, the Nobel Laureate, and the 30-Year Patent War* (Simon & Schuster, New York, 2000)
- Charles H. Townes, *How the Laser Happened: Adventures of a Scientist* (Oxford University Press, New York, 1999)

Laser Applications

- Gerhard K. Ackermann and Jürgen Eichler, *Holography: A Practical Approach* (Wiley-VCH, Hoboken, NJ, 2007)
- Paul R. Berman and Vladimir S. Malinovsky, *Principles of Laser Spectroscopy and Quantum Optics* (Princeton University Press, 2011)
- Jean-Claude Diels and Wolfgang Rudolf, *Ultrashort Laser Pulse Phenomena*, 2nd ed. (Academic Press, San Diego, CA, 2006)
- Judith F. Donnelly and Nicholas Massa, *Light: Introduction to Optics and Photonics* (Photonics, 2018)
- Jeff Hecht, *Beam Weapons: Roots of Reagan's Star Wars* (Laser Light Press, 2016)
- Jeff Hecht, *Lasers, Death Rays, and the Long, Strange Quest for the Ultimate Weapon* (Prometheus Books, 2019)
- Jeff Hecht, *Understanding Fiber Optics*, 5th ed. (Pearson/Prentice-Hall, Upper Saddle River, NJ, 2006)
- John C. Ion, *Laser Processing of Engineering Materials: Principles, Procedure, and Industrial Application* (Butterworth-Heinemann, 2011)
- National Research Council, *Optics and Photonics: Essential Technologies for Our Nation* (National Academies Press, Washington, DC, 2014)
- Markolf H. Niemz, *Laser-Tissue Interactions: Fundamentals and Applications*, 3rd ed. (Springer, New York, 2007)
- John F. Ready, *Industrial Applications of Lasers* (Academic Press, San Diego, CA, 1997)
- William M. Steen and Kenneth Watkins, *Laser Material Processing*, 4th ed. (Springer, New York, 2010)

Laser Safety

- Ken Barat, *Laser Safety in the Lab* (SPIE Press, 2012)
- Ken Barat, *Laser Safety: Tools and Training*, 2nd ed. (CRC Press, 2014)
- Roy Henderson and Karl Schulmeister, *Laser Safety* (Institute of Physics Publishing, Bristol and Philadelphia, 2004)
- Laser Institute of America, *Laser Safety Guide* (Laser Institute of America, 2015)

Selected Scholarly Journals

Advances in Optics & Photonics (reviews and tutorials)

IEEE Journal of Quantum Electronics

IEEE Journal of Selected Topics in Quantum Electronics

IEEE LEOS Newsletter

IEEE Photonics Technology Letters

Journal of Lightwave Technology

Laser & Photonics Reviews

Optica (open access)

Optical Engineering

Optics & Photonics News

Optics Express (open access)

Optics Letters

Professional Societies

European Optical Society, <http://www.myeos.org/>

Institute of Electrical and Electronic Engineers, <http://www.ieee.org>

IEEE Photonics Society, <https://www.photonicsociety.org/>

Laser Institute of America, <https://www.lia.org/>

Optical Society of America, <http://www.osa.org>

SPIE, <http://www.spie.org>

GLOSSARY

- Achromatic** Made to be without color. Achromatic lenses are compound lenses including lenses made from two or more different glass compositions arranged so they focus all colors of visible light to the same point.
- Acousto-optic** Interaction between an acoustic wave and a light wave, used in beam deflectors, modulators, and Q switches.
- Active medium** The light-emitting material in a laser, sometimes used to identify a specific element in a crystal.
- Active species** The atoms or molecules producing stimulated emission in a laser.
- Alexandrite** A synthetic crystal doped with chromium to form a tunable solid-state laser that emits near-infrared light.
- Amplification** Increasing the amount of light or the strength of a signal. When done optically, it usually occurs by stimulated emission and grows exponentially.
- Amplified spontaneous emission** Amplification of spontaneous emission without laser oscillation in a cavity without the feedback needed to reach laser threshold. This can be a problem in optical amplifiers. It produces less coherent light than laser emission, which has advantages for some applications. Called *superluminescence* in LEDs.
- Amplifier** An optical device that increases the power of an input optical signal by stimulated emission, but which lacks resonator mirrors.

- Angstrom (Å)** A unit of length, 0.1 nanometer or 10^{-10} meter, abbreviated Å. It is not a standard SI unit, but is often used to measure wavelength in the visible spectrum.
- Arc lamp** A high-intensity lamp in which an electric discharge continuously produces light.
- Attenuation** Reduction of light intensity or loss. It can be measured in units of optical density or decibels as well as percentage or fraction of light lost.
- Attenuator** An optical element that transmits only a given fraction of incident light.
- Average power** The average level of power in a series of pulses. It equals pulse energy times the number of pulses divided by the time interval.
- Band gap** The gap between valence band and conduction band in a semiconductor.
- Beam diameter** The distance between the edges of the beam, which are defined as the distance from the center where power drops to a certain level, often $1/e^2$ of the central power.
- Beam divergence** The angle at which a beam spreads.
- Beam profile** The pattern of power across a laser beam.
- Beam splitter** A device that divides incident light into two separate beams, one reflected and one transmitted.
- Birefringent** Has a refractive index that differs for vertically and horizontally polarized light.
- Brewster's angle** The angle at which a surface does not reflect light of one linear polarization.
- Cavity dumping** Inserting a mirror into a high-Q laser cavity to dump the power circulating in the cavity out of the laser.
- Chalcogenide** A material containing compounds of sulfur, selenium, or tellurium.
- Chemical laser** A laser that is excited by a chemical reaction. The most common type produces hydrogen fluoride or deuterium fluoride.
- Chirped pulse** A pulse that changes in wavelength during its length, so longer wavelengths precede shorter wavelengths or vice versa, depending on how the refractive index of the material changes with wavelength.
- Chirped pulse amplification** An amplification technique that deliberately introduces chirp to stretch the duration of pulses to reduce their peak power so they can be amplified without to higher energy. Optics can be added to reverse the chirp and

compress the amplified pulse to reach extremely high peak power.

Chromatic aberration Bringing light of different colors to different focal points, an optical aberration arising from the variation of refractive index with wavelength.

Coating Material applied in one or more layers to the surface of an optical element to change the way it reflects or transmits light.

Coherence Alignment of the phase and wavelength of light waves with respect to each other. If the waves are perfectly aligned, the light is *coherent*.

Collimate Make light rays parallel.

Compound semiconductor A semiconductor that is a chemical compound of two or more elements, such as gallium arsenide (GaAs) or gallium indium nitride (GaInN).

Concave Curving inward, so that the central parts are deeper than the outside, like the inside of a bowl.

Conduction band Energy level in a solid in which electrons are not bound to individual atoms but are free to carry current through the solid.

Confocal Having the focal point of two mirrors at the same place. Confocal resonators have concave end mirrors that have the same focal point in the middle of the laser cavity.

Continuous wave Emitting a steady beam.

Convex Curving like the outside of a ball so the outer parts are lower than the center.

Crystals Solids in which atoms are arranged in regular structures of fixed composition.

Cycles per second Number of oscillations a wave makes; the frequency. 1 hertz = 1 cycle per second.

Cylindrical lens A lens that has a cylindrical surface profile so it refracts light in one direction but not in the perpendicular direction.

Decibel A logarithmic comparison of power levels, abbreviated dB and defined as the value $10 \log(P_2/P_1)$, or 10 times the base-10 logarithm of the ratio of the two power levels.

Detector A device that generates an electric signal when illuminated by light.

Dielectric Electrically nonconductive or insulating.

Difference frequency An output wave with frequency that equals the difference of the frequencies of two input waves.

- Diffraction** Scattering or spreading of light waves when they pass an edge.
- Diffraction limit** The minimum possible spreading of light, a function of the wavelength divided by emitting area (λ/D).
- Diffuse (reflection)** Reflection of light by a surface that is uneven on an atomic level, like a painted white wall, which scatters the light and does not yield a mirror-like image.
- Diode** An electronic device that preferentially conducts current in one direction but not in the other. Semiconductor diodes contain a p - n junction between regions of different doping, which lets current flow in one direction but not the other. Diodes can emit light (e.g., laser diodes) or detect it (photodiodes).
- Diode laser** A semiconductor laser in which a current flowing through the device causes electrons and holes to recombine at the junction of p - and n -doped regions, where stimulated emission (laser action) takes place.
- Diode pumping** Using a diode laser as the pump source for an optically pumped laser.
- Direct bandgap** A band gap that allows electrons to drop from the conduction band to the valence band without changing their momentum, which is essential for efficient light emission.
- Direct modulation** Control of diode laser output by modulating the drive current.
- Dispersion** The spreading out of light by color, arising from the variation of refractive index with wavelength.
- Distributed feedback** Using structures such as a series of etched lines in a substrate rather than a simple mirror to provide feedback in a laser.
- Divergence** The angular spreading of a laser beam with distance.
- Doppler broadening** Spreading of absorption or emission lines caused by motion of gas atoms or molecules.
- Duty cycle** The fraction of time a laser spends emitting light.
- DWDM** Dense wavelength-division multiplexing, WDM with closely spaced wavelengths.
- EDFA** Erbium-doped fiber amplifier.
- Electro-optic** The interaction of light and electric fields, typically changing the light wave. Used in some modulators, Q switches, and beam deflectors.
- Electromagnetic radiation** Waves moving at the speed of light and made up of oscillating electrical and magnetic fields

perpendicular to one another. Also behaves as photons or quanta of electromagnetic energy. Electromagnetic radiation includes radio waves, microwaves, infrared, visible light, ultraviolet, X-rays, and gamma rays.

Electromagnetic spectrum The range of wavelengths or frequencies at which electromagnetic radiation is emitted.

Electronic transition Change in energy level of an electron.

Etalon An optical resonator, typically with a short distance between reflective surfaces, placed inside laser resonators to restrict the range of wavelengths.

Excimer laser A pulsed ultraviolet laser in which the active medium is a short-lived molecule containing a rare gas such as xenon and a halogen such as chlorine.

Exciton An electron-hole pair in a semiconductor; the two are bound to each other, but the electron has not dropped from the conduction band to fill the hole in the valence band.

Extreme ultraviolet Wavelengths from 10 to 200 nm, absorbed very strongly by air.

Far-infrared laser One of a family of gas lasers emitting light at the far-infrared wavelengths of 30 to 1000 micrometers.

Fiber laser A solid-state laser in which the gain medium is an optical fiber with a light-guiding core and one or more cladding layers.

Flashlamp A lamp designed to produce intense flashes of light, can be used for optical pumping of lasers.

Focal length Distance from the center of a lens to the point where it focuses parallel light rays.

Forward bias Voltage applied across a diode so it carries current easily.

Front-surface mirror A mirror with a reflective coating on its front surface so light reflects directly off the coating without first passing through glass.

Free-electron laser A laser in which stimulated emission comes from electrons passing through a magnetic field that varies in space.

Frequency For light waves, the number of wave peaks per second passing a point. Measured in hertz, or cycles per second.

Fused silica Synthetic silica (SiO_2) formed from highly purified materials.

Gain The amount of amplification, usually measured per unit length.

- Gain bandwidth** Range of wavelengths over which a laser medium has gain.
- Gallium aluminum arsenide** A semiconductor used in LEDs, diode lasers, and certain detectors. Chemically, $\text{Ga}_{1-x}\text{Al}_x\text{As}$, where x is a number less than one.
- Gallium arsenide** A semiconductor used in LEDs, laser diodes, detectors, and electronic components; chemically, GaAs.
- Gallium nitride (GaN)** A semiconductor used in blue, violet, and ultraviolet lasers and LEDs; often contains indium.
- Gas laser** A laser in which the active medium is a gas contained inside a hollow tube, usually powered by an electric discharge through the gas.
- Glass** An amorphous solid, typically made mostly of silica (SiO_2) unless otherwise identified. Silica glasses transmit visible light.
- Harmonic generation** Multiplication of the frequency of a light wave.
- Hertz** Frequency in cycles per second.
- Heterojunction** A boundary between semiconductors that differ in composition, such as GaAs and GaAlAs.
- Index of refraction** The ratio of the speed of light in vacuum to the speed of light in a material, a crucial measure of a material's optical characteristics; usually denoted as n .
- Indium gallium arsenide phosphide** A semiconductor used in lasers, LEDs, and detectors. The band gap and, hence, the wavelength emitted by light sources and detected by detectors depends on the mixture of the four elements. Abbreviated InGaAsP.
- Inertial confinement fusion** Nuclear fusion produced by focusing powerful laser pulses to implode targets containing hydrogen fuel.
- Infrared** Invisible wavelengths longer than 700 nm and shorter than about one millimeter. Wavelengths from about 700 to 2000 nm are called the near infrared.
- InGaAsP** Indium gallium arsenide phosphide, a semiconductor compound used in light sources and detectors. The composition is often written $\text{In}_{1-x}\text{Ga}_x\text{As}_{1-y}\text{P}_y$, where x and y are numbers less than one.
- Injection laser** Another name for the semiconductor laser, derived from the fact that current carriers are injected to produce light.
- Integrated optics** Optical elements analogous to integrated electronic circuits, with multiple devices on one substrate.

- Intensity** Properly, power per unit solid angle, but often used instead of *irradiance* for power per unit area.
- Interference** The addition of the amplitudes of light waves. In destructive interference the waves cancel; in constructive interference, they combine to make more intense light.
- Inversion** see **Population inversion**.
- Ion laser** A laser in which the active laser species is an ionized gas. Often used to describe argon-ion lasers.
- Ionization** Removal of one or more electrons from a neutral atom, leaving it with a positive charge.
- Irradiance** Power per unit area.
- Joule** Unit of energy equal to one watt of power delivered for one second.
- Junction** The boundary between *p*- and *n*-type materials in a semiconductor, where positive and negative carriers recombine.
- Junction laser** A semiconductor diode laser.
- Laser** Acronym of light amplification by stimulated emission of radiation; one of the wide range of devices that generate light in a resonant cavity by that principle. Laser light is directional, covers a narrow range of wavelengths, and is more coherent than ordinary light.
- Laser diode** see **Diode laser**
- Lattice constant** Atomic spacing in a semiconductor crystal.
- LED** Light-emitting diode; a semiconductor diode that emits incoherent light by spontaneous emission.
- Light** Strictly speaking, light is electromagnetic radiation visible to the human eye. Commonly, the term is applied to electromagnetic radiation close to the visible spectrum that acts similarly, including the near infrared and near ultraviolet.
- Light-emitting diode (LED)** A semiconductor diode that produces incoherent light by spontaneous emission.
- Longitudinal modes** Oscillation modes of a laser along the length of its cavity, so twice the length of the cavity equals an integral number of wavelengths. Distinct from transverse modes, which are across the width of the cavity.
- Lower laser level** The lower level in a laser transition.
- Maser** Microwave analog of a laser, an acronym for microwave amplification by stimulated emission of radiation.
- Master oscillator power amplifier (MOPA)** A multistage pulsed laser, with a master oscillator the first stage, followed by one

or more amplifiers to produce higher powers than an oscillator can generate on its own.

Metastable An excited energy level that has an unusually long lifetime.

Micrometer One millionth of a meter, abbreviated μm . The μ is a Greek mu.

Mode A manner of oscillation in a laser. Modes can be longitudinal or transverse.

Mode locking The locking together of many longitudinal modes in a laser cavity, producing a series of short laser pulses. Mode locking can be visualized as clumping together a group of photons that bounce back and forth in the laser cavity.

Monochromatic Containing only a single color, wavelength, or frequency.

Monochromator A light source that selects a narrow slice of the spectrum from a light source with a broad spectrum.

Multimode Containing multiple modes of light. Typically refers to lasers that operate in two or more transverse modes.

***n* region** Part of a semiconductor doped so it has an excess of electrons as (negative) current carriers.

Nanometer A unit of length equal to 10^{-9} meter. Its commonest use is to measure visible wavelengths.

Nanosecond A billionth of a second, 10^{-9} second.

Near infrared The part of the infrared nearest the visible spectrum, typically 700 to 1500 or 2000 nm, but not rigidly defined.

Negative lens A lens that spreads out or diverges light.

Nonlinear effect An effect proportional to the second or higher power of the input, as distinct from linear effects proportional to the first power of the input.

Nonlinear optics Optics that produce nonlinear effects.

Normal (angle) Perpendicular to a surface.

Optical amplifier A laser without a resonant cavity; a device that amplifies light by stimulated emission but does not resonate.

Optical density Attenuation of light passing through a material, defined as $-\log_{10}(\text{Output/Input})$.

Optical parametric amplifiers Broadly tunable nonlinear light sources that amplify light by a three-wave-mixing process.

Optical parametric oscillators Broadly tunable nonlinear oscillators that generate light by a three-wave-mixing process.

Optical pumping Using light to excite atoms or molecules to higher energy levels.

- Optically pumped semiconductor laser (OPSL)** A semiconductor laser that lacks the internal structures needed to control current flow and is pumped optically rather than electrically.
- Orthogonal** Perpendicular.
- Oscillator** A laser cavity with mirrors so stimulated emission can oscillate within it. To many purists, only an oscillator can be a laser.
- Output mirror** The mirror through which a laser emits its beam.
- Oxide glass** A glass made from oxide compounds, usually including silica.
- p region** Part of a semiconductor doped with electron acceptors in which positive “holes” (vacancies in the valence electron level) are the dominant current carriers.
- Peak power** Highest instantaneous power level in a pulse.
- Phase** Position of a wave in its oscillation cycle.
- Phase matching** Precisely matching the phases of fundamental and harmonic frequencies in a nonlinear material to increase efficiency of the process and output power.
- Photoconductive** An optical detector that conducts electric current in proportion to the amount of light incident on it.
- Photodetector** A light detector.
- Photodiode** Usually, a semiconductor diode that produces an electrical signal proportional to light falling upon it. There are also vacuum photodiodes that can detect light.
- Photometer** An instrument for measuring the amount of light visible to the human eye.
- Photons** Quanta of electromagnetic radiation. Light can be viewed as either a wave or a series of photons.
- Photonics** The science and technology of using photons, in the same sense that electronics is the science and technology of using electrons. It often refers to active optical devices such as lasers and detectors, which manipulate light when it is acting more as a photon than as a wave, but can also be used as a synonym for optics.
- Polarization** Alignment of the electric and magnetic fields that make up an electromagnetic wave. Normally, this refers to the electric field. If light waves all have a particular polarization pattern, they are called polarized.
- Polarization vector** A vector indicating the direction of the electric field in an electromagnetic wave.
- Polarizer** A device that transmits light of only one polarization.

Population inversion The condition when more atoms are in an upper energy level than in a lower one. A population inversion is needed for laser action.

Positive lens A lens that focuses light to a point.

Power The flow of energy per unit time. It is measured in watts.

Pump efficiency Fraction of the energy pumping a laser that emerges in the output beam. For optical pumping, this means the fraction of the pump light, not counting losses in turning input electricity into the pump light.

Pumping The way a laser gets the energy to produce a population inversion.

Q factor Quality factor of a resonant laser cavity, a measure of loss within cavity.

Q switch A device that changes the Q (quality factor) of a laser cavity to produce a short, powerful pulse.

Quantum Divided into discrete pieces or levels. A photon is a quantum of light energy. Quantum levels are discrete states with specific values of energy.

Quantum cascade laser A semiconductor laser in which energy is extracted in a series of steps as electrons pass through a series of quantum wells.

Quantum defect In optical pumping, the fraction of the energy in the pump photon that is lost and does not emerge in the laser photon.

Quantum efficiency For optical pumping, the fraction of the energy of a pump photon that emerges in the laser photon.

Quantum well A thin layer in a semiconductor diode with smaller bandgap than the layers above and below it, which traps electrons that lack the energy needed in the adjacent higher-bandgap layers.

Quartz A natural crystalline form of silica (SiO_2).

Quaternary A compound made of four elements, for example, InGaAsP.

Radian A unit of angular measure; 2π radians equals 360° , a circle.

Radiant flux Instantaneous power level in watts.

Radiometer An instrument to measure power (watts) in electromagnetic radiation. Distinct from a photometer, which measures light perceived by the human eye.

Raman scattering Scattering of light by an atom that changes its vibrational state during the scattering process, so the scattered wavelength is shifted from that of the input photon.

- Rayleigh range** The distance over which light rays remain parallel after exiting a laser.
- Rays** Straight lines that represent the path taken by light.
- Real image** An image that can be projected onto a surface.
- Recombination** In a semiconductor laser, dropping of a conduction-band electron into a vacancy in the valence band of an atom. Free electrons can also recombine with gas atoms by dropping into a bound energy state.
- Refraction** The bending of light as it passes between materials of different refractive index.
- Refractive index** The ratio of the speed of light in vacuum to the speed of light in a material, a crucial measure of a material's optical characteristics. Abbreviated n .
- Repetition rate** The number of pulses per second.
- Resonator** A region with mirrors on the ends and a laser medium in the middle. Stimulated emission from the laser medium resonates between the mirrors, one of which lets some light emerge as a laser beam.
- Retroreflector** An optical device that reflects incident light back in precisely the direction from which it came. It is typically a prism or a cluster of three mirrors forming the corner of a cube.
- Reverse bias** Voltage applied across a diode so it does not carry current.
- Second harmonic** A wave at twice the frequency (or half the wavelength) of the fundamental wave.
- Semiconductor laser** A laser in which the active medium is a semiconductor. Most but not all semiconductor lasers are semiconductor diode lasers.
- Semiconductor diode laser** A laser in which recombination of current carriers at a p - n junction generates stimulated emission.
- Silica** Silicon dioxide, SiO_2 , the major constituent of ordinary glass.
- Silica glass** Glass in which the main constituent is silica.
- Single-frequency laser** A laser that emits only a very narrow range of wavelengths, nominally a single frequency but actually a very narrow range small enough to be considered a single frequency.
- Single mode** Containing only a single mode. Beware of ambiguities because of the difference between transverse and longitudinal modes. A laser operating in a single transverse mode often does not operate in a single longitudinal mode.

- Slope efficiency** The fraction of each additional watt of drive power turned into laser output above a threshold.
- Small-signal gain** Amount of amplification or gain at low power levels, below saturation.
- Solid-state laser** A laser made of a nonconductive solid that contains atoms that produce stimulated emission when excited by light from an external source. Different from semiconductor lasers.
- Speckle** Coherent noise produced by laser light. It gives a mottled appearance to holograms viewed in laser light.
- Spectral beam combination** Combining laser beams at closely spaced wavelengths to generate beams of higher power and better beam quality from fiber or diode lasers.
- Spectral reflection** Reflection from a mirror-like smooth surface.
- Spectroscope** An instrument that spreads out the spectrum of light from an external source.
- Spectroscopy** Study of the wavelengths emitted and absorbed by materials.
- Specular** Mirror-like reflection.
- Spherical aberration** The imperfect focusing of lenses with spherical surfaces that focus light from different parts of the surface to different focal points.
- Spontaneous emission** Emission of a photon without outside stimulation when an atom or molecule drops from a high-energy state to a lower one.
- Stable resonator** An optical resonator in which light rays bounce indefinitely between the two mirrors if there are no losses.
- Stimulated Brillouin scattering** The nonlinear effect with the lowest threshold in fiber lasers and amplifiers, arising from laser photons losing a small amount of their energy to vibrations of atoms in the glass.
- Stimulated emission** Emission of a photon that is stimulated by another photon of the same energy; the process that makes laser light.
- Supercontinuum** A broad range of wavelengths produced by nonlinear effects combine to spread the spectrum of a laser pulse with high peak power.
- Superluminescence** Amplified spontaneous emission in an LED; a lack of cavity mirrors prevents the feedback needed for laser action.
- Ternary** Compound made of three elements, for example, GaAlAs.

- Thermodynamic equilibrium** A nominally steady-state condition, when a system is thermally in balance.
- III–V semiconductor** A semiconductor compound made of one (or more) elements from the IIIA column of the periodic table (Al, Ga, and In) and one (or more) elements from the VA column (N, P, As, or Sb). Used in LEDs, diode lasers, and detectors.
- Three-wave mixing** A nonlinear process involving the mixing of three electromagnetic waves, with the sum of two of the frequencies equaling the third frequency. Used in optical parametric sources.
- Threshold** The excitation level at which laser emission starts.
- Threshold current** The minimum current needed to sustain laser action in a diode laser.
- Time response** The time it takes to react to a change in signal level.
- Total internal reflection** Total reflection of light back into a material when it strikes the interface with a material having lower refractive index at a glancing angle.
- Transition** Shift between energy levels.
- Transverse modes** Modes across the width of laser. Distinct from longitudinal modes, which are along the length.
- Transverse wave** A wave that oscillates perpendicular to the direction in which it travels, such as light and other electromagnetic waves.
- Tunable** Adjustable in wavelength.
- Ultraviolet** Part of the electromagnetic spectrum at wavelengths shorter than 400 nm to about 10 nm, invisible to the human eye.
- Unstable resonator** An optical resonator in which light rays reflected between the two cavity mirrors would leak out the sides of the cavity.
- Upper laser level** The upper energy level of a laser transition.
- Valence band** Energy levels of outer electrons in an atom, which form bonds to other atoms in a solid. Electrons are bound to the atom.
- VCSEL** Vertical-cavity surface-emitting laser, an important type of semiconductor laser.
- Vibronic** A electronic transition accompanied by a change in vibrational energy level.
- Virtual image** An image that can be seen by the eye but cannot be projected onto a surface.

- Visible light** Electromagnetic radiation visible to the human eye, at wavelengths of 400 to 700 nm.
- Wall-plug efficiency** The efficiency of a laser calculated as the fraction of input electrical power (i.e., coming through the wall plug) that emerges in the laser beam.
- Waveguide** A structure that guides electromagnetic waves along its length. An optical fiber is an optical waveguide.
- Wavelength** The distance an electromagnetic wave travels during one cycle of oscillation. Wavelengths of light usually are measured in nanometers (10^{-9} meter) or micrometers (10^{-6} meter). The standard symbol is λ (Greek lambda).
- Wavelength-division multiplexing** Transmitting signals simultaneously at many wavelengths through the same optical fiber. A similar technique called spectral beam combination can be used to generate high power beams from fiber or diode lasers.
- WDM** Wavelength-division multiplexing
- X-ray** A photon emitted when an electron drops into an inner shell of an atom. Typically, wavelengths are 0.03 to 10 nanometers, but the boundaries are not well defined.
- YAG** Yttrium aluminum garnet, a crystalline host for neodymium lasers.
- YLF** Yttrium lithium fluoride (YLiF_4), a solid-state laser host.

INDEX

- aberrations, optical 128–130
- absorption 36, 45–46
 - spectroscopy 460–461
 - surface 478–479
- acceleration, with laser 534
- achromatic lenses 130
- acousto-optic beam deflector 146
- acousto-optic modulators 159
- active layer 349
- active Q switches 187–188
- active species 258–261, 271–284
- adaptive optics 525
- additive manufacturing 489–491
- advantages
 - fiber laser 305
 - high-power lasers 476–477
 - laser light 432–433
- Airborne Laser 245–246
- Akasaki, Isamu A. 517
- alexandrite lasers 292–293
- Alferov, Zhores 361, 517
- all-fiber beam path 313
- Amano, Hiroshi 517
- ammonia maser 5
- amplification 4, 7, 67–72
- amplified spontaneous emission 205–206, 332
- amplifiers
 - fiber 175, 205, 304–305
 - fiber Raman 332–335
 - optical 202–203
- anti-lasers 530
- antireflective coatings 142
- Apollo laser ranging of the Moon 449
- argon-fluoride lasers 240–242
- argon ion lasers 229–232
- arsenide diode lasers 388
- articulated arms 146
- aspheric lenses 130
- asteroid defense 538
- astronomy 525–526
- atmospheric data links, laser 447
- atom lasers 522–524
- atomic clocks 519–520
- atomic transitions 34–36
- attenuation 45–46
- attosecond pulses 414, 532
- audio disc players 427–428
- avalanche photodiodes 160
- average power 11, 117–118
- bandgaps, 255, 350–351, 384
 - nitride lasers 385
 - semiconductor lasers 382–383
- bandwidth, 74–75, 98–99
 - diode laser 381
 - fiber amplifier 304–305
 - fiber laser 320
 - LEDs 355
 - solid-state lasers 288–293
 - ytterbium fiber lasers 323–324
- bar code scanners 434–436

- Barish, Barry 517
- bars, diode laser 369–370
- Basov, Nikolai, 516–517
- battle simulation, laser 457
- beam, laser 12–13, 73–80, 103–107
 - choppers 158–159
 - combination 318
 - deflectors 146
 - delivery 53, 147, 479–480
 - diameter 73, 104, 379–380
 - divergence 73, 103–106, 379–380
 - expander 128–129
 - hazards 548–552
 - modulation 156–159
 - parameter product 163
 - profile 107
 - propagation factor (M^2) 163
 - quality 107–108, 163–164, 305, 479–480
 - splitters 136
 - waist 105
- Bennett, William R. Jr. 217
- Betzig, Eric 517
- biomedical lasers 493–500
- birefringence 135
- Bloembergen, Nicolaas 517
- Blu-Ray 428, 438–439
- Bose–Einstein condensates 522–524, 526
- Bose, Satyendra Nath 522
- BPP, *see* beam parameter product
- brazing 486
- Brewster window 120–121
- broad-area diode laser 368–369
- broadening, line 100
- bulk solid-state lasers 273–277
- bundles, fiber optic 53
- buried heterojunction diode
 - laser 367–368
- cancer treatment 500
- Cao, Hui 530
- carbon dioxide lasers 103, 235–239, 488
- carbon monoxide lasers 246
- cataract surgery 496–497
- cavity
 - dumping 120, 181, 190–192
 - laser 81–85, 222–223
 - maser 66
 - optics 66–72, 312–313
 - resonance 66–72, 75
 - wavelength selection 100–103
- C band fiber amplifiers 329
- CD players 428
- ceramics 140, 259, 262
- chalcogenide materials, 139, 284, 328
- charge carriers in semiconductors 346–347
- chemical lasers 243–246, 508
- chemical oxygen-iodide lasers 245–246
- chirped pulse amplification 204–205, 323–324, 406, 531
- chromatic aberration 129–130
- chromium-doped chalcogenide lasers 284
- Chu, Steven 517
- cladding, *see* optical fibers
- classical optics 127–136
- cleaning with laser 487
- clocks, atomic 519–520
- coatings, 141–145
- Cohen-Tannoudji, Claude 517
- coherence 13–14, 95–98
 - and interference 97–98
 - length 97, 229
- collimators 128–129
- color filters 144–145
- colors, visible 25
- compositions, semiconductor lasers, table 382
- compound lenses 128
- compound semiconductors 342, 351–354
- computer printers 440–441
- concave lens 48
- concave mirror 50–51
- conduction 254–255, 344–347
- conduction bands 344–346
- confinement of light and current in semiconductor 89, 359–370
- construction alignment 448
- constructive interference 27–28
- continuous-wave emission 182
- conversion factors 553

- convex lens 48
- convex mirror 50–51
- cooling, laser 221, 521–522
- copper vapor lasers 234
- core, fiber 52, 302. *See also* optical fibers
- Cornell, Eric 517, 522
- corner cubes, *see* retroreflectors
- Corning Glass Works 429
- countermeasures, laser 458
- crystals 138–139, 259–260, 262
- current flow, diode laser 359–370
- curvature, lens 47, 50
- cutting with lasers 484
- cylindrical lenses 130–131

- data storage, optical 437–439
- data transmission, optical 442–447
- dazzlers, laser 458–459
- deep space communications 537
- defense lasers, low-power 456–459
- density, threshold current 359
- dentistry 498–499
- dermatology 497–498
- designators, target 456–457
- destructive interference 27–28
- detached retina 496
- detectors, light 160–161
- DFB laser, *see* distributed feedback
- D’Haenens, Irnee 6, 426
- diabetic retinopathy 496
- diameter, beam 73, 104–106
- dielectrics 254
- dielectric solids 174–175
- difference frequency generation 153
- diffraction 46
- diffraction gratings 133–134
- diffraction limit 103, 106–107
- diode lasers 175–176, 341–390
 - bars 369–370
 - compared to LEDs 353–359
 - confinement of light and current 359–370
 - history 342–343
 - structure 359–370
 - threshold current 358–359
- diode-pumped alkali lasers 247
- diode pumping 87, 257, 264, 267–271, 314, 320–322
- diodes, electrical 348–350
- direct bandgap 350–351
- direct diode lasers 481, 488–489
- direct modulation 157, 183, 381
- discharge pumping 179, 221–222, 414
- dispersion 130
- dispersive optics 132–134
- distributed Bragg reflection laser 373–374
- distributed feedback diode lasers 372–373
- divergence, beam 12–13, 73, 103–106
 - diode lasers 379–380
- dopants, laser 259
- doped fibers 302, 307–308, 311–312, 318–328
- doping, semiconductors 346–347
- Doppler effect 100, 223
- Doppler-free spectroscopy 518
- double-heterojunction diode laser 362–363
- drilling with lasers 483–484
- driverless cars 430–431
- dual-clad optical fibers 307, 314–315
- duration, pulse 117
- duty cycle 118, 184
- DVDs 428, 438–439
- dye lasers 399–404

- edge emitting diodes 366–375, 379–380
- effective mode area 323
- efficiency,
 - calculating 110–111
 - laser 81–82, 85–89, 110–115, 232, 272, 276–277, 316–317, 320–322, 381
 - quantum efficiency 112
 - wall-plug 85–86
 - waste heat, table 317
- Einstein, Albert 2, 5, 29, 37, 522, 524
- electrical excitation 87–88, 179–180
- electromagnetic radiation 4, 21–26
- electron accelerators 417
- electron-hole recombination 348–350
- electron shells 29–32

- electronic transitions 34–35, 255
- electrons, in semiconductors 346–350
- electro-optic devices 158
- emission, light 36–37, 353–359. *See also* stimulated emission, spontaneous emission
- energy extraction from laser cavity 113–114
- energy levels 29–36, 60–66, 255, 273, 278–279, 282, 292, 319, 400–401, 413
- energy, photon 23–24
- energy storage 180–181, 184–190
- equilibrium, thermodynamic 37–38, 60
- erbium-doped fiber amplifiers 328–332, 445–446
- erbium-doped fiber lasers 325–326
- erbium lasers 281–283
- etalon 101–102
- excimer lasers 173–174, 220, 240–242
- excitation, laser 61, 85–89, 177–181, 221–222
- excitons 349, 357
- exclusion principle 31
- external cavity diode laser 374–375
- external modulation 158–159
- Extreme Light Infrastructure 533–534
- extremely high power 530–534
- extreme ultraviolet photolithography 491–492
- extreme ultraviolet sources 410–416
- eye hazard to retina 426
- eye response to colors 455
- eye safety 162–163, 281–284, 426, 455, 477–478, 547–552
- eye-safe lasers 281–284
- eye surgery 495–497

- Fabry-Perot cavity 75–77
- Fabry-Perot diode lasers 371–372
- facet reflection, diode laser 357–358
- fake holograms 468
- fall time 116–117
- far-field beam 104
- far-infrared lasers 246
- feedback in diode laser 357–359
- femtosecond lasers, *see* ultrafast
- femtosecond spectroscopy, 519
- fiber amplifiers 304–305, 328–332
- fiber beam delivery 53, 147
- fiber Bragg gratings 313
- fiber lasers 175, 301–335
 - design 306–318
 - efficiency 316–317
 - history 303–305
 - materials working 488
 - Raman lasers 332–335
- fiber optics 51–54
 - bundles 53
 - communications 428–429, 442–447
 - flexible bundles 53
 - surgery 499
- fiber Raman amplifiers 332–335
- fibers
 - doped 302, 307–308, 311–312, 318–328
 - fluoride 328
 - microstructured 309–310
- filters, optical 141–145
- flashlamp pumping, solid-state lasers 263
- flowing gas laser 221
- fluorescence spectroscopy 459–460
- fluoride glasses 139, 328
- focal point 47–48
- fossil scanning 450
- four-level laser 62–64
- Fourier transform 118–119, 193, 409
- Franken, Peter 516
- free-electron lasers 177, 416–420
- frequency 22–24
 - combs 408–410, 519–520
 - doubling 516
- Fresnel reflection 312
- full width at half maximum 117
- fused silica 139
- fusion with lasers 503–505

- Gabor, Dennis 464, 517
- gain 82–83
 - bandwidth 74, 98–99
 - cavity 67–72
 - erbium amplifier 332
 - saturation 181, 311–312
 - small-signal 68

- gain-guided diode laser 367–368
- gallium arsenide 343, 351–352, 354, 357, 361–363, 366–368, 377, 387
- gas lasers 217–247
 - wavelength table 224
- gas replacement, in lasers 220–221
- gas-dynamic carbon-dioxide lasers 238
- Gaussian mode 78–79
- glass 139
- glass lasers 259, 262, 275–276
- Glauber, Roy 517
- Gordon, James 5
- Gould, Gordon 5, 7
- gravitational waves 524–525
- green laser dazzlers 458–459
- ground state 63, 112–113, 116
- group II and VI semiconductors 352–353
- group III and V semiconductors 351–353
- guide stars, laser 525–526
- gun sights, laser 456
- hair removal 498
- Hall, John 408, 517
- Hall, Robert 343
- Hänsch, Theodor 408, 517–518, 529
- harmonic generation 84–85, 151–152, 198–199, 210, 277–278, 286–287, 324, 326, 413–414, 516
- Haroche, Serge 517
- Harris, Steve 526
- Hau, Lene 526
- hazards
 - electrical 547–548
 - laser 162–163, 281–284, 426, 455, 477–478, 547–552
- heat dissipation, laser 262, 312
- heat treating 486–487
- helium–cadmium lasers 232–234
- helium–neon lasers 217, 225–229
- Hell, Stefan W. 517
- Herriott, Donald 217
- heterojunctions 361–363
- high harmonic generation 413–414, 532
- high power lasers
 - applications 475–511
 - chemical lasers 244–246
 - diode lasers 368–370
 - pulsed 204–205
- hollow optical waveguides 146–147
- holmium lasers 283, 327
- holographic interferometry 467
- holographic optics 146, 468
- holography 464–468
- Holonyak, Nick 343
- homojunctions 361, 363
- horizontal cavity surface emitting lasers 379
- host materials 254, 258–264, 280
- Hubble Space Telescope 525
- Huygens, Christian 21
- hybrid silicon photonics 389
- hydrogen atom 29–32
- hydrogen-fluoride chemical lasers 244–245
- idler beam 405–407
- illumination LEDs 355–356
- image, real 49
- impurities, semiconductor 346–347
- index-guided diode lasers 367–368
- indirect bandgap 350–351
- indium gallium arsenide diode lasers 387
- industrial laser market 481
- Industrial Laser Solutions* 481
- inertial confinement fusion 503–505
- infrared lasers, *see* carbon-dioxide, laser, semiconductor lasers, solid-state lasers
- infrared optics 238–239
- InGaAsP diode lasers 352, 372–373, 378, 382–383, 387–388
- injection lasers, *see* diode lasers
- inner cladding, optical fiber 307–308
- inner core, optical fiber 307–308
- insulators 254, 344–346
- integrated photonics 389, 536–537
- intensity 162
- interband cascade lasers 392–393
- interference 26–28
 - and coherence 97–98
- interferometry 451–452, 467
- Internet communications 428–429

- intracavity power 108
- intrinsic semiconductor 348
- inversion, population 37–38, 60–67
- ion lasers 172, 229–232
- ionization 31
- ionizing radiation 412
- IPG Photonics 322
- iron-doped lasers 284
- irradiance 162
- isotope separation 501–503

- Javan, Ali 217
- Jell-O laser 171, 529
- junction laser 342
- junction plane 366–370
- junctions, semiconductor 348–350

- Kao, Charles 429, 517
- Ketterle, Wolfgang 517
- kidney stone surgery 499
- Kimble, Jeff 528
- Kroemer, Herbert 361, 517
- Krupke, William 247
- krypton fluoride lasers 240–242
- krypton-ion lasers 229–232

- lamp pumping 256–257, 263–264, 266–268
- large-core fiber lasers 318
- laser applications 16–17
 - acceleration 534
 - acupuncture 500
 - communication through the air 447
 - cooling 521–522
 - countermeasures 458
 - cutting 484
 - dazzlers 458–459
 - fusion 503–505
 - guide stars 525
 - light shows 453
 - low power 425–469
 - high power 475–510
 - marking 484
 - measurement 447–452
 - paint removal 487
 - pointers 454, 550–551
 - printing 440–442
 - producing extreme ultraviolet plasmas 414
 - propulsion 537–538
 - radar 430–431
 - ranging 448–449
 - research 515–538
 - spectroscopy 516–520
 - weapons 505–510
 - welding 485–486
- laser components and properties
 - bandwidth 99
 - beam quality 107
 - beams 73–80, 103–107
 - cavities 66–72
 - diodes 341–390
 - energy storage 180–181
 - materials 8–9, 171–177
 - medium, shape 68–70
 - modulation 110
 - oscillation 7, 201–204
 - polarization 120–121
 - power 108–110
 - pulse control 119–120
 - rods 266–268
 - safety 162–163, 547–552
 - transition efficiency 112–113
 - types 169–170
- laserdisc players 427
- laser-like sources 209–211
- laser-pumped X-ray lasers 415
- lasers, listed by type e.g., gas lasers
- lasers and light bulbs 432–433
- LASIK surgery 495–496
- lattice constant 360, 383–384
- lattice matching 364–366
- Lawrence Livermore National Laboratory 504, 530–531
- L band fiber amplifiers 329–330
- lead salt diode lasers 389
- LEDs 343, 350, 353–359
- Leith, Emmett 465
- length, coherence 97, 229
- lenses 47–51, 127–131
 - compound 128
 - holographic 146, 468
- lidar 430–431, 448–449
- lifetime, excited state 37, 62–63
- light bulbs and lasers 432–433

- light detection 160–161
- light-emitting diodes, *see* LEDs
- light guiding 51–54, 301–303
- light shows, laser 453
- LIGO 524–525
- limits on pulse duration 118–119
- linear arrays, diode laser 369–370
- linear polarization 120–121
- line broadening 100
- line narrowing 99
- lines, spectral 32–36
- linewidth, laser 82–83
- living lasers 529
- longitudinal discharge, gas laser 221
- longitudinal flow carbon-dioxide lasers 237
- longitudinal modes 74–75, 101, 192–194
- Losev, Oleg 343
- loss in laser cavity 72
- lower laser level 60
- low-power laser applications 425–469
- machining 481–489
- Madey, John M. J. 417
- magnification 49
- Maiman, Theodore 6, 8, 426, 515
- marking, with lasers 484
- Mars Orbital Laser Altimeter 449
- masers 5, 61–62, 66
- master oscillator power amplifier (MOPA) 203–204
- materials 136–141, 171–177
 - absorption 478–479
 - coatings 142–145
 - diode lasers 381–389
 - dye laser 399–404
 - LED 354
 - microstructured 141, 309–310
 - nonlinear 152, 156
 - semiconductor laser 381–389
 - solid-state lasers 258–264, 271–284
- Matthews, Dennis 415
- measurements in optics 159–164
- measurement with lasers 447–452
- medical treatment, laser 493–500
- meniscus lens 48
- metal-vapor lasers 232–234
- metalworking with lasers 487
- metastable energy levels 62–63
- metric units 25, 555
- microstructured materials 140–141, 309–310
- mid-infrared lasers 283–284, 390–393
- mirrors 44–45, 50–51, 131–132
 - coatings 142
 - fiber Bragg gratings 313
 - front surface 44
 - laser cavity 67–72
 - output 68–70
- missile defense 505–510
- mixed-gas ion lasers 232
- mode hopping 372
- modelocking 120, 192–194
- modes
 - laser 73–75, 381
 - transverse 78–80, 107
- modulation 110, 156–159
- Moerner, William E. 517
- molecular gas lasers 173, 219–220, 235–239
- molecular transitions 35–36
- monochromator 134
- mounts, optical 148–149
- Mourou, Gerard 204, 517
- multikilowatt lasers 480
- multilayer interference coatings 143–144
- multimode fiber 53, 309
- multi-stage amplification 322–323
- multistriple diode lasers 369–370
- multi-transition lasers 207
- multiwavelength operation 102–103
- Nakamura, Shuji 517
- nanoscale lasers 527–528
- nanosecond pulses 184–190
- National Ignition Facility 504
- natural masers and lasers 65
- Navy LaWS (weapon system) 509–510
- near-field beam 104
- negative lenses 50
- negative temperature 39
- neodymium lasers 273–277
 - energy levels 31–33
 - glass lasers 275–276

- neodymium lasers (*Continued*)
 YAG 274
 YLF 275
- networks, optical 442–447
- neutral-density filters 145
- Newton, Isaac 1–2, 21
- nitride semiconductor lasers 384–386
- nitrogen lasers 243
- Nobel Prizes 516–517
- noise from amplified spontaneous emission 332
- noncontact machining 482–483
- nonlinear effects 317, 323
- nonlinear optics 84–85, 149–156, 198–199, 404–408
- normal (angle) 42
- n*-type semiconductor 346, 348–349
- nuclear fusion with lasers 503–505
- nuclear X-ray lasers 415–416
- OPA, 284–288, *see also* optical parametric amplifiers
- opacity 39–40
- OPO, *see* optical parametric oscillators
- OPSLs, *see* optically pumped semiconductor lasers
- optical aberrations 128–129
- optical amplifiers 202–204, 328–329, 445–446
- optical coatings, *see* coatings
- optical coherence tomography 464
- optical communications 442–447
- optical damage 318
- optical data storage 437–439
- optical density 145
- optical disc players 427–428
- optical fibers 51–54, 301–310
- optical materials 136–141
- optical measurements 159–164
- optical modulators 158–159
- optical molasses 522
- optical to optical conversion 86–87, 115, 316
- optical parametric amplifiers 405–406
- optical parametric oscillators 406–407
- optical properties of diode lasers 379–381
- optical pumping 5, 8, 86–87, 178, 197–198, 256–257
 semiconductor lasers 284–288
- optical SETI 537
- optical signal amplifiers 205
- optical tables 148–149
- optical trapping 521
- optical tweezers 521
- optical waveguides 146–147
- optics 1–2, 39–40, 47–51
 for carbon-dioxide lasers 238–239
 laser cavity 83–84
- orbits, electron 30
- organic dye lasers 176–177, 399–404
- organic semiconductors 353
- oscillation, fiber Raman laser 334–335
- oscillator-amplifiers 202–204, 272, 276, 328–329, 445–446
- oscillators, laser 7, 70–72, 201–204
- outer core, optical fiber 307–308
- output mirror 68, 70, 106
- output optics, laser 106
- overall efficiency 110–111
- oxide glasses 139
- pain therapy 500
- paint removal, with laser 487
- parametric gain 153–154
- parametric sources, *see* optical parametric
- passive Q switches 188–189
- Pauli exclusion principle 31
- Payne, David 304
- peak power 11, 116–117
- penetration depth, laser drilling 483–484
- petawatt laser 530–531
- Phillips, William 517
- photochemistry 501–503
- photodetectors 160–161. *See also* detectors
- photodiodes 160
- photodynamic therapy 500
- photoelectric effect 28–29
- photolithography 360, 491–492
- photomasks 491
- photometry 161–162
- photon defect 112

- photon energy 23–24
- photonic bandgap fiber 309–310
- photonic crystal fibers 309–310, 527
- photonic integration 536–537
- photonics 1–2
- photons 2, 21–24, 28–29
- planar optical waveguides 147
- Planck, Max 23
- Planck's constant 23
- pointer hazards 550–551
- pointers, laser 454–455
- polarization 120–121, 134–135
- population inversion 37–38, 60–67
- positioning equipment 148–149
- positive lenses 47–49
- power beaming 537
- power in the bucket 164
- power density 109
- power, intracavity 108
- power supply efficiency 111–112
- prefixes, metric 25
- prepress equipment 442
- pressure broadening 100
- printers, computer 440–441
- prisms 133–134
- profile, beam 107
- Prokhorov, Alexander 516–517
- p*-type semiconductor 346, 348–349
- pulse characteristics, laser 115–120
 - cavity dumping 190–191
 - Q switching, 184–190
 - timing 115–120
- pulsed lasers 12, 182–194, 294–295, 323–325
- pumping 177–181
 - dye lasers 403–404
 - efficiency 85–86, 88
 - fiber lasers 303–304
 - laser rods 266–268
 - neodymium 278–279
 - optical 5, 8, 86–87, 178, 197–198, 256–257
 - parametric sources 405–407
 - semiconductor lasers 284–288
 - slab lasers 268–269
 - solid-state lasers 262–264
 - thin-disk lasers 269–271
- Q switching 119, 181, 184–190, 294–295
- quality factor 184
- quanta 23, 29
- quantum cascade lasers 390–393
- quantum computing 536
- quantum cryptography 536
- quantum defect 315–316, 321
- quantum dots 364–365
- quantum efficiency 112
- quantum mechanics 29–32
- quantum wells 363–365
- quantum wires 364–365
- quasi-continuous-wave 183–184
- quasi three level laser 320
- quaternary semiconductors 352

- radar, laser 430–431, 448–449
- radiation, electromagnetic 4, 21–26
- radiation, ionizing 412
- radiometry 161–162
- radius of curvature 47, 50
- Raman amplification 332–334
- Raman fiber lasers 332–335
- Raman scattering 154–156
- Raman shift 154–156, 200, 210, 324
- rangefinders, laser 449
- rare-earth doped fibers 318–332
- rare-gas-halide lasers 173–174, 240–242
- Rayleigh range 104
- Rayleigh scattering 46
- reading with lasers 433–436
- Reagan, Ronald 505
- real image 49
- recombination, electron-hole 348–350, 354
- red diode lasers 386
- Rediker, Robert 343
- reflection 39–40, 44–45, 140–141
- reflection, total internal 43, 51–54
- reflectivity 39–40
- refraction 41–43, 45, 47
- refractive index 40–41
 - and wavelength 45
 - differences, fiber cores 307–309
 - semiconductor 358
- refractive surgery 495–496

- remote sensing 461–462
- repetition rate (pulse) 117
- replacement of laser gas 220–221
- research with lasers 515–538
- research spectroscopy 463–464
- resonators 66–80, 303
- retroreflectors 131–132
- reverse-bias 348–349
- ring resonator 78–80
- rise time 116–117
- rods, laser 68–70, 266–268
- rotational transitions 35–36
- Round, Henry J. 343, 353
- ruby lasers 6, 8, 63, 256, 271–272

- safety requirements for lasers 477–478, 547–552
- sampled grating distributed Bragg reflector laser 374–375
- saturation 181, 188–189, 311–312
- scanners, bar code 434–436
- scanners, beam 146
- scattering, Rayleigh 46
- Schawlow, Arthur 5, 7, 517–518, 529
- sealed-tube carbon-dioxide lasers 236–237
- second harmonic generation, *see* harmonic generation
- selective laser sintering 490–491
- selenium energy levels 413
- self-terminating lasers 194–195
- semiconductor fabrication 360, 491–493
- semiconductor lasers 175–176, 341–393
- semiconductor optical amplifiers 375
- semiconductors 254–255, 344–345
 - compound 351–353
- sensing, laser 459–464
- side pumping, fiber laser 314
- signal amplifiers, fiber 304–305, 331–332
- signal beam 405–407
- silica 46
- silicon 346–347
- silicon carbide 352
- silicon photonics 389
- simple lenses 47–51
- single-atom lasers 527–528
- single-clad optical fibers 307
- single heterojunction diode laser 361–363
- single-mode fiber 53, 309
- single-wavelength operation 102–103
- slab lasers 268–269
- slope efficiency 113
- slow light 526
- small core fiber lasers 317
- small-signal gain 68
- Snell's law 42–43
- Snitzer, Elias 301
- soft X-rays 412
- soldering 486
- solid-state lasers 8–9, 171, 174–175, 253–295, 488
 - defined 253–258
 - materials listed 261
- solution looking for a problem, 426
- space laser communications 447
- spasers 528
- spatial coherence 97
- spatial filters 145
- speckle 98
- spectral beam combination 318
- spectral lines 32–36
- spectroscopy 459–462, 516–520
- spectrum, electromagnetic 24–26
- specular reflection 44
- speed of light 22–23, 40–41
- spherical aberration 130
- spontaneous emission 36–37, 68–69, 354–357
- spot size 13, 128
- stable resonators 76–77
- stacks, diode laser 369–370
- stereolithography 489
- stimulated Brillouin scattering 317
- stimulated emission 4–5, 9, 36–37, 68–71
- Stone, Douglas 530
- strained layer superlattice 366
- Strickland, Donna 204, 517
- stripe-geometry diode laser 366–367
- sum-frequency generation 153
- supercontinuum sources 210–211, 408
- supermarket checkout 434–435

- superradiance 206, 211
- surface absorption, table 478
- surface emitting diode lasers 375–379
- surface plasmons 528
- surface profiling 450
- surgical lasers 493–495
- surveying 448

- table, solid-state laser materials 261
- target designators, laser 456–457
- tattoo bleaching 497
- telecommunications 428–429
- Teller, Edward 415
- TEM modes 78–80
- temporal coherence 97
- ternary semiconductors 352
- thermodynamic equilibrium 37–38, 60
- thin disk lasers 269–271, 284–288
- Thorne, Kip 517
- II–VI compounds 351–353, 388
- III–V compounds 351–353, 377, 382–384, 388
- three-dimensional printing 489–491
- three-dimensional scanning 450
- three-level laser 62–63
- threshold current 347, 351, 357–359, 377, 381
- threshold, laser 113–114, 264, 311–312, 317, 357–359, 377, 381
- thulium lasers 282, 326
- timing, pulse 115–120
- tissue interactions 493–495
- titanium-sapphire lasers 290–291, 414
- total internal reflection 43, 51–54, 302, 307
- Townes, Charles H. 5, 7, 516–517
- transform limit 118–119, 193
- transitions
 - carbon dioxide 235–236
 - energy level 32–36
 - excimer lasers 241
 - extreme ultraviolet 412
 - helium-cadmium 233
 - helium-neon 226
 - ion laser 230–231
 - ruby lasers 272
 - solid-state lasers 195, 260–261
 - vibronic 288
 - wavelength 81–82
 - ytterbium 278–279, 321
- translucent 43
- transmission, light 39–40
- transmission, optical materials 138
- transparency 39, 43, 138–140, 261–262
- transverse discharge, gas laser 221–222
- transverse flow carbon-dioxide lasers 237–238
- transversely excited carbon-dioxide lasers 238
- transverse modes 78–80
- transverse waves 21–22
- tube, helium-neon laser 227–228
- tubes for gas lasers 222–223
- tunable lasers 81–85, 207–209
 - diode lasers 374–375
 - dye lasers 399–404
 - free-electron 418–420
 - semiconductor lasers 390–393
 - solid-state lasers 288–293
 - titanium sapphire 290–291
 - tweezers, optical 521
 - ytterbium 279–280
- two-dimensional bar codes 435–436
- two-level laser 61–62
- types of lasers 169–170

- ultrafast lasers 192–194, 204–205, 209, 290–291, 324–325, 408–410, 420, 530–535
- ultrafast laser surgery 499
- ultrafast machining 481
- ultrafast spectroscopy 519
- ultraviolet bands, defined 411
- ultraviolet lasers, *see* excimer lasers, nitrogen lasers, harmonics, ion lasers
- ultraviolet sources, extreme 410–416
- uncertainty principle 118–119
- undulator magnets 417
- Universal Product Code 434–435
- unstable resonators 76–77, 80
- Upatnieks, Juris 465
- upper laser level 60, 62
- uranium enrichment 501–503

- vacuum ultraviolet 411
- valence bands 344–346
- VCSELs 376–378, 380
- VECSELs 284–288, 378–379
- vertical cavity surface emitting
 - lasers 376–378, 380
- vertical external cavity surface
 - emitting lasers, *see* VECSELs
- vibration measurements 467
- vibrational transitions 35–36, 236
- vibronic lasers 288–290, 292–293
- video disc players 427–428
- violet laser pointers 454
- visual range, human 25

- wall-plug efficiency 85–86, 114–115, 305
- war games 457
- waste heat 85–86, 114–115, 117
- waveguides 52
- wave interference 26–28
- wavelength-division multiplexing 329, 331, 444–445
- wavelengths 10–11, 22–24, 45, 98–103
 - conversion 84–85, 178, 195–200
 - free-electron lasers 419–420
 - gas lasers 224 (table)
 - OPSLs 286–287
 - optical parametric amplifier 406
 - optical parametric oscillator 407
 - quantum cascade lasers 392
 - and refractive index 45
 - selection 81–85, 100–103, 207–209
 - semiconductor lasers, tables 382, 385
 - ultraviolet 411
 - vibronic lasers 289
 - ytterbium 280
- wave-particle duality 2
- wave plates 135
- waves, electromagnetic 21–26
- waves, transverse 21–22
- weapon-class lasers 508–509
- Weiss, Rainer 517
- welding, laser 485–486
- white-light LEDs 355–356
- Wieman, Carl 517, 522
- wiggler magnets 417
- Wineland, David J. 517
- wound healing 500

- X-ray free-electron lasers 420, 535
- X-ray lasers, pumped by laser
 - 415
- X-ray lasers, pumped by nuclear
 - bomb 415–416

- Young, Thomas 21
- ytterbium lasers
 - broadband 293–294
 - doped fiber amplifiers 330
 - doped fiber lasers 320–325
 - doped lasers 278–280
 - energy levels 321
- yttrium–aluminum garnet, as Nd
 - host 274
- yttrium lithium fluoride, as Nd
 - host 275

- Zewail, Ahmed 517, 519