FIGURE 19.13

High-speed signal assembled from data packets.

Incoming Data Packets

Headers

Loaded onto high-speed signal as fast as they come.

Empty Interval

Data Transmission
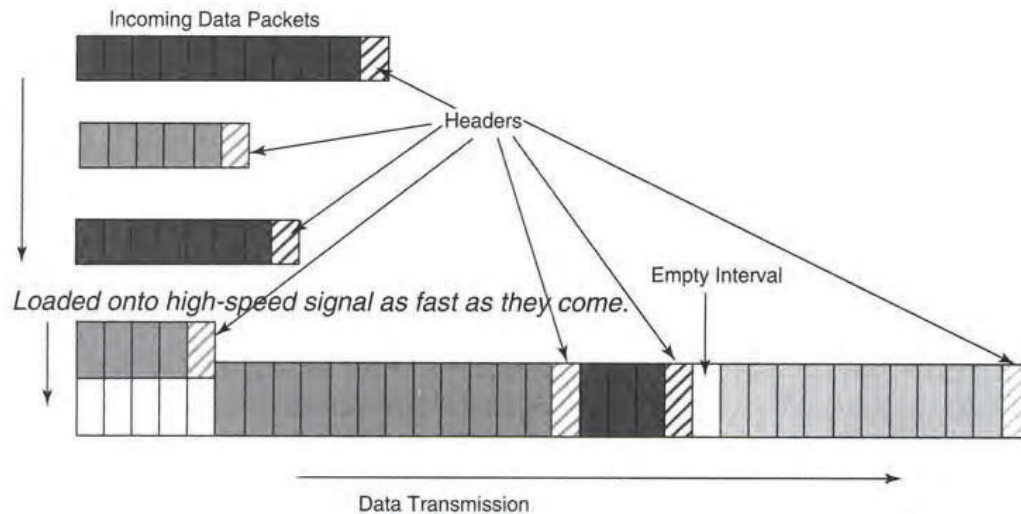
directly from data bundled in packets, as shown in Figure 19.13. The packets do not have to arrive at a particular rate or contain a particular number of bits, although they must have headers that specify their destination. Figure 19.7 compares this to TDM circuit switching.

The data rate is set by the clock speed of the transmitter, which generates a specified number of pulses per second. Data packets are lined up for transmission in the sequence of their arrival. The transmitter sends one packet at a time, starting with the header, which contains address information, and proceeds through the rest of the bits in sequence. If another packet is waiting, the transmitter starts sending it; otherwise it sends blank intervals.

It's like lining people up to get onto a moving staircase where each step holds one person. The people line up at the base, and one gets onto each passing step. As long as people are in line, every step is filled. Steps only go empty if no one is left in line. To make the analogy better, you could imagine tour groups lining up, with a single leader acting as the "header" of each group.

● Statistical multiplexing combines multiple packet streams that may have different data rates.

Multiple packet streams are combined by a technique called *statistical multiplexing*, rather than by interleaving bits. Incoming data packets on each channel accumulate in separate storage buffers. The multiplexer takes packets from each buffer in turn, keeping track of the traffic on each channel, then allocates more time to the busiest channels.

Statistical multiplexing is used for transmitting Internet Protocol (IP) traffic. Like time-division multiplexing, it combines data streams from multiple sources into a single faster signal. However, it does not require a rigid hierarchy of incoming data rates. Because statistical multiplexers average bursty traffic over many channels, the total capacity of the input channels may be higher than their output capacity. That is, a statistical multiplexer might have 10 inputs delivering up to 100 Mbit/s, but one output able to transmit only 600 Mbit/s. That design can work as long as the average input is below the peak capacity. For example, if each input channel averages only 20 Mbit/s, the combined inputs should be safely below 600 Mbit/s most of the time. (This sort of averaging is common in telephone networks, which don't have output connections for every input line because most lines are used only a small fraction of the time. Problems only arise when everyone makes long-distance calls on Mother's Day.)

Statistical multiplexing requires a set of priorities because too many packets may queue up while waiting to be transmitted. One approach is to allow packets to "time out" after a certain interval, like a person who gives up after waiting on hold for customer service. This method is used in the traditional version of the Internet Protocol (IPv4). A refinement is to set higher priority for delay-sensitive traffic, such as telephone conversations or streaming video, so it goes to the head of the transmission queue, while lower-priority data packets wait; this is used in IPv6, which is not yet implemented across the entire Internet.

Wavelength-Division Multiplexing

Wavelength-division multiplexing (WDM) is the transmission of different signals on multiple wavelengths through the same fiber. It closely resembles electronic frequency-division multiplexing, but is done at the much higher frequencies of light waves, as shown in Figure 19.14. The optical channels at 193.1, 193.2, 193.3, and 193.4 THz are the optical counterparts of the radio-frequency channels at 174, 180, 186, and 192 MHz shown in Figure 19.13.

Channel spacing may be defined in terms of frequency or wavelength, but if the carriers are in the optical or infrared part of the spectrum, the process is generally called wavelength-division multiplexing. Standards have been set for two types: *dense-WDM (DWDM)* and *coarse-WDM (CWDM)*.

Dense-WDM packs signals closely together, usually for long-distance transmission, and usually in the *erbium-amplifier band* between about 1530 and 1610 nm. Standards specify DWDM spacings in frequency units as 50, 100, and 200 GHz, which correspond to about 0.4, 0.8, and 1.6 nm in wavelength units near 1550 nm. The same standards also allow for spacing at 400 and 1000 GHz, which usually are not considered DWDM. The packing density of DWDM systems is limited by the modulation bandwidth of the signals, so 2.5-Gbit/s signals can be packed more tightly together than can 10-Gbit/s signals.

●
WDM transmits signals at multiple wavelengths through one fiber.

FIGURE 19.14
Wavelength-division multiplexing.

Coarse-WDM is used in metro networks, which don't require amplification.

Coarse-WDM spaces channels more loosely to reduce equipment costs. Current specifications space 18 channels at 20-nm intervals between 1270 and 1610 nm. The primary applications of coarse-WDM are in metro networks, which can carry heavy traffic but don't require optical amplification.

Total fiber capacity is limited by the range over which WDM is possible. In practice, amplifier bandwidth limits the range in amplified systems. The range is larger in unamplified systems, where the limits come from fiber attenuation. Potential capacities can reach staggering levels. The entire erbium band can accommodate more than 100 channels at 10 Gbit/s, allowing commercial systems to offer total capacity of more than 1 Tbit/s. Experimental systems have transmitted more than 10 Tbit/s a few hundred kilometers through a single fiber.

Most WDM systems have few populated channels.

Today's DWDM systems generally populate only a few optical channels because actual transmission requirements fall far short of total fiber capacity. However, carriers reserve the extra channel capacity for inexpensive future expansion.

Signal Management and Optical Networking

Multiplexing is involved in signal management.

Multiplexing is closely related to the management of signals transmitted through a telecommunications network. Multiplexing combines signals so they can be transmitted through fiber or other media. Multiplexed signals may be demultiplexed or otherwise broken up when they have to be switched, redirected, or otherwise managed. Electronically multiplexed signals are broken up into their component bit streams (or analog channels). WDM signals are broken up into their component wavelengths.

Much publicized during the telecommunications bubble, optical networking is signal management in the optical domain based on wavelength. An important advantage of optical networking is that WDM channels are entirely independent of each other, unlike TDM signals that must be in the same family of formats. A single fiber can carry several wavelengths, each transmitting in different formats. One wavelength can carry Gigabit Ethernet, another a 2.5-Gbit/s stream of Internet data, a third a single digitized high-definition video signal, and a fourth an analog cable-television signal. All you need are transmitters, receivers, fibers, optical amplifiers, and WDM optics that can handle the required data rates.

Signal management in the optical domain also requires *optical switches* that can select one or more optical channels in a WDM system. Ideally, optical networks also should be able to convert signals from one wavelength to another for transmission in other parts of the network.

Granularity measures how finely a signal can be broken up into its component parts.

The ability to switch individual optical channels can aid signal management by easing access to individual data streams. Suppose 40 streams of 2.5-Gbit/s data are multiplexed in a single fiber. Optical networking could easily isolate one data stream if the signals were 40 WDM optical channels. Isolating one data stream from a 100-Gbit/s TDM signal is much harder. The ease of breaking up multiplexed signals into their components is called *granularity*. The more granular the signal, the better carriers can meet customer needs and the more efficiently they can manage their networks.

Optical networking is still evolving, and it is too early to predict its final shape. Issues to be resolved include the technologies for wavelength conversion and other functions, what services should be offered, and what customers require.

Transmission Distance

The telecommunications network is global in extent, but transmission distance remains an important element of system performance. Two types of transmission span are important in fiber optics—the distance between individual optical amplifiers, and the distance between regenerators or repeaters, which are separated by a series of amplifiers.

The span between a pair of amplifiers is limited by the input power and the fiber loss. When signal power diminishes to an unacceptable level, you need an amplifier. Amplifiers aren't needed in short systems, where the required power reaches the receiver without amplification.

The end-to-end distance is limited by the degradation the signal experiences as it passes through a series of amplifiers. Once noise and dispersion have degraded the signal sufficiently, it must be regenerated. Typically regenerators are built into terminal nodes of a system, so the space between regeneration points is the length of a system, such as that of a submarine cable link from New York to London, or a terrestrial cable from Chicago to Saint Louis.

As you will learn, trade-offs are involved in the two distances. Increasing the spacing between amplifiers also increases the noise added to the signal and reduces the end-to-end distance. That's why repeaters have to be spaced much closer together in cables that cross the Atlantic ocean than in cables that run a few hundred kilometers between large cities.

● Signal degradation limits end-to-end distance.

Cost and Reliability

Network cost and reliability are critical concerns for carriers that provide telecommunication services. A cutting-edge system that sets transmission speed records is fine in the laboratory, but no carrier will want to install it if it is out of service half the time or requires a small army of top-level engineers to maintain.
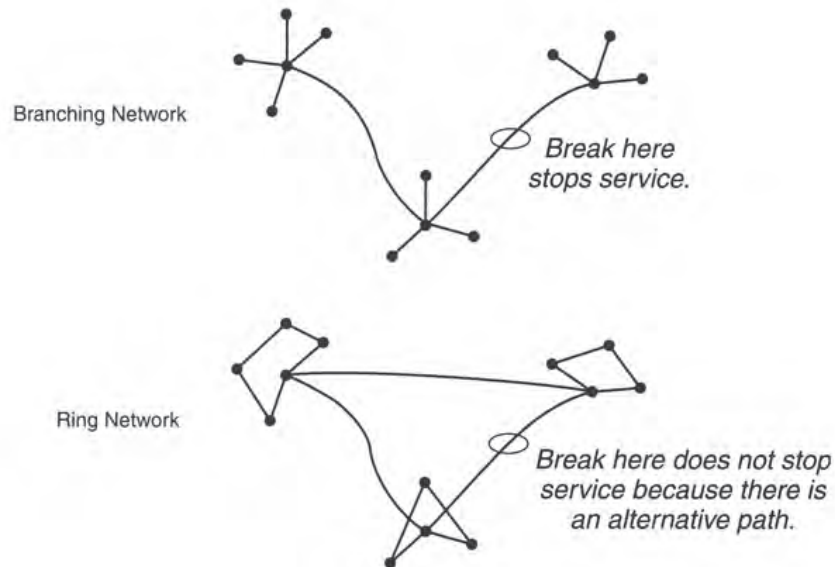
Costs fall into two broad categories: *capital expense (capex)* and *operating expense (opex)*. Capital expense is the cost of buying new equipment. Operating expense is the cost of running the network, including power, normal maintenance, required adjustments, and repairs. Often they are charged to separate departments, but choices are made based on trade-offs between the two. For example, an optical add-drop multiplexer that can be reconfigured from a central facility will incur more capital expense than one that requires manual adjustment. But the manually operated model will have higher operating expense because technicians have to travel to the site to make adjustments. Engineers decide between the two types of systems based on their expectations of operating requirements. Typically spending an extra $50,000 on the latest, greatest, automated system isn't reasonable if the system is only 50 miles from the control center and will require only one adjustment every year. But it's another matter if the site is 500 miles from the control center and requires adjustment twice a week.

Construction and installation costs typically count as capital expense and are another important factor. Normally carriers plan capacity for future expansion, so their network designs include spare fibers and extra wavelength slots in fibers to provide that capacity. Adding an extra pair of fibers to a cable being installed now is much cheaper than installing an additional cable later.

● Costs are divided into capital expense and operating expense.

FIGURE 19.15

Branch and ring topologies for connecting network nodes.

Branching Network

Break here stops service.

Ring Network

Break here does not stop service because there is an alternative path.

Ring and mesh networks can better withstand cable failures than branching networks.

Reliability depends on factors such as network topology and operating environment as well as the choice of equipment. Traditional networks branched out from a few central points, as shown in Figure 19.15. This reduced cable and installation costs, but left the network vulnerable to interruption by a single cable break. Modern systems that carry heavy traffic are arranged in rings or meshes, so an alternative path remains between any two points in the network if a cable breaks.

The operating environment also is important for reliability. Much equipment is outdoors, where it must withstand moisture and extreme temperatures. Remote optics and electronics are installed in shielded cases, or "repeater huts," which provide some protection. Sensitive equipment usually is installed in temperature-controlled office-like environments.

What Have You Learned?

1. The global telecommunications network is evolving rapidly in response to changes in business and technology.

2. The network includes copper, wireless (radio), and fiber links.

3. Optical networking manages light signals, usually in the heart of the network.

4. Telecommunications networks are made of links between nodes. They carry voice, video, and data signals, which have different transmission requirements.

5. The telephone system is circuit-switched. The Internet is packet-switched.

6. Standards or protocols allow different equipment to exchange signals in a common format.

7. Point-to-point transmission links pairs of terminals permanently. Point-to-point multipoint or broadcast transmission sends signals from one transmitter to many terminals.

8. Networks link many terminals with each other in various configurations.

9. Packet switching uses routers to deliver signals in a network; the routers read data headers on the packets to determine their destination. The Internet uses packet switching.

10. Circuit switching establishes temporary connections between pairs of terminals; it is used for telephone systems.

11. A signal modulates a carrier wave at a higher frequency. It may modulate amplitude, frequency, or phase of the carrier. Amplitude modulation is most common in fiber optics.

12. No return to zero (NRZ) coding is the most common way to code digital signals for fiber-optic transmission. Signals in NRZ form do not return to zero after a 1 bit is transmitted.

13. Error-correction codes operate on blocks of data, correcting errors that arise in transmission. The number of corrections depends on the coding scheme.

14. Transmission capacity is the bandwidth or data rate of a communication system.

15. Time-division multiplexing interleaves data in several incoming bit streams to generate a composite bit stream containing all the data. All input is at the same rate, and the output rate is a multiple of the input rate.

16. Statistical multiplexing combines multiple streams of data packets, which may arrive at different rates, to generate a single stream of data packets. It is used on the Internet.

17. Wavelength-division multiplexing and frequency-division multiplexing combine signals at different wavelengths or frequencies. WDM applies to optical transmission through a fiber; FDM applies to radio or microwave transmission through air or coaxial cable.

18. Most WDM systems have few populated channels.

19. Granularity measures how finely signals can be broken up into their component parts. It is important for signal management.

20. Transmission distance may be measured as spacing between amplifiers or regenerators. Amplifier spacing depends on power and attenuation; regenerator spacing depends on noise and dispersion.

21. Costs are divided into capital expense and operating expense.

What's Next?

In Chapter 20, you will learn about the standards developed for fiber-optic systems.

Further Reading

Roger L. Freeman, *Fundamentals of Telecommunications* (Wiley InterScience, 1999)

Gary M. Miller, *Modern Electronic Communications,* 6th ed. (Prentice Hall, 1999)

Jean Walrand and Pravin Varaiya, *High-Performance Computer Networks,* 2nd ed. (Morgan Kaufmann, 2000)

Questions to Think About

1. Internet routers read headers on data packets, and use that information to direct the packets toward the proper destination. What type of network architecture would you expect to be used to connect routers in the Internet backbone? Why?

2. Can you use headers to control the flow of data packets around a ring network in which signals pass through all the terminals anyway?

3. An electronic time-division multiplexer generates signals at 1 Gbit/s. If it has eight inputs, what are their data rates?

4. How would an optical time-division multiplexer work?

5. How is wavelength-division multiplexing analogous to frequency-division multiplexing?

6. What are the prime limitations on amplifier spacing and regenerator spacing? How do the two affect each other?

Chapter Quiz

1. What type of telecommunications is packet-switched?

    a. cable television

    b. standard telephones

    c. the Internet

    d. all

    e. none

2. What type of telecommunications is circuit-switched?

    a. cable television

    b. standard telephones

    c. the Internet

    d. all

    e. none

3. What protocol covers the arrangement of data for packet switching?

    a. the Internet Protocol

    b. SONET

    c. Asynchronous Transfer Mode

    d. Plesiochronous Digital Hierarchy

    e. NRZ coding

4. A router does which of the following?

    a. It makes circuit-switched connections between terminals.

    b. It broadcasts signals to many points.

    c. It optically directs light signals to their destinations.

 d. It reads packet headers and directs signals to their destinations.

 e. It is equivalent to a switch.

5. A switch does which of the following?

 a. It makes circuit-switched connections among terminals.

 b. It broadcasts signals to many points.

 c. It converts packet headers to circuit-switching directions.

 d. It reads packet headers and directs signals to their destinations.

 e. It is equivalent to a router.

6. Amplitude modulation is used for

 a. digital transmission of a single 2.5-Gbit/s optical channel over fiber.

 b. analog transmission of cable-television signals.

 c. AM radio broadcasting.

 d. optical transmission of Internet data.

 e. all of the above

7. What is the proper name for digital coding in which a strong signal means a 1 and a low or zero signal means a 0?

 a. no return to zero (NRZ)

 b. return to zero (RZ)

 c. Manchester coding

 d. frequency-division multiplexing

 e. phase modulation

8. Interleaving incoming bit streams to produce a faster output signal is called

 a. packet switching.

 b. frequency-division multiplexing.

 c. time-division multiplexing.

 d. statistical multiplexing.

 e. wavelength-division multiplexing.

9. Simultaneously transmitting separate signals through an optical fiber at different wavelengths is called

 a. packet switching.

 b. frequency-division multiplexing.

 c. time-division multiplexing.

 d. statistical multiplexing.

 e. wavelength-division multiplexing.

10. What type of multiplexing requires all incoming signals to be at the same data rate?

 a. packet switching

 b. frequency-division multiplexing

 c. time-division multiplexing

d. statistical multiplexing

e. wavelength-division multiplexing

11. Transmission capacity of an optical fiber is the

a. total amount of information the fiber can transmit.

b. distance between amplifiers.

c. number of wavelengths the fiber can transmit, regardless of data rate.

d. distance from end to end.

e. data rate that can be transmitted at 1550 nm.

12. A fiber-optic system can transmit 2.5 Gbit/s on each of 40 optical channels, with an amplifier spacing of 100 km. The company that operates the system has installed transmitters and receivers at only 4 wavelengths. What is the data rate of the installed system?

a. 2.5 Gbit/s

b. 10 Gbit/s

c. 40 Gbit/s

d. 100 Gbit/s

Fiber System Standards

About This Chapter

For two people to communicate, they must speak the same language. Communication systems likewise work only if transmitters and receivers attached to them speak the same language. Communication engineers have devised standards to assure that equipment from different companies will be able to interface properly.

This chapter will introduce you to the system-level standards most important for fiber-optic systems. Some are specific to fiber optics; others also cover other communications technologies. The topic of standards is complex and continuously evolving, so I will not go into much depth, especially for standards with little direct impact on fiber-optic systems. However, you should at least learn to recognize the most important standards and their functions.

Why Standards Are Needed

As you learned earlier, signals can be transmitted in a variety of ways, with different types of digital or analog coding. However, those differences only scratch the surface of the immense potential for variations. You can think of those physical differences as being similar to the distinctions among the media you use to communicate with other people—speech, the written word, sign language, pictures, and so on.

There are many other levels of variations in signal formats, just as people speak many different languages or computer programs store data in different formats. Unlike human languages, signal formats are designed by engineers to transmit signals efficiently and economically. Their choices depend on the types of signals being carried, the distances

and types of terminals involved, and the hardware and software they have available. The results can vary widely with factors such as time and network scale.

These differences become a problem when you want networks to connect to each other or when you want to combine two or more generations of equipment, such as existing telephones with new digital switches and transmission lines. Then you need common languages and ways of translating signals between formats. That's when you need standards.

Standards have evolved considerably over the years, changing with both the marketplace and the technology. In the 1970s, AT&T was effectively America's telephone monopoly, so it set the standards for telecommunication systems. Since the 1984 breakup of AT&T, industry groups have come to set the standards. Many standards have become international, so you can make phone calls to Brazil, send faxes to India, and dispatch e-mail to Indonesia.

Changing standards have accommodated changes in industry practice. In the 1960s, telephone lines carried only analog voice telephone conversations. By the 1970s, the telephone network started to convert to digital voice transmission between switching centers. In the 1980s, the telephone network started to handle more computer data and video transmission, plus fax signals. Today, the high-speed lines operated by long-distance carriers are digital data highways that transmit a wide variety of signals, all digitized into a common form that can be reconverted to other formats at the receiving end. Standards make this multipurpose system possible.

●
Standards serve diverse functions.

Standards serve a variety of functions. They establish common physical features as well as common languages, so devices can fit together as well as understand each other. They may also assure that equipment meets the requirements of specific customers such as telephone companies or the military.

This chapter covers standards that apply to fiber-optic systems. Some important examples include:

- Physical standards for connectors, to make sure they mate optically and mechanically.
- Optical, electronic, and mechanical interface standards for packaged devices, so they connect properly to each other.
- Standard formats for data transmission.
- Telephone industry standards for data transmission and device performance.

Other standards also affect fiber-optic equipment. Fire-safety standards affect the choice of materials for cables. Component standards assure that connectors, transmitters, and laser modules are interchangeable. We won't talk about these standards here because they affect system design only indirectly.

Standards may be industry-wide, proprietary, or some combination of the two. For example, industry committees set standard data formats for Ethernet transmission, but software usually stores data it generates in a proprietary format. In some cases proprietary formats become de facto standards, like Microsoft Word format for word processing files. Usually proprietary standards are optimized for a particular company's equipment, while industry-wide standards are a compromise that works reasonably well on everyone's equipment.

Companies often battle over standards as they try to gain a competitive edge. Sometimes two or more groups create competing standards. However, agreed-upon standards are crucial to the function of an open, deregulated market involving many vendors and equipment that must interconnect.

You need to be sure you can plug any phone you buy today into the telephone jack in your wall and use it with any local or long-distance carrier. Most standards take into account the existence of old equipment and can accommodate much of it. You can use your digital PCS cell phone to call your grandmother on the heavy black dial phone she has used since 1952. Neither you nor your grandmother should notice the automatic electronic conversion between the two formats. You should remember, however, that some new standards do not accommodate old equipment, such as standards for digital television transmission, which make no effort to talk with the "old" analog set you bought brand new in 2004.

Families of Standards

Families of standards have been developed to meet specific transmission requirements. Typically this means sending voice, video, or data signals in different environments. For example, there are standards for how to digitize voice telephone calls and how to interleave the data streams from individual conversations to give higher and higher data rates. Other standards specify formats for data transmission over certain types of computer networks.

Standards for voice, video, and data signals evolved separately and remain somewhat distinct. Many standards are administered by different organizations, or by different subgroups within large organizations such as the International Telecommunications Union. Table 17.2 lists many of the organizations involved in fiber-optic system standards.

Telecommunications standards evolved to meet specific industry requirements. In general, they specify interfaces rather than internal operations. In other words, what matters is not what goes on inside the box, but what goes into and comes out of the box. For example, it doesn't matter if your home telephone looks like a beer can, Mickey Mouse, or an antique pay phone as long as it sends and receives standard voice signals. This provides industry the flexibility to improve the technology as long as it stays within performance standards.

Traditional circuit-switched telephone transmission has well-established standards that range from requirements for wire-line telephones to high-speed data-transmission rates. The telephone industry designed this family of standards to work together, and many of these standards are used in long-distance transmission.

Standards for transmitting computer data now center on the packet-switching formats used on the Internet. These standards include Ethernet and the Internet Protocol, and can handle data flow at uneven rates. Packet-switching standards continue to evolve to handle other services. An example is *Voice over Internet Protocol (VoIP)*, which converts voice signals into packets for Internet transmission.

Video transmission standards are largely industry-specific, developed by television broadcasters and cable companies. Broadcast signal formats require approval by government

agencies that regulate broadcasting, such as the Federal Communications Commission in the United States. Video transmission is evolving slowly from old analog formats to new digital formats, producing completely new standards that are incompatible with old equipment.

We will explore specific standards later, but first you should learn how modern standards are structured.

Layers of Standards

Modern standards
are structured in
layers that serve
distinct functions.

Modern standards are developed as a series of *layers*, each of which serves a distinct function. Essentially, each layer provides a set of interfaces for users of that layer, which effectively covers over deeper layers that the users don't need to worry about. The layered structure comes from the *Open System Interconnection (OSI)* model developed by the International Organization for Standardization. Many older standards, like those for digital telephone transmission, have been modified to fit the OSI layered model. Analog services appear only as inputs to the top layer of the stack.

You can see how a layered standard operates by considering the voice layer of the modern telephone network, shown in Figure 20.1. When you make a telephone call, you dial a phone number and hear the phone ring and a person answer, just as if the call were traveling over standard telephone wires, whether you're calling across town or across the country. You shouldn't notice that your voice is chopped into bits that are interleaved with bits from other telephone calls. All these operations are in the "cloud" of the telecommunications system (shown at the bottom of the figure). The system acts like you're calling one of your neighbors over dedicated local phone lines (shown at the top of the figure). The network reaches both cell phones and regular phones, and you can tell them apart only by their sounds. Faxes also travel over voice telephone lines, and they respond the same way whether you call a fax machine or a computer with fax software.
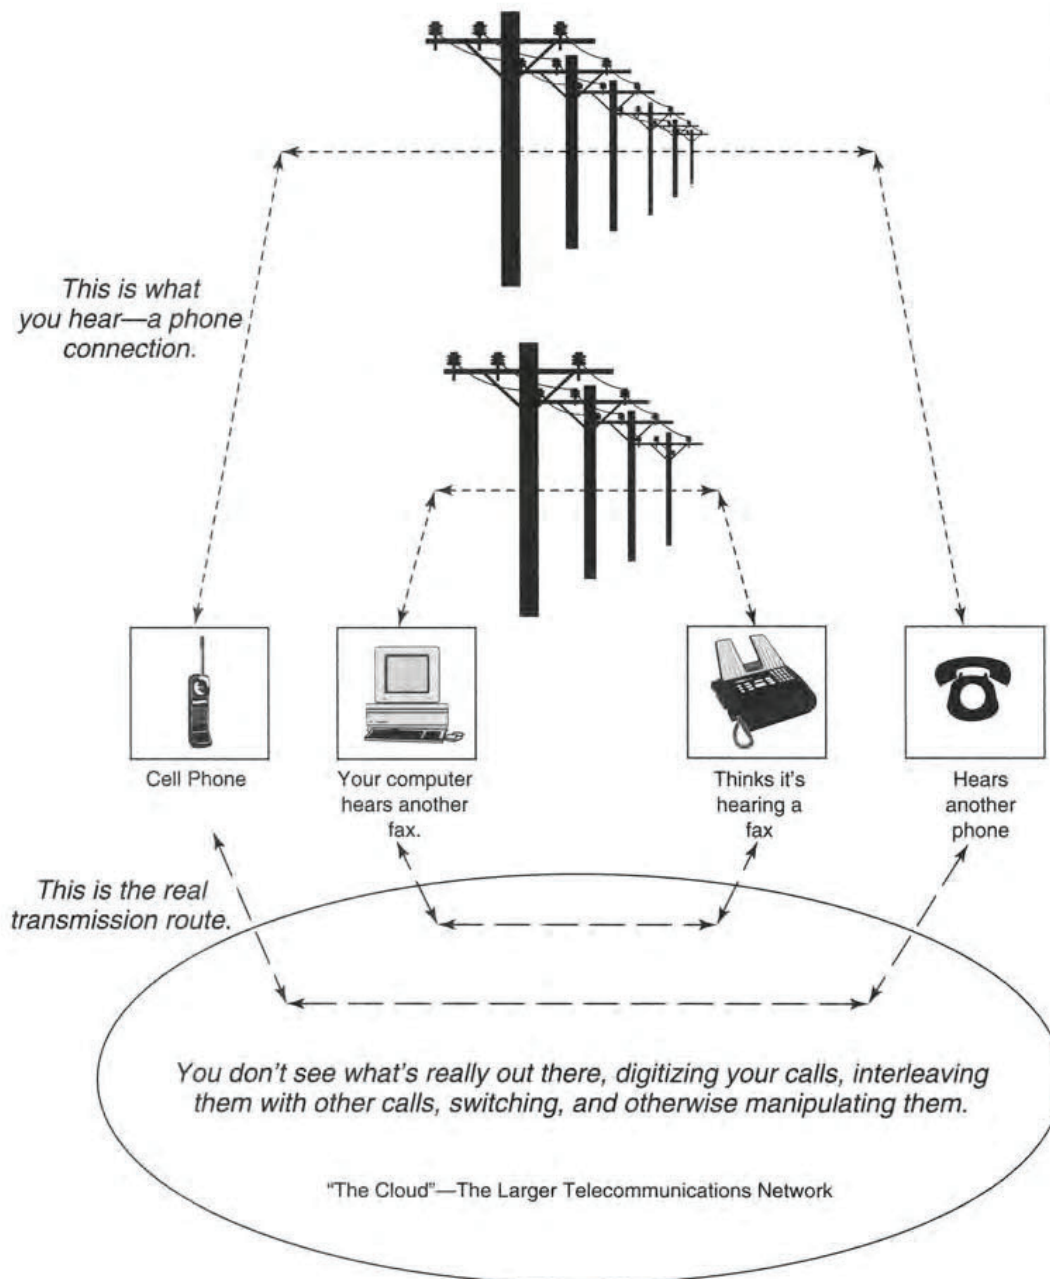
A layer in a layered standard can be viewed as providing interfaces to black boxes that represent other layers. The engineer who designs a phone, for example, works only on the voice layer of the telephone network, with no need to know about the other layers. The engineer is concerned only with the standard interfaces the phone has with the user and the network, which are other layers in the OSI model.

Layers in the Network

The services most
users see are the
top of a stack of
layers.

Figure 20.2 shows the overall layered standard structure in a somewhat simplified form that neglects data-transmission details normally handled by software. It shows the layers used in the telephone network and the Internet. There has been serious discussion about simplifying this layered structure to reduce equipment costs, but change has been slow. Note the layers were numbered before the emergence of wavelength-division multiplexing.

Each service shown in the top layer has its own standard format. Some of these services actually perform multiple functions for the user. Standard analog voice telephone lines can transmit signals from faxes and dial-up modems as well as voice. The faxes and modems send digital data as audio tones that the phone network processes like ordinary voice phone

FIGURE 20.1
Voice layer of the telephone system.

This is what you hear—a phone connection.

| Cell Phone | Your computer hears another fax. | Thinks it's hearing a fax | Hears another phone |

This is the real transmission route.

You don't see what's really out there, digitizing your calls, interleaving them with other calls, switching, and otherwise manipulating them.
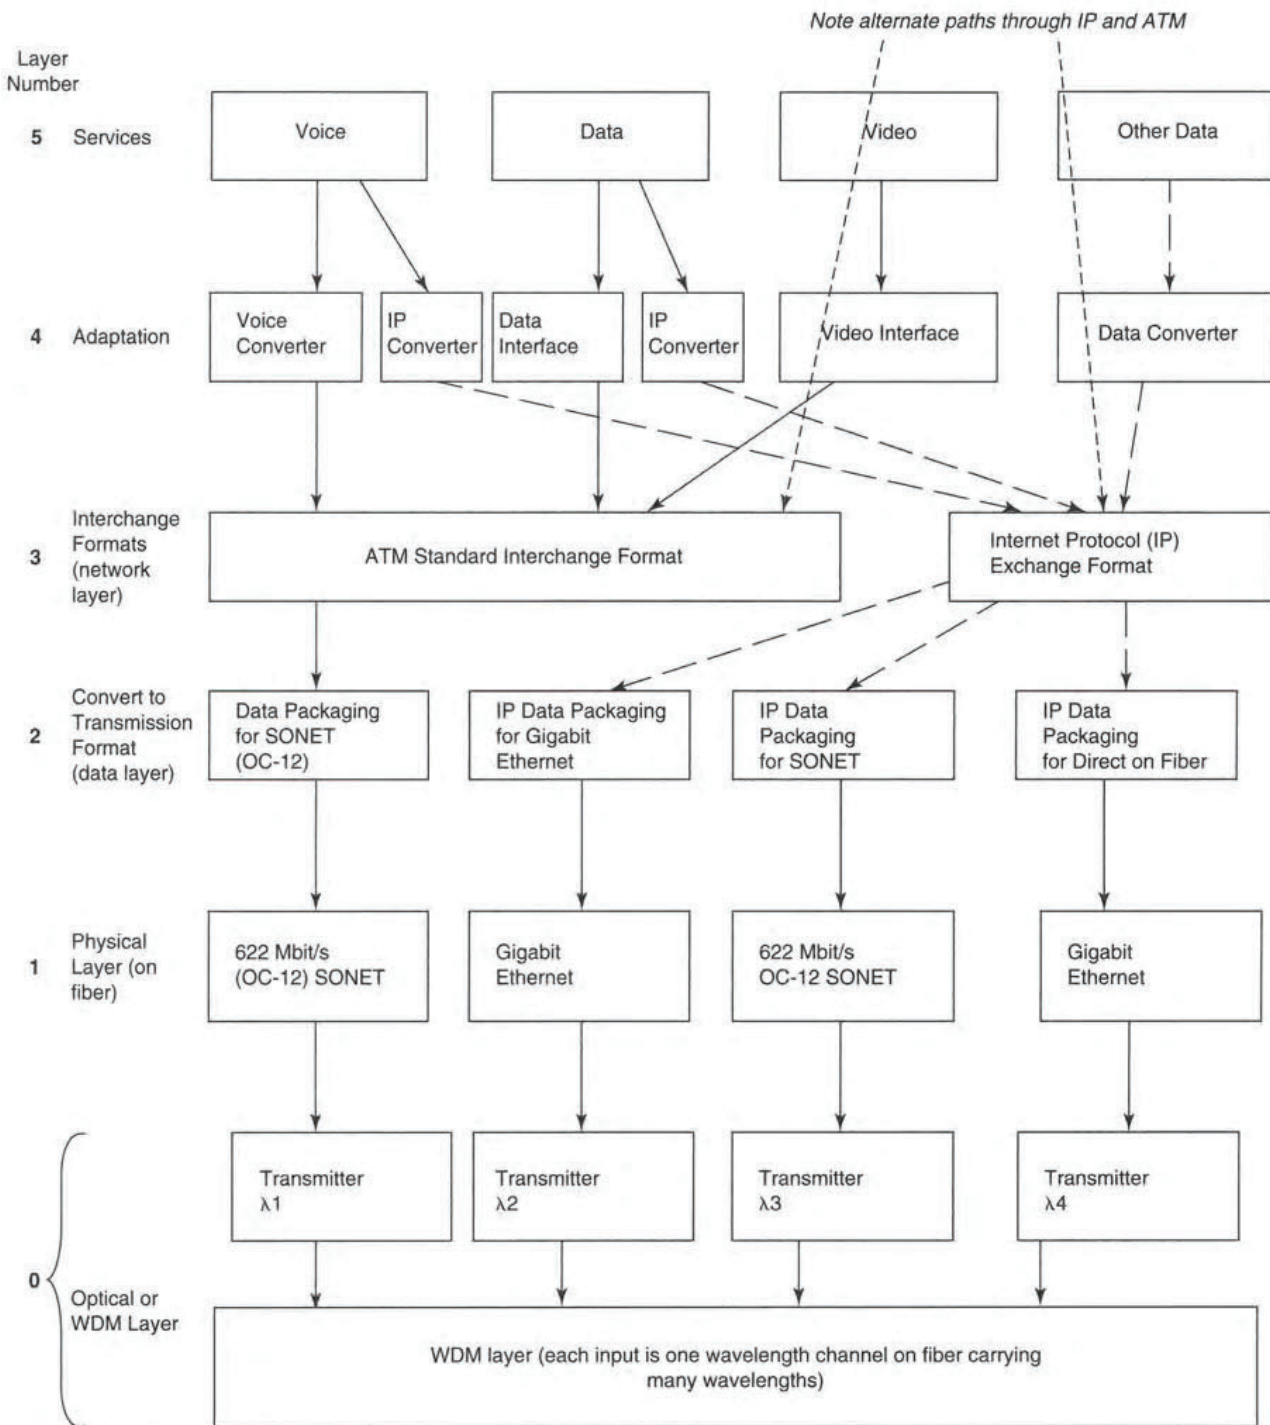
"The Cloud"—The Larger Telecommunications Network

calls. Broadband Internet connections provide service directly in digital form. This top layer is called the *services layer* or *layer 5*.

The next layer down, layer 4, is the *adaptation layer*, which converts signals into the formats needed for data interchange. This includes converting analog signals into digital format, and converting digital signals from special service formats into those used for data interchange.

The *network* or *interchange layer*, known as layer 3, covers the interchange of data packaged in standard formats. The two most important standards are the *Internet Protocol (IP)*,

The interchange layer uses IP and ATM formats.

Note alternate paths through IP and ATM

Layer Number		
5	Services	

Voice   Data   Video   Other Data

| 4 | Adaptation | |

Voice Converter   IP Converter   Data Interface   IP Converter   Video Interface   Data Converter

| 3 | Interchange Formats (network layer) | |

ATM Standard Interchange Format

Internet Protocol (IP) Exchange Format

| 2 | Convert to Transmission Format (data layer) | |

Data Packaging for SONET (OC-12)

IP Data Packaging for Gigabit Ethernet

IP Data Packaging for SONET

IP Data Packaging for Direct on Fiber

| 1 | Physical Layer (on fiber) | |

622 Mbit/s (OC-12) SONET

Gigabit Ethernet

622 Mbit/s OC-12 SONET

Gigabit Ethernet

Transmitter $\lambda 1$

Transmitter $\lambda 2$

Transmitter $\lambda 3$

Transmitter $\lambda 4$

| 0 | Optical or WDM Layer | |

WDM layer (each input is one wavelength channel on fiber carrying many wavelengths)

FIGURE 20.2

Standard layers for communications services, data interchange and transmission.

developed for Internet data exchange, and *Asynchronous Transfer Mode (ATM)*, developed for data and voice transmission on telephone networks. Interchange standards specify how bits are arranged, but not what medium transmits the bits. Many different signals intermingle in the interchange layer, where they are merely packages of data bits accompanied by label bits that identify them and describe their purpose. (ATM and IP formats do not serve identical purposes, and sometimes IP signals are repackaged into ATM format for transmission on telephone lines, although this is not shown in Figure 20.2.)

The *data conversion layer* is below the network layer. Called *layer 2*, it repackages data from the interchange layer into standard forms suitable for transmission on various media.

The *physical layer* or *layer 1* is the medium that carries signals, such as the optical fiber. The network transmits bits through this medium in specific formats. Figure 20.2 shows three of these standard formats:

- *Synchronized Optical Network (SONET)*, a set of standard time-division multiplexed data rates developed by the telephone industry, one of which is 622 Mbit/s.
- *Gigabit Ethernet*, an Ethernet format carrying 1 Gbit/s.
- Raw IP format data at 1 Gbit/s.

Each data stream in the physical layer corresponds to a series of light pulses passing through an optical fiber. The physical layer is usually called *layer 1*, and was originally the bottom layer for data transmission. Once in this format, signals were sent through an optical fiber to their destination.

The development of WDM created another possibility. Data streams generated in the physical layer can modulate signals at separate wavelengths transmitted through a single fiber; that is, as shown in Figure 20.2, each channel in the physical layer can correspond to one wavelength in an optical fiber carrying a WDM signal. In this case the *optical* or *WDM layer* is the bottom layer (sometimes called *layer 0*) of a telecommunications network.

WDM adds another layer beneath the physical layer.

Implementation of Layered Networks

As shown in Figure 20.2, signals may follow different paths through a layered network. Standard wire-line telephone calls may follow the path shown by the solid line at left. They first are converted into ATM format for interchange, then into SONET format for transmission on the physical layer. Internet data may follow the path shown by the dashed line at right, first put into IP format for interchange, then transmitted on the physical layer in Gigabit Ethernet format.

Other paths are possible if you follow other series of arrows through the stack of layers. For example, voice may be converted first to IP format, then transmitted over the Internet as Voice over Internet Protocol. Conversely, Internet data may remain in IP format on the interchange layer, then be packaged for SONET transmission in the physical layer.

Actual network implementation may differ from the nominal layered standard model. The layered structure is designed for modularization, so engineers can plug in modules from other layers knowing only interface specifications, not design details. Modularization can be a good thing if different organizations are building or operating different parts of the network, because it encourages competition among many vendors of services and

equipment. However, implementing each layer requires costly equipment, and some layers add extra bits to manage the transmitted signal, building up overhead. If signals travel only within a single network that uses the same standards throughout, it's often cheaper and easier to avoid conversion to interchange formats and go directly to the format used on the physical layer. Signals require interchange formats only when passing through two or more networks that use different formats. Signals may even remain on the services layer if they do not travel beyond the service network, such as a local-area network for computer data or a local telephone exchange for voice.

The details of layered standards are beyond the scope of this book. Much of what happens in the service and adaptation layers is more in the realm of software than fiber-optic hardware. The standards that are important for fiber optics deal with the interchange, physical, and optical layers. These layers differ in important ways, particularly in how they direct or switch signals.

Circuit and Packet Switching

Routers operate in the interchange layer.

In Chapter 19, you learned the fundamental difference between circuit and packet switching. The two types of switching also are used on different layers, as shown in Figure 20.3. Routers provide packet switching for the interface layer and circuit switches direct transmission on the physical layer. This reflects the different functions of the two layers.
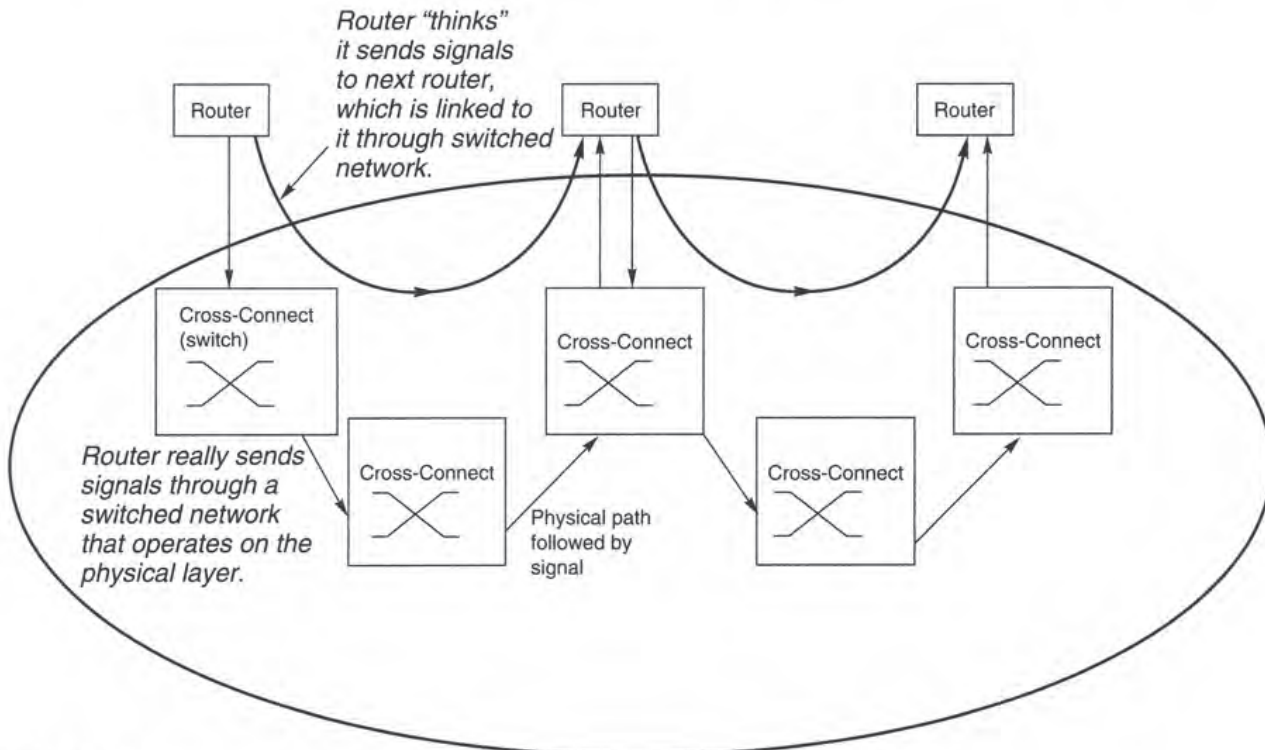
FIGURE 20.3
The interchange layer uses routers; the physical layer uses circuit switches.

The interface or networking layer organizes input signals for efficient transmission over the physical layer. Packet switching is well suited for that purpose. A router reads the packet headers, which specify each packet's destination, and may assign a transmission priority based on the type of signal in the packet. The router uses this information, as well as factors such as network traffic, to manage the data flow in the interface layer, and places the packet in a queue for transmission to the physical layer. From the viewpoint of the interchange layer, the signals appear to be switched from router to router, although they actually follow a more complex path through the physical layer, as shown in Figure 20.3.

The packet-switching standards for the interface layer don't specify a data transmission rate. That's because the actual transmission rate depends on the physical layer. The interchange layer is like a machine that reads labels and stacks parcels in a queue, and the physical layer resembles a conveyor belt that removes parcels at a fixed rate. The two work together, but have different functions. The interchange layer organizes the data packets, but the speed of the physical layer determines the data rate.

The physical layer transmits data through a physical medium—an optical fiber, one wavelength on a WDM fiber, or some other medium. The physical layer does not sort individual packets; it carries a stream of data from point to point at a constant rate. The physical layer uses circuit switches that make connections without examining the packets. The switches thus make a virtual circuit between two points that may go through a series of circuit switches—like the path between the first two routers in Figure 20.3.

Circuit switching is on the physical layer.

This distinction between circuit and packet switching is sharpest for Internet data transmission, where routers "groom" data flow on the interchange layer, and circuit switches are hidden below in the physical layer. Telephone traffic does not have to pass through routers at all, but if it does, the routers must make the service seem circuit-switched.
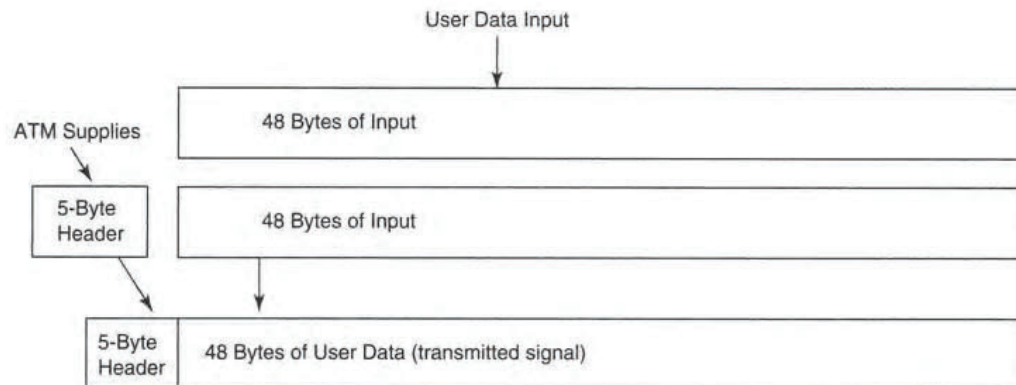
Interchange Standards

The two primary interchange standards are *asynchronous transfer mode* (ATM) and the *Internet Protocol* (IP). Both package data into packets and append headers containing the data's destination and other important information. They differ in their intended purposes. ATM was developed by the telecommunications industry to handle voice, video, and data, where priority is essential, and it behaves as if it were circuit-switched. IP was developed for data transmission, and initially did not assign priority to data packets; it acts packet-switched.

Asynchronous Transfer Mode (ATM)

The ATM standard breaks incoming data into 48-byte blocks of user data, to which it prefixes a 5-byte header, producing fixed-length *cells* of exactly 53 bytes, as shown in Figure 20.4. The header contains the destination address, information on the data stream, and a priority code assigned to the data. Two-way voice (telephone) traffic receives the highest priority, because it requires a constant bit rate with no delay to provide adequate service. Video also receives a high priority because it also is vulnerable to delays. Data receives a lower priority, because delays are not considered critical. The short cell length was chosen for maximum throughput for the

ATM creates virtual circuits.

FIGURE 20.4

Asynchronous transfer mode cells.

User Data Input

ATM Supplies

| 48 Bytes of Input |

5-Byte Header | 48 Bytes of Input |

5-Byte Header | 48 Bytes of User Data (transmitted signal) |

mixture of voice, video, and data that was expected when the ATM standard was developed many years ago, but is not ideal for modern levels of data transmission because it increases overhead.

ATM designers wanted to combine the cost advantages of packet switching with the "quality of service" needed for voice traffic. Voice and video both require a guaranteed minimum capacity. The capacity of the transmission line limits the total number of virtual voice and video circuits, but ATM guarantees each virtual voice and video circuit the capacity it needs. Data receives a lower priority, and essentially fits into the spaces between the voice and video signals. That isn't as bad as it sounds, because voice conversations are not continuous, but data may have to wait for quiet intervals or other events that make transmission slots available.

Strictly speaking, ATM is packet-switched, but the telephone user shouldn't notice its presence. ATM is a resource-allocation scheme that helps data signals share phone lines without degrading the voice signals.

Internet Protocol (IP)

● IP uses variable-length packets.

The Internet Protocol (IP) is another packet-switched standard that breaks incoming data into packets of variable length instead of into fixed-length cells. The header contains information on the packet's length as well as its destination.

The IP standard has evolved over the years to interconnect networks transmitting digital data. It was originally designed for "internetworking"—the interchange of data between networks at different institutions, some with different architectures and data formats. This purpose firmly entrenched it in layer 3, the interchange or networking layer of the standard system.

● IP packets are stored in a buffer until transmission.

The flow of computer data fluctuates much more widely and quickly than the volume of voice telephone calls. Computers can tolerate transmission delays as long as the data packets can be reassembled in the proper order, so data networks average the flow over time to fill available transmission capacity more evenly. They do this by putting packets in a buffer storage until outgoing lines are available to transmit them. Figure 20.5 shows how this works in a simple system that can transmit 1000 bits per microsecond. If the packets that arrive during the first microsecond total only 100 bits, blank bits are transmitted along with the 100 data bits. If the packets arriving during the next microsecond contain 1500 bits, 1000 bits are transmitted and the other 500 bits sit in a buffer, to be transmitted during the next microsecond. If 600 bits
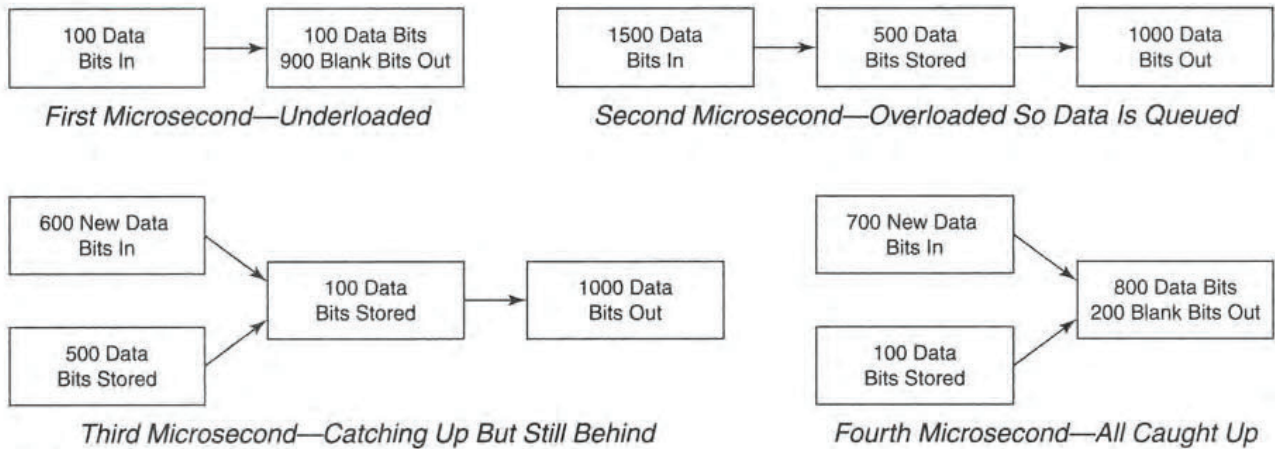
FIGURE 20.5

Packet switch stores extra bits until it can catch up on transmission.

arrive in the following microsecond, the 500 bits in the buffer and the first 500 new bits are transmitted, with 100 bits left over in the queue. If 700 bits come in the next microsecond, the 100 leftover bits are transmitted, then the 700 new bits, emptying the buffer.

Designers originally assumed that the Internet would handle only computer data, so IP headers did not assign priority levels to each packet, unlike ATM. This is still the case with the most widely used version of Internet Protocol, IPv4. With no priority codes, each IP node attempts to transmit each packet in the order in which it arrived on a "best effort" basis, so packets are transmitted in the order they were processed. The headers include counters that record how long the packet has been waiting for transmission. Packets that are stuck for a long time eventually time out and are discarded. That means your browser sometimes waits for a packet that will never come. A newer version of Internet Protocol, IPv6 (there is no version 5), assigns priority codes to packets, but is not yet widely implemented.

The lack of priority codes limits how well the Internet can transmit high-priority signals such as voice. Telephone calls can be made over the Internet using a protocol called Voice over IP (VoIP), but the quality of the service depends on Internet traffic. If traffic is low to moderate, packets experience few delays and the quality is good. If traffic is high, packets can be delayed and signal quality suffers. Implementation of IPv6 is expected to improve voice quality by reducing delays.

> IPv4 lacks priority codes for data packets.

An important feature of IP is that the total capacity of input lines to a router can be greater than the total capacity of the output lines. This is possible because the average load on the input ports normally is well below the peak capacity. For example, the average load on 100-Mbit/s lines might be only 5 Mbit/s, so 20 input lines could be connected to a router linked to a single Gigabit Ethernet output without causing serious data traffic jams.

Fiber Transmission Standards

Fiber-optic transmission standards function as physical layer standards, although some of them don't fit into the simplified picture of Figure 20.3, which shows the physical layer as the core of the network between routers. The physical layer also exists outside the network

> Physical layer standards cover transmission on fibers.

core, delivering signals to remote equipment or servers. Physical-layer transmission differs in fundamental ways from the interchange or network layer.

The physical layer is essentially a set of transmission lines that deliver data between pairs of points. In the network core, these points may be individual devices or circuit switches, which direct streams of data along dedicated transmission lines or virtual circuits. In this sense the physical layer works like traditional telephone lines in the core of the network, although it may include routers on the periphery.

Physical layer standards cover long- and short-distance links.

Physical layer standards structure digital transmission for specific media, specifying data transmission rates and formats. Some of these standards, such as SONET, SDH, and the digital telephone hierarchy, are designed for long-distance transmission in the network core. Others such as *Fibre Channel* and *Ethernet* are designed for specific types of networks on the periphery. All of these standards provide particular functions needed by the applications they cover. SONET was developed by the telephone industry to guarantee integrity of long-distance transmission lines, so it monitors network operation. Fibre Channel was designed specifically for storage-area networks, which connect data banks to computers and provide backup data storage. Other formats, including raw IP data streams, are also used in the physical layer, although those formats may lack important monitoring functions.

We won't go into the details of these standards and their relationship with the layered model. You don't need to worry about them while you're learning about fiber optics. But you should know how these standards work, so we'll look briefly at the most important ones.

Digital Telephone Hierarchy

Digital telephone lines are based on a hierarchy of TDM rates.

Back when AT&T monopolized telephone service in most of the United States, it devised standards for digital telephone transmission. This digital telephone hierarchy remains in use at data rates to 45 Mbit/s. This standard is called the *plesiochronous digital hierarchy* or PDH, because many independent clocks provide timing for a hierarchy of data rates. The International Telecommunications Union devised a similar—but not identical—standard used in much of the rest of the world. (You'll learn more about these standards in Chapter 23, which explains the global telecommunications network; Appendix C lists the data rates.)

The digital telephone hierarchy is a sequence of these time-division multiplexed data rates. It starts with the electronics that convert a standard analog voice telephone signal into digital format at 64,000 bit/s. The bit streams from 24 digitized phone lines are interleaved to give a single stream of 1.5 Mbit/s (called the *DS1 rate*), which is transmitted through T1 lines that may be fiber but often are copper. T1 lines, in turn, feed multiplexers that generate 6-Mbit/s T2 and 45-Mbit/s T3 lines. The top of the original hierarchy, the T4 rate of 274 Mbit/s, is no longer used. This hierarchy provides virtual circuits in the form of reserved slots in the data stream. The data streams were picked to match the needs of voice transmission, and they are not broken into packets.

The digital telephone hierarchy has some drawbacks. One is that the only way to extract a lower-speed signal from one at higher speed is to break the entire signal into its component parts. Another drawback is that the data stream lacks the control signals needed in a modern network.

SONET/SDH

The *Synchronous Optical Network* (SONET) standard calls for long-distance transmission over fibers at a hierarchy of digital rates higher than those of the old digital telephone hierarchy. SONET is the North American standard; its international counterpart (set by ITU) is the *Synchronous Digital Hierarchy* (SDH). SONET and SDH differ only slightly, to accommodate differences in lower-speed telephone standards. Both standards interface with the ATM interchange format and provide features not available in the old digital telephone hierarchy.

SONET organizes data into 810-byte blocks called *frames.* Each frame includes 27-byte headers containing information on signal routing and destination in addition to the signal data. The system generates and inserts headers as it packages input signals into SONET frames. The frames can be switched individually without breaking the signal up into its component parts.

Developed to carry mixed traffic over fiber, SONET/SDH explicitly defines a series of time-division multiplexed transmission rates (designated *OC-x*). The SONET base rate (OC-1) is 51.84 Mbit/s, which with overhead accommodates the widely used T3 rate of the North American digital telephone hierarchy. Next is the 155.52 Mbit/s OC-3 rate, nominally produced by interleaving frames from three OC-1 signals; this rate is the base level of SDH (SDH-1). Above OC-3 are OC-12 at 622 Mbit/s, OC-48 at 2.5 Gbit/s, OC-192 at 10 Gbit/s, and OC-768 at 40 Gbit/s. Appendix C lists these transmission rates.

The frame structure in SONET allocates a fixed number of slots per second for input signals. Although the frames may look like packets, SONET uses them to provide circuit-switched TDM services, which guarantee transmission capacity. SONET frames also can handle packet-switched data, but they do so by creating a virtual circuit to transmit the signal on the physical layer.

The SONET/SDH standards do more than specify data rates and frame sizes. They also specify that the network be arranged in the ring topology shown in Figure 20.6, rather than in the older hub-and-spoke or branching system. A SONET ring includes a complete set of redundant fibers, which enables the SONET hardware to reroute signals if a cable is broken. Thus the system continues operating when a cable fails.

> SONET/SDH is a hierarchy of digital rates for fiber links.

> SONET frames guarantee transmission capacity.

Ethernet Standards

The *Ethernet* family of standards transmit data to and from personal computers on local-area networks and are the most familiar physical layer transmission standards used in fiber optics. Ethernet operates on the physical layer, formatting data to be transmitted through networking cables, but it also works as a packet-based standard on the data layer (layer 2) and is switched by routers.

On the physical layer, Ethernet transmits two error-correcting bits with each eight-bit byte, a total of 10 bits per byte. The original standard, just called *Ethernet,* supported 10 Mbit/s data transmission. Later versions are *Fast Ethernet* at 100 Mbit/s, *Gigabit Ethernet* at 1 Gbit/s, and *10 Gigabit Ethernet* at 10 Gbit/s. In practice, fiber is common for Gigabit Ethernet and standard for 10 Gigabit Ethernet (although copper may be used over short distances). Standard Ethernet and Fast Ethernet usually are transmitted on copper.

> Ethernet is a family of physical layer standards.
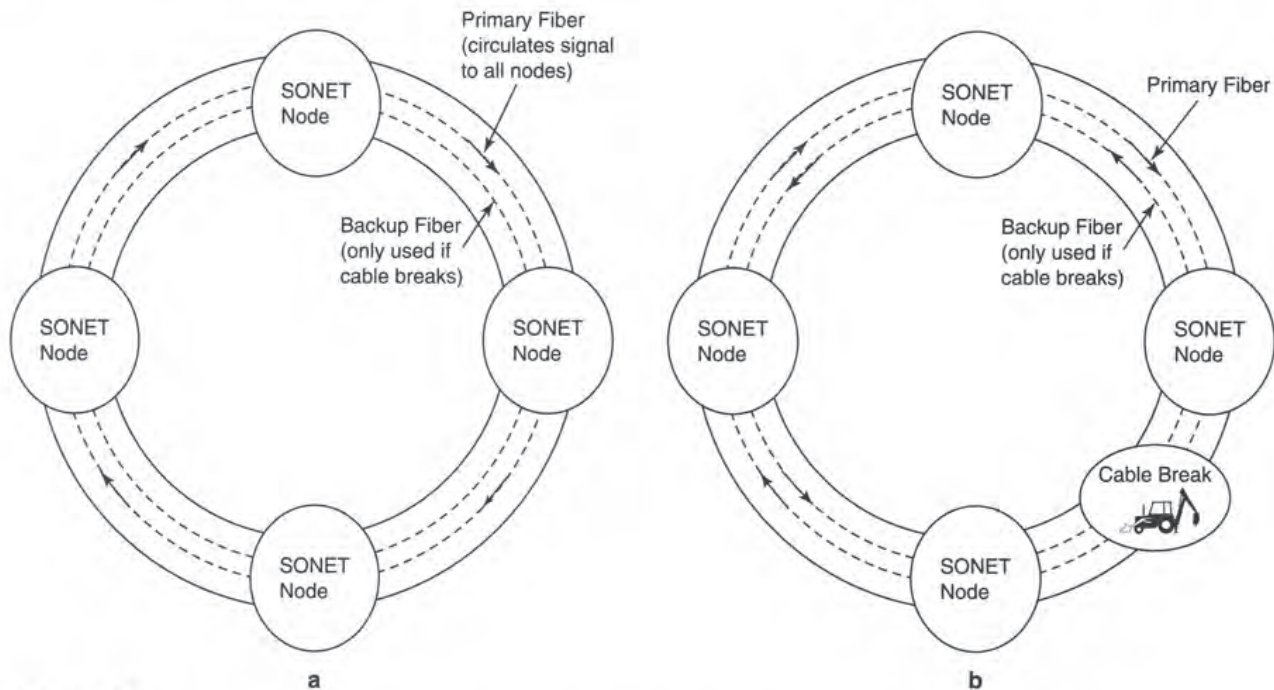
FIGURE 20.6

Intact SONET ring links all nodes through one fiber. Broken SONET ring still links all nodes through primary and backup fibers.

Although data rates increase by a factor of 10 in each step to higher speed, Ethernet is not a time-division multiplexed system that interleaves data bits or frames. It's a packet-switched system that merges streams of packets, so the potential input data rate does not have to match the specified output rate; for example, you can have more than 10 Fast Ethernet inputs to a Gigabit Ethernet link. You will learn more about Ethernet transmission in Chapters 25 and 26.

Fibre Channel

> ● Fibre Channel was developed for storage-area networks.

The *Fibre Channel* standard was developed to relay data from computer networks to data banks on site or in remote locations. This application is called a *storage-area network*, and generally does not require spanning long distances. The British spelling, *fibre*, was chosen to indicate that this standard applies to metal cables in addition to single-mode and multimode fibers. Fibre Channel transmits at data rates of 1, 2, 4, and 10 Gbit/s, and like Ethernet adds two error-correction bits to each eight-bit data byte.

Fibre Channel groups data into frames, encodes and decodes the frames, and physically transports them through various media. These functions correspond to the interchange, data, and physical layers shown in Figure 20.2. Fibre Channel can handle point-to-point transmission between a computer and storage devices, transmission around a loop, or transmission through a switched network.

Current Standards Issues

So far we've discussed mature standards that have been codified and approved. Because fiber-optic technology changes fast, some standards-related issues remain in a state of flux. The telecommunications bubble raised unrealistic expectations about the rate of traffic growth, and venture-funded companies charged ahead with over-optimistic visions of the future. Many new ideas have not been formally codified and approved as standards, leaving important technical and commercial issues.

Optical Layer Standards

The OSI model for telecommunications was developed before wavelength-division multiplexing became practical. The emergence of practical WDM technology led to the possibility of an optical layer as an extension of the physical layer, with individual wavelengths serving as physical channels that are merged into a single optical signal. Signals could be processed in the optical domain using wavelength conversion and all-optical switching.

The optical layer exists "under" the physical layer.

From a standards viewpoint, the optical layer is below the physical layer, as shown in Figure 20.7. Standards specify important features needed for wavelength-division multiplexing, such as channel wavelengths and spacing. It also is possible to split the optical layer into sublayers for optical channels (which carry physical-layer signals), optical multiplexing, and the actual optical transmission.

Currently, many issues need to be resolved before the optical layer can become a reality. Should an "optical wrapper" of overhead data be added in the optical layer, just as headers are added to data packets in the interchange layer? Should the optical layer be a "transparent" network, where switching and transmission are all-optical? Or can optical signals be converted into electronic form for some operations, creating an "opaque" network (because no single photon can pass through the entire network)? How much signal processing needs to be done optically?

Video Standards and Copy Protection

Video standards are in a state of flux as the industry moves from analog to digital format. Existing analog equipment doesn't work with the digital technology used in other types of

FIGURE 20.7

Optical networking layers.

THINGS TO THINK ABOUT

Market Pressures and Standards

Standards are compromises that companies reach in order to serve market needs. Sometimes market needs conflict so sharply that these compromises are very difficult to reach. Such is the case with *digital rights management*, technology that restricts the copying of digital media.

Entertainment and consumer electronics companies make money in different ways. Entertainment companies want to sell as many programs as possible to as many people as possible, so they don't want free copies circulating. They're glad to sell you copies of their programs, but they don't want you giving copies of them away to other people. In fact, they'd really like to charge you each time you watch or listen to them. They'd like to ban all copying, as was evident two decades ago, when they tried unsuccessfully to outlaw analog video recorders.

On the other hand, consumer electronics companies want to sell video recorders and home networking equipment used to record and distribute programs. They make money when you buy hardware, but not when you buy programs, so they defend *limited* copying. (Nobody openly advocates mass-producing unauthorized copies of programs.) A big problem today is striking a balance between the rights of customers who own legitimate copies and the interests of entertainment companies.

---

communications, so analog signals have long been transmitted separately. Yet efforts to integrate video with other communications services have run into serious commercial problems.

Controversy surrounds efforts by the entertainment industry to ban the sale and use of hardware and software that allows digital copying and distribution of audio and video programs. This is a serious loss to consumers who make copies strictly for personal use, like recording television programs for later viewing or mixing their own audio recordings. Limits on distribution of digital programs could hurt development of home entertainment networks that could record programs from the family-room set and send them to a bedroom set. The electronics industry worries that restrictions on program use will hurt its sales, and telecommunications carriers don't want to monitor their customers' activities. This conflict has stalled development of video standards on multiple fronts.

Proprietary Technologies and Market Pressures

Proprietary equipment doesn't follow standards.

Standards are agreements made by companies to manufacture equipment that provides specified functions and uses specified formats. The goal of standards is to create an open market for vendors who agree to follow them. However, some companies decide not to follow standards and offer proprietary equipment that is unavailable from other companies.

Depending on your viewpoint, proprietary technologies are either the best systems companies can offer to their customers, or an effort to trap customers into buying more equipment in the future. In practice, all systems are somewhat proprietary on the inside, as modern standards cover interfaces, not inner workings. For example, some carriers may not convert IP signals into SONET for transmission on their networks if their physical fiber-optic connection takes IP signals directly to the next router. These carriers wouldn't need

physical layer interfaces inside their own networks, so they could just transmit raw IP format signals at a data rate the link can handle.

Market pressures ultimately determine what standards are accepted and widely used by industry.

What Have You Learned?

1. Standards specify coding techniques so different systems can understand each other. The importance of standards increases with the scale of the system.

2. Families of standards exist for different types of telecommunications, and for different telecommunication functions. They typically specify interfaces and formats that equipment can use to connect and transfer signals.

3. Data transmission standards rely largely on packet switching. Packet headers contain information needed to route packets. Circuit switching is used in many telephone standards.

4. Modern standards consist of a series of layers, each serving its own function. Signals pass through the layers, but users don't see the lower layers. Major layers (from top down) are the service, adaptation, interchange (or network), data conversion, and physical layers.

5. Internet Protocol (IP) and Asynchronous Transfer Mode (ATM) standards apply to the interchange layer. Both standards group data into packets. Routers operate on the interchange layer.

6. SONET (Synchronized Optical Network), the digital telephone hierarchy, and Ethernet apply to the physical layer. The physical layer includes the core of the network, as well as links between devices such as computers and external backup storage. Circuit switches are found in the physical layer.

7. The digital telephone hierarchy is a series of time-division multiplexed digital data rates to 45 Mbit/s. It was originally developed in the 1960s but is still in use.

8. Asynchronous Transfer Mode creates virtual circuits although it is a packet-based protocol. It has priority codes.

9. Internet Protocol was designed for internetworking. It has variable-length packets and is packet-switched. IPv4 does not assign priority codes to packets, but IPv6 does.

10. Wavelength-division multiplexing adds an optical layer below the physical layer. The final shape of the optical layer is still in development.

11. Queued packets are transmitted in the order they were received, as fast as the output port can take them. If they stay in the queue too long, they time out.

12. SONET/SDH specifies digital data transmission in frames at an ordered series of rates starting at 51.84 Mbit/s through optical fibers. It specifies a ring topology to guarantee that service continues if a cable is broken.

13. Fibre Channel is a standard for storage-area networks.

14. The Ethernet family includes standard Ethernet at 10 Mbit/s, Fast Ethernet at 100 Mbit/s, Gigabit Ethernet at 1 Gbit/s, and 10-Gigabit Ethernet at 10 Gbit/s.

15. Analog video standards are treated separately from digital standards. Digital video standards are still in development; a key issue is digital rights management.

What's Next?

In Chapter 21, you will learn the basic elements of designing optical systems that transmit at a single wavelength. Chapter 22 will cover optical networking design.

Further Reading

Vivek Alwayn, *Optical Network Design & Implementation* (Cisco Press, Indianapolis, 2004)

John C. Bellamy, *Digital Telephony*, 3rd ed., (Wiley InterScience, New York, 2000)

Roger L. Freeman, *Fundamentals of Telecommunications* (Wiley InterScience, New York, 1999)

Gil Held, *Voice & Data Internetworking* (McGraw-Hill, New York, 2000)

Jean Walrand and Pravin Varaiya, *High Performance Communication Networks* (Morgan Kaufmann, San Francisco, 2000)

Web Resources

ATM: *http://www.atmforum.com/*

Fibre Channel: *http://www.fibrechannel.com*

Gigabit Ethernet: *http://www.gigabit-ethernet.org/*

Questions to Think About

1. Follow the voice signals in Figure 20.2 through the layers in the diagram. What function does Internet Protocol serve?

2. Some makers of telecommunications equipment propose to transmit IP signals directly on fiber, without going through ATM or SONET coding. What advantages might this have?

3. What difference between ATM and IPv4 formats is most important for voice transmission?

4. Packet switching has a major advantage in that it combines signals that arrive at uneven rates to use transmission capacity efficiently. Suppose you have four packet-switched input signals, which can arrive at peak rates of 1 Gbit/s. However, on average the packets account for only about 20% of the peak capacity. If all goes well, can you squeeze those four input channels through a 1-Gbit/s output?

5. You can pack 24 voice channels on one T1 carrier, four T1 carriers into a T2 channel, and seven T2 carriers into a T3 signal. How many voice channels can a T3 signal carry?

6. How many voice channels can an OC-192 signal carry, assuming an OC-1 carrier transmits the equivalent of one T3 carrier?

Chapter Quiz

1. Which of the following are *not* defined by telecommunications industry standards?

   a. data transmission formats on optical fiber

   b. transmission speeds in digital telecommunications

   c. interchange formats for signals sent to other countries

   d. data transmission in local-area networks

   e. monthly telephone service charges

2. What kind of standard is Asynchronous Transfer Mode (ATM)?

   a. data-interchange format

   b. fiber transmission

   c. analog television

   d. time-division multiplexing

   e. financial transfer for banking

3. Which of the following is a SONET data rate?

   a. T3, 45 Mbit/s

   b. 100 Mbit/s

   c. OC-3, 155 Mbit/s

   d. 1 Gbit/s

   e. none of the above

4. How does packet switching combine signals from different sources?

   a. It assigns each one a different wavelength.

   b. It packages them into a series of packets, with headers indicating destinations.

   c. It assigns each one a different series of time slots in a sequence of bits.

   d. It transmits them simultaneously at different frequencies.

   e. None of the above

5. On what layer(s) of the telecommunication network are routers used?

   a. services layer

   b. interchange layer

   c. physical layer

   d. optical layer

   e. all layers

6. Which of the following statements is true?

　　a. Circuit switching has become obsolete and is no longer used.

　　b. Circuit switching is used for most data traffic; packet switching is used in telephone systems.

　　c. Packet switches can create a virtual circuit in the physical layer between two routers.

　　d. Circuit switches can create a virtual circuit in the physical layer between two routers.

　　e. Circuit switches and packet switches cannot be used in the same stack of layers.

7. How many ATM cells can be stacked into the payload of a SONET frame, neglecting reserved fields other than the header?

　　a. 14

　　b. 15

　　c. 16

　　d. 17

　　e. 20

8. How many IP packets can be stacked into the payload of a SONET frame, neglecting reserved fields other than the header?

　　a. 1

　　b. 4

　　c. 14

　　d. 16

　　e. impossible to tell

9. How does an IPv4 router process data packets?

　　a. It transmits them as fast as they come in because the input must equal the output.

　　b. It deletes packets that do not have a priority code above a specified level.

　　c. It queues packets, then transmits them in the order in which they arrived.

　　d. It gives first priority to time-sensitive packets and delays other packets.

　　e. It rejects packets containing time-sensitive data.

10. How does a SONET circuit switch process SONET frames?

　　a. It transmits them as fast as they come in because the input must equal the output.

　　b. It deletes packets that do not have a priority code above a specified level.

　　c. It queues packets, then transmits them in the order in which they arrived.

　　d. It gives first priority to time-sensitive packets and delays other packets.

　　e. It rejects packets containing time-sensitive data.

11. Time-division multiplexing at rates to 45 Mbit/s is used in which of the following?

   a. SONET

   b. Internet Protocol

   c. ATM

   d. Plesiochronous telephone hierarchy

   e. Fibre Channel

12. The cable that carries data from your computer to a backup storage device is part of the

   a. services layer.

   b. interchange layer.

   c. physical layer.

   d. optical layer.

   e. no layer.

Single-Channel System Design

About This Chapter

Now that you have learned the ideas behind fiber-optic communication systems, it's time to look at how they are designed. Design is a big topic, so it is split into two chapters. This chapter covers design of single-channel systems to meet loss and bandwidth requirements. Chapter 22 covers design considerations for wavelength-division multiplexing and optical networking.

Loss budgets and transmission capacity, or bandwidth, are crucial in both single- and multichannel systems, but the basic principles are the same for both. You calculate loss budget to be sure that enough signal reaches the receiver to give adequate performance. Likewise, you must calculate pulse dispersion, or bandwidth, to be sure the system can transmit signals at the proper speed. Some simple guidelines will give you rough assessments. In the real world, you also must consider cost-effectiveness and make trade-offs among various approaches to find the best performance at the most reasonable cost. Single-channel design techniques can be applied to each channel in a multiwavelength system.

These two chapters will not prepare you for heavy-duty system design. However, they will give you the basic understanding of design concepts and technical trade-offs you need to assess fiber-optic systems.

Variables

Design of a fiber-optic system is a balancing act. You start with a set of performance requirements, such as sending 2.5 Gbit/s through a 5-km cable. You add some subsidiary goals, sometimes explicitly, sometimes implicitly. For example, you may demand cost as low as possible, less than another alternative, or no more than a given amount. Your

system might need an error rate of no more than 10^{-15} and should operate without inter-ruption for at least 5 years.

You must look at each goal carefully to decide how much it is worth. Suppose, for instance, you decide that your system absolutely must operate 100% of the time. You're willing to pay premium prices for transmitters, receivers, and super-duper heavily armored absolutely gopher-proof cable. But how far should you go? If that is an absolute must because of national security and you have unlimited quantities of money, you might buy up the entire right of way, install the cable in ducts embedded in a meter of concrete, and post guards armed with tanks and bazookas to make sure no one comes near the cable with a backhoe. If its purpose is just to keep two corporate computers linked together, you might be satisfied with laying a redundant gopher-proof cable along a second route different enough from the first that no single accident would knock out both.

That facetious example indicates how many variables can enter into system design. In this chapter, I will concentrate on the major goals of achieving specified transmission distance and data rate at reasonable cost in the simple case of a fiber carrying only one optical channel. Many design variables can enter into the equation, directly or indirectly. Among them are the following:

- Light source output power (into fiber)
- Coupling losses
- Spectral linewidth of the light source
- Response time of the light source and transmitter
- Signal coding
- Splice and connector loss
- Type of fiber (single- or multimode)
- Fiber attenuation and dispersion
- Fiber core diameter
- Fiber NA
- Operating wavelength
- Optical amplifiers
- Direct versus indirect modulation of transmitter
- Switching requirements
- Receiver sensitivity
- Bit error rate or signal-to-noise ratio
- Receiver bandwidth
- System configuration
- Number of splices, couplers, and connectors
- Type of couplers
- Costs
- Upgradability of design

FIGURE 21.1

Three ways to interconnect computer terminals.

Many of these variables are interrelated. For example, fiber attenuation and dispersion depend on operating wavelength as well as the fiber type. Coupling losses depend on factors such as fiber NA and core diameter. Some interrelationships limit the choices available. For example, the need to achieve low fiber loss may require operation at 1300 or 1550 nm, and the need for optical amplification may dictate 1550 nm.

Some variables may not give you as many degrees of freedom as you might wish. For example, you may need to interconnect several computer terminals. You have enough flexibility to pick any of the possible layouts in Figure 21.1, but you still have to connect all the terminals, and that requires enough optical power to drive them all.

The type of system dictates the features you consider. If your goal is to connect computer terminals on a single floor of an office building, coupling loss will be more important than fiber attenuation. If your goal is to span the Pacific Ocean, fiber attenuation, dispersion, and optical amplifiers will be your main concerns. Designers of a transpacific cable must carefully consider how to achieve maximum transmission capacity, but the office-network designer must try to minimize coupling losses.

Often you can reach similar performance goals in more than one way. As you will learn in Chapter 22, a typical example is whether to transmit one signal at 10 Gbit/s or four WDM signals at 2.5 Gbit/s. Both choices can deliver a total of 10 Gbit/s, so the choice must depend on other factors, such as the cost, expected reliability, and potential for future upgrades.

Real fiber-optic system design is inherently a complex task, like trying to solve many simultaneous equations in algebra. The best way to understand the concepts is to look at them one at a time before you worry about how they interact. For that reason, the examples that follow are kept simple, without worrying about the complex trade-offs that affect real design decisions.

Power Budgeting

Power budgeting is much like making sure you have enough money to pay your bills. In this case, you need enough light to cover all optical transmission losses and to deliver enough light to the receiver to achieve the desired signal-to-noise ratio or bit error rate.

Power budgeting verifies that enough light reaches the receiver for proper system operation.

That design should leave some extra margin above the receiver's minimum requirements to allow for system aging, fluctuations, and repairs, such as splicing a broken cable. However, it should not deliver so much power that it overloads the receiver.

One note of warning: be sure you know what power is specified where. You can lose 3 dB if the transmitter manufacturer specifies output as peak power but the receiver manufacturer specifies average power.

In simplest form, the power budget is

$$\text{Power}_{\text{transmitter}} - \text{total loss} + \text{amplification} = \text{margin} + \text{receiver sensitivity}$$

> The difference between transmitter output and receiver input equals the sum of system losses and margin.

when the arithmetic is done in decibels or related units such as dBm. The simplicity of these calculations is a main reason for using decibel units.

Remember that optical amplification can offset loss in the system budget. Optical amplifiers are expensive, but that high cost is justified in some cases. You wouldn't buy a $3000 optical amplifier so you could replace a $100 laser source with a $10 LED, but you would if you could avoid spending $10,000 on an electro-optic regenerator.

All factors that affect power in the system must be considered. These include:

- Light source output power
- Loss in transferring light from source into fiber
- Connector loss
- Splice loss
- Coupler loss
- Fiber loss
- Loss coupling to receiver
- Receiver sensitivity at data rate
- Forward error correction

Some of these factors have been covered in detail before, but others deserve more explanation.

Light Coupling Losses

> Light is lost coupling into fibers.

Most light coupling losses come in transferring light from a source into a fiber; little is lost delivering light to a receiver. The main problem is matching the source emitting area to the fiber core.

TYPES OF SOURCES

Many LEDs emit from areas on their surfaces that are larger than the cores of 62.5/125 or 50/125 graded-index multimode fibers. As shown in Figure 21.2, when the emitting area is larger than the fiber core, some light enters the fiber core and escapes. The smaller the fiber core and the numerical aperture, the larger the loss. Large-area LEDs couple

FIGURE 21.2

Light losses in transferring LED output into a fiber.

Some light misses the fiber.

Some light goes into the cladding and leaks out.

Some light gets into the core, but is beyond the acceptance angle and leaks out.

LED

Fiber

Only a little light is coupled into the fiber.

efficiently into large-core step-index multimode fibers, but their coupling losses can be up to 13 dB. Losses are lower for edge-emitting LEDs, which have smaller emitting areas than surface emitting LEDs. Power coupled into a graded-index fiber typically is tens of microwatts.

Semiconductor lasers have smaller emitting areas than surface-emitting LEDs, and deliver more light into fibers. VCSELs have light-emitting areas tens of micrometers across, which couple well to graded-index multimode fibers. Edge-emitting diode lasers have emitting areas only a few micrometers across, which match well with single-mode fiber cores. One drawback of edge emitters is that their small emitting area gives them a relatively large beam divergence, so some of the light that enters the core may leak into the cladding. Both VCSELs and edge emitters can transfer a milliwatt or more into a fiber.

FIBER CHOICE

Another factor affecting light collection is the fiber. A fiber core diameter larger than the emitting area will decrease losses, although increasing core diameter further gives no benefit.

If the emitting area is larger than both fiber cores, you can estimate the difference in efficiency of light collection by a pair of fibers—fiber 1 and fiber 2—using the equation

●

Fibers with larger cores or NAs collect more light.

$$\Delta \text{Loss (dB)} = 20 \log_{10}\left(\frac{D_1}{D_2}\right) + 20 \log_{10}\left(\frac{\text{NA}_1}{\text{NA}_2}\right)$$

where the Ds are core diameters and the NAs are numerical apertures of the two fibers. You can use the formula as long as no optics are used to change the effective size of the emitting area or the effective NA of the source.

The difference can be significant with a large source. Consider, for example, the difference between a step-index fiber with 100-μm core and 0.3 NA and a graded-index fiber with 50-μm core and 0.2 NA. Substituting the numbers gives:

$$\Delta \text{Loss} = 20 \log\left(\frac{100}{50}\right) + 20 \log\left(\frac{0.3}{0.2}\right) = 9.6 \text{ dB}$$

That difference is nearly a factor of 10, well worth considering if you have run out of loss budget.

FIGURE 21.3

Laser output couples easily into a single-mode fiber core.

Light-emitting area is smaller than fiber core.

Single-Mode Fiber

Laser

On the other hand, remember that a larger core reduces transmission bandwidth as well as increasing light-collection efficiency. The sacrifices are largest in moving from single-mode to multimode graded-index fiber and from graded-index multimode to step-index multimode.

SINGLE-MODE FIBERS

• Single-mode fibers are used with laser sources.

Single-mode fibers normally are used only with laser sources. Edge-emitting lasers work best because their emitting area is smaller than that of a single-mode fiber core, as shown in Figure 21.3. VCSELs generally have larger light-emitting areas, so they typically are used with 50/125 or 62.5/125 µm graded-index multimode fibers. Fabricating a microlens on the fiber tip, or adding a microlens to the laser package, can enhance light coupling.

LEDs are a poor match to single-mode fibers. Carry the preceding example for multimode fibers a step further to a single-mode fiber with a nominal 10-µm core and NA of 0.11, and you will find it collects 19.2 dB less light from an LED than does a 50/125-µm fiber. Microlenses can reduce this excess loss, but in general single-mode fibers are not used with large-area LEDs.

Fiber Loss

• Fiber loss equals attenuation times distance, sometimes plus transient losses.

Fiber loss nominally equals the attenuation (in decibels per kilometer) times the transmission distance:

$$\text{Total loss} = (\text{dB/km}) \times \text{length}$$

However, this is only an approximation for multimode fibers. One problem is that measurements of fiber attenuation in dB/km do not consider transient losses that occur near the start of a multimode fiber. An LED with a large emitting area and high NA excites high-order modes that leak out as they travel along the fiber. Typically this transient loss is 1 to 1.5 dB, concentrated in the first few hundred meters of fiber following the transmitter. This loss becomes less significant after you go a kilometer or two, but graded-index multimode fibers are rarely used over much longer distances. Thus, it's important to remember transient loss and allow for it in your system margin.

An additional problem that can occur with graded-index fiber is uneven and unpredictable coupling of modes between adjacent lengths of fiber. These concatenation effects make loss of long lengths of spliced graded-index fiber difficult to calculate; fortunately, such systems are extremely rare.

Single-mode fibers are much better behaved because they carry only one mode, avoiding differential mode attenuation.

Fiber-to-Receiver Coupling

One of the rare places where the fiber-optic engineer wins is in coupling light from a fiber to a detector or receiver. The light-sensitive areas of most detectors are larger than most fiber cores, and their acceptance angles are larger than those of multimode fibers. Of course, if you're determined to screw things up, you can find a detector with a light-collecting area smaller than the core of large-core multimode fibers. But that isn't likely.

Losses are normally small in transferring light from fibers to receivers.

Receiver Sensitivity and Error Correction

The rest of this chapter assumes that receivers have a fixed sensitivity. That is, I assume that a system has to deliver a minimum power to the receiver to ensure proper function. Things aren't really that simple, however, because a number of factors determine receiver sensitivity.

As you learned in Chapter 11, receiver sensitivity as measured by power level drops as data rate increases because the bits are shorter at higher speeds, so fewer photons arrive during one bit interval. Thus, as the bit rate increases, the error rate increases if the power is held constant.

Forward error correction, described in Chapter 19, enhances receiver sensitivity by transmitting extra bits to detect and correct transmission errors. Forward error correction can also greatly improve system performance; but for simplicity we will ignore its effects and just assume the receiver requires a minimum input power to operate.

Forward error correction enhances receiver sensitivity.

Other Losses

Splices, connectors, and couplers contribute significant losses in a fiber-optic system. Fortunately, these losses generally are easy to measure and calculate. Each device has a specified loss that you can multiply by the number of devices to estimate total loss.

The amount of loss varies, particularly among connectors. A data sheet may specify a connector as having maximum loss of 1.0 dB and "typical" loss of 0.5 dB. The maximum loss is the specified upper limit for the device; none of these devices should have higher loss if installed properly. The typical value is nominally an average value over a large number of connectors. Usually the typical value is a conservative figure, so if you had 100 connectors with typical loss of 0.5 dB, the total might be 35 dB rather than the 50 dB you would expect. That makes it unlikely that anyone will wind up with a batch of 10 connectors with total loss of 7 dB, or 0.7 dB per connector, well above the "typical" loss.

The typical value generally is most important in a system containing multiple connectors, but sometimes the maximum may be important. Overestimating loss and overloading the receiver, as well as underestimating loss, can cause problems.

The accumulation of transient losses is a potential problem in multimode systems. Coupling light into a fiber from a light source or another length of fiber shifts some of the light into higher-order modes, which leak out along the length of the fiber. This can increase loss depending on the system configuration.

Total connector loss is the number of connectors times "typical" loss.

Margin

One quantity that always figures in the loss budget is system *margin*, a safety factor for system designers. This allows for uncertainties in calculating losses, for minor repairs, and for minor degradation of system components. Uncertainties are inevitable because component losses are specified within ranges and because components change as they age and are used. Margin also allows for repairs in case of cable damage, which typically add to cable loss.

Depending on the application, the performance requirements, the cost, and the ease of repair, the loss margin added by designers may be 3 to 10 dB.

Optical Amplifiers

Optical amplifiers boost signal strength.

Optical amplifiers can overcome losses by boosting signal strength, but practical concerns may offset this advantage. Optical amplifiers are expensive, and as analog devices inevitably amplify background noise as well as signal. On the other hand, they can amplify gigabit signals and multiple wavelengths in their operating ranges. Thus they are mainly used in high-performance and WDM systems, where their high cost can be spread among many signals being amplified.

You can use optical amplifiers in several places:

- As postamplifiers after transmitters, to generate high-power signals in fibers
- In the middle of long transmission systems, to boost signal strength for further transmission
- As preamplifiers before receivers, to raise signal strength to the proper level for the receiver
- Before or after couplers, which divide input signals among many outputs.

Examples of Loss Budgeting

To see how *loss budgeting* works, I'll step through three simple examples. Example A, shown in Figure 21.4, is a short system transmitting 100 Mbit/s between two points in a building. Example B, shown in Figure 21.5, is a telephone system carrying 2.5 Gbit/s between two switching offices 300 km apart. Example C, shown in Figure 21.6, is an intrabuilding network linking 10 terminals with each other at a signal speed of 100 Mbit/s. The examples are arbitrary and are intended to show how design works rather than to illustrate actual systems. Note that in considering only the loss budget, you don't directly address whether or not the system can carry the data rate listed. We'll look at that issue later in this chapter.

Connectors can contribute the dominant losses if several are in a short system.

Example A

In Figure 21.4, designers need to transmit signals through 200 m of fiber already installed in a building. That means that they must route the signal through patch panels with

FIGURE 21.4

Example A: Point-to-point link in a building.

connectors. In this example, they have six connector pairs, three on each floor: one linking the terminal device to the cable network for that floor, and one pair on each end of a short cable in the patch panel. (Connectors also attach the fiber to transmitter and receiver, but their losses are included under LED power transfer and receiver sensitivity.) The 50/125 graded-index multimode fiber has attenuation of 2.5 dB/km at the 850-nm wavelength of the LED transmitter. The loss budget is as follows:

LED power into fiber	−16.0 dBm
Connector pairs (6 @ 0.7 dB)	− 4.2 dB
Fiber loss (200 m @ 2.5 dB/km)	−0.5 dB
System margin	−10.0 dB
Required receiver sensitivity	−30.7 dBm

The calculation shows that the dominant loss is from the connectors. The fiber loss may underestimate transient loss, but the large system margin leaves plenty of room.

The calculated receiver sensitivity is a reasonable level, and system margin could be improved by picking a more sensitive receiver. This calculation started with a given loss, system margin, and input power, but you could start by specifying receiver sensitivity, system margin, and loss to calculate the needed input power. Note that the LED transmitter provides a low input power, but that is adequate for this short system.

Example B

Loss sources in the telephone system shown in Figure 21.5 are quite different. The system spans 300 km, with one splice every 10 km in a single-mode fiber with loss of 0.25 dB/km

Fiber attenuation dominates loss in a 300-km fiber system.

FIGURE 21.5

Example B: 300-km single-mode fiber system.

at 1550 nm. It links two rural areas, carrying a single wavelength at 2.5 Gbit/s. The high speed and long distance demand a laser source and a more sensitive receiver. Signals go through two connector pairs in patch panels at each end. Although the fiber loss is very low at 1550 nm, the long distance makes fiber attenuation the dominant loss. The sample calculation shows the laser transmitter alone does not deliver enough power to span that distance:

Laser power into single-mode fiber	0.0 dBm
Fiber loss (300 km × 0.25 dB/km)	−75.0 dB
Splice loss (29 × 0.1 dB)	−2.9 dB
Connector pairs (4 × 0.8 dB)	−3.2 dB
Power at receiver	−81.1 dBm
Receiver sensitivity	−32.0 dBm
System power deficit	−49.1dB

The output power falls far short of system requirements. You need optical amplifiers. Suppose you add a pair of optical amplifiers with 30-dB gain, one at the 100-km point and the second at 200 km. This requires four extra connector pairs (one on each end of each optical amplifier), which replace two splices. The loss budget then becomes:

Laser power into single-mode fiber	0.0 dBm
Fiber loss (300 km × 0.25 dB/km)	−75.0 dB
Optical amplifier gain	60.0 dB
Splice loss (27 × 0.1 dB)	−2.7 dB
Connector pairs (8 × 0.8 dB)	−6.4 dB
Power at receiver	−24.1 dBm
Receiver sensitivity	−32.0 dBm
System power margin	7.9 dB

That looks much better. To verify that the loss budget works for the whole system, you should check the budget for each segment.

Segment 1:

Laser power into single-mode fiber	0.0 dBm
Fiber loss (100 km × 0.25 dB/km)	−25.0 dB
Splice loss (9 × 0.1 dB)	−0.9 dB
Connector pairs (3 × 0.8 dB)	−2.4 dB
Power at optical amplifier	−28.3 dBm
Gain of optical amplifier	30.0 dB
Output of segment 1	1.7 dBm

Segment 2:

Optical amplifier into single-mode fiber	1.7 dBm
Fiber loss (100 km × 0.25 dB/km)	−25.0 dB
Splice loss (9 × 0.1 dB)	−0.9 dB
Connector pairs (2 × 0.8 dB)	−1.6 dB
Power at optical amplifier 2	−25.8 dBm
Gain of optical amplifier	30.0 dB
Output of segment 2	4.2 dBm

Segment 3:

Laser power into single-mode fiber	4.2 dBm
Fiber loss (100 km × 0.25 dB/m)	−25.0 dB
Splice loss (9 × 0.1 dB)	−0.9 dB
Connector pairs (3 × 0.8 dB)	−2.4 dB
Power at receiver	−24.1 dBm
Receiver sensitivity	−32.0 dBm
System power margin	7.9 dB

For our purposes, 7.9 dB seems an adequate power margin. In practice, the system margin probably will be better, because I have assumed a relatively high connector loss of 0.8 dB and a low laser output of 0 dBm.

Segment-by-segment calculations both check your result and make sure that placement of optical amplifiers doesn't get you into trouble. In practice, optical amplifiers

saturate at high powers, so you may get only 25 dB of gain with 20 dBm input. In this example, suppose you put the second optical amplifier at the 170-km point because you happen to have a storage building at that point. Then the calculations for segments 2 and 3 are as follows:

Segment 2:

Optical amplifier into single-mode fiber	1.7 dBm
Fiber loss (70 km × 0.25 dB/km)	−17.5 dB
Splice loss (6 × 0.1 dB)	−0.6 dB
Connector pairs (2 × 0.8 dB)	−1.6 dB
Power at optical amplifier 2	−18.0 dBm
Reduced gain of optical amplifier	25.0 dB
Output of segment 2	7.0 dBm

Segment 3:

Laser power into single-mode fiber	7.0 dBm
Fiber loss (130 km × 0.25 dB/m)	−32.5 dB
Splice loss (12 × 0.1 dB)	−1.2 dB
Connector pairs (3 × 0.8 dB)	−2.4 dB
Power at receiver	−29.1 dBm
Receiver sensitivity	−32.0 dBm
System power margin	2.9 dB

Although the receiver power is above the required level, a system margin of 2.9 dB is inadequate for contingencies. This is a reminder that you can't simply add up the losses and gains of all components without considering the input conditions to components such as optical amplifiers.

Example C

Coupling losses are largest in a network distributing signals to many terminals.

Complications also arise when you have to divide input signals among many outputs, as shown in Figure 21.6. In this case, you need to connect 10 terminals so the output of each one is divided among the receivers of all 10 terminals. You can do this with a 10 × 10 directional star coupler, which divides an input signal from any of the 10 incoming fibers (one from the transmitter end of each terminal) among the 10 output fibers going to the receiver end of each terminal. Assume the coupler divides the signals equally and has excess loss of 3 dB. Because the data rate is a modest 100 Mbit/s, let's calculate the loss budget with an LED source.

FIGURE 21.6

Example C: 10-terminal network.

LED transmitter (850 nm)	−16.0 dBm
Fiber loss (100 m @ 2.5 dB/km)	−0.25 dB
Connector pair loss (2 @ 0.5 dB)	−1.0 dB
Coupler loss (includes its own connections)	−13.0 dB
Power at receiver	−30.25 dBm
Receiver sensitivity	−30.0 dBm
System power deficit	−0.25 dB

The calculations show we are in trouble, with a negative system margin of –0.25 dB. The system might work, but we're right at the threshold of receiver sensitivity. As the system ages, performance will only get worse. We need to do something else.

A single component, the coupler, dominates the loss budget, but it is not easy to eliminate because it is needed to distribute the signal to all terminals. One way to overcome the power deficit is by replacing the LED with a more powerful VCSEL source. In this case, the power budget becomes

VCSEL transmitter (850 nm)	0.0 dBm
Fiber loss (100 m @ 2.5 dB/km)	−0.25 dB
Connector pair loss (2 @ 0.5 dB)	−1.0 dB
Coupler loss (includes its own connections)	−13.0 dB
Power at receiver	−14.25 dBm
Receiver sensitivity	−30.0 dBm
System power margin	15.75 dB

This may be too much of a good thing, depending on the power level that overloads the receiver. If the receiver overloads at −15 dBm, you can add a 3-dB attenuator to bring the receiver power down to −17.25 dBm, well within the operating range of −15 to −30 dBm. An alternative is to use a lower-power VCSEL, with output of −5 dBm. In this, the power budget becomes:

VCSEL transmitter	−5 dBm
Fiber loss (10 m @ 2.5 dB/km)	−0.25 dB
Connector pair loss (2 @ 0.5 dB)	−1.0 dB
Coupler loss (includes its own connectors)	−13.0 dB
Power at receiver	−19.25 dBm
Receiver sensitivity	−30.0 dBm
System power margin	10.75 dB

This should avoid overloading the receiver without the need for an attenuator.

Transmission Capacity Budget

Bandwidth or bit rate depends on fiber, source, and receiver characteristics.

The transmission capacity of a fiber-optic system is the total analog bandwidth or digital data rate it can carry. With wavelength-division multiplexing, total capacity of a fiber is the sum of the capacities of all optical channels the fiber carries. In this section, we will cover only the transmission capacity of a single optical channel; Chapter 22 will cover multichannel systems.

Single-channel capacity depends on how fast all the parts of the system respond to changes in signal intensity. In practice, transmission speed is mainly affected by properties of the transmitter, fiber, and receiver. For simplicity, I will ignore secondary effects such as noisy amplifiers and nonlinear effects that also can restrict data rates.

The simplest way to calculate transmission capacity is from the time response or rise time of the signal in the important components. This corresponds to the rise time of a transmitter or receiver, and dispersion in a fiber. As you will see below, this response is cumulative, and as you saw in Chapter 5, it is the sum of the squares of the responses of the various components.

Bandwidth and data rate can be calculated from rise time.

Both analog bandwidth and digital bit rate are related to the time response or rise time, although the relationships are approximate and depend on details.

For an analog system, the 3-dB bandwidth B in megahertz is inversely proportional to the rise time Δt in nanoseconds

$$B(\text{MHz}) = \frac{350}{\Delta t \, (\text{ns})}$$

Thus an analog system with 1-ns rise time has a roughly 350-MHz bandwidth.

For NRZ-coded digital signals, the rise time can be as large as 0.7 times the bit interval. If rise time is in nanoseconds, the maximum data rate in Gbit/s is

$$\text{NRZ data rate (Gbit/s)} = \frac{0.7}{\Delta t \, (\text{ns})}$$

Thus a digital system with 1-ns rise time can support an NRZ data rate of roughly 700 Mbit/s. RZ codes require double the bandwidth or rise times twice as fast:

$$\text{RZ data rate (Gbit/s)} = \frac{0.35}{\Delta t \text{ (ns)}}$$

These relationships are not exact, but they are useful for rough calculations.

The speed limits on digital signals are often easier to understand. You can get an intuitive feel for the impact of response time on a digital signal by thinking about pulse detection. If pulses follow each other at a certain speed, you need to be able to detect a pulse in less time than the interval between pulses.

Overall Time Response

The choice of time response or rise time simplifies calculations. The overall time response of a system is the square root of the sum of the squares of the response times of individual components:

$$\Delta t_{\text{overall}} = \sqrt{\sum (\Delta t_i^2)}$$

where $\Delta t_{\text{overall}}$ is the overall time response and Δt_i is the time response of each component.

Connectors, splices, couplers, and optical amplifiers do not affect time response significantly in current systems. The important response times are those of the transmitter, receiver, and fiber:

$$\Delta t_{\text{overall}} = \sqrt{\Delta t_{\text{transmitter}}^2 + \Delta t_{\text{receiver}}^2 + \Delta t_{\text{fiber}}^2}$$

That is, the overall response time is the square root of the sum of the squares of transmitter rise time, receiver rise time, and the pulse spreading caused by fiber dispersion.

Transmitter and receiver rise and fall times are listed on data sheets, ready to plug into the formula. Fiber response times must be calculated from the fiber's dispersion properties. You can see how transmitter and receiver properties affect data rate if we assume dispersion is small. Suppose we have a short data link spanning only 20 m of fiber, so we can neglect dispersion. If transmitter and receiver rise times are both 1 ns, the overall response time is:

$$\Delta t_{\text{overall}} = \sqrt{\Delta t_{\text{transmitter}}^2 + \Delta t_{\text{receiver}}^2} = \sqrt{2} = 1.414 \text{ ns}$$

Going back to the earlier formula for maximum data rate, we see it is:

$$\text{Data rate (Gbit/s)} = \frac{0.7}{1.4 \text{ (ns)}} = 0.5 \text{ Gbit/s}$$

If one component is much slower than the other, it dominates the response time. For example, if the transmitter had 1-ns rise time but the receiver had a 10-ns response, the overall response time would be 10.05 ns, limiting data rate to 70 Mbit/s.

Fiber dispersion becomes important in longer systems and deserves a closer look.

Fiber Dispersion Effects

Fiber response times must be calculated from the fiber length, the characteristic dispersion per unit length, and the source spectral width. As you learned earlier, fibers show modal, material, and polarization-mode dispersion. Which types are most important depends on

Overall time response is the square root of the sum of the squares of component rise times.

Fiber response must be calculated from dispersion.

the type of fiber. In multimode fibers, modal dispersion and chromatic dispersion are important. In single-mode fibers, modal dispersion is zero, but chromatic dispersion and polarization-mode dispersion are significant. (Remember that chromatic dispersion is the sum of material and waveguide dispersion.)

You estimate total pulse spreading caused by dispersion with a sum-of-squares formula similar to that for overall time response:

$$\Delta t = \sqrt{\Delta t_{modal}^2 + \Delta t_{chromatic}^2 + \Delta t_{PMD}^2}$$

In practice, this can be simplified. For multimode fibers, polarization-mode dispersion is insignificant, so only modal and chromatic dispersions are considered. Single-mode fibers do not suffer modal dispersion, so only chromatic and polarization-mode dispersions are considered. In practice, polarization-mode dispersion is negligible for data rates below 2.5 Gbit/s.

The equation can be converted to units of dispersion by substituting values of the pulse spreading for each type of dispersion. Modal dispersion is a characteristic value D_{modal} (specified in ns/km) times fiber length (in km). Chromatic dispersion is a characteristic value $D_{chromatic}$ (specified in ps/km-nm) times the fiber length (in km) and the spectral width of the transmitter (in nm). Polarization-mode dispersion is more complex because it varies randomly in time, but an average value is obtained by multiplying the nominal characteristic value D_{PMD} (specified in ps/root-km) by the square root of fiber length (in km).

Plugging these quantities into the pulse-spreading equation gives different formulas for multimode and single-mode fibers, based on fiber-dispersion properties, fiber length L, and spectral width of the light source $\Delta\lambda$. For multimode fiber, the pulse spreading is:

$$\Delta t_{multimode} = \sqrt{(D_{modal} \times L)^2 + (D_{chromatic} \times L \times \Delta\lambda)^2}$$

For single-mode fiber, this formula is:

$$\Delta t_{single\ mode} = \sqrt{(D_{chromatic} \times L \times \Delta\lambda)^2 + (D_{PMD} \times \sqrt{L})^2}$$

A couple of examples will show how dispersion calculations work.

MULTIMODE DISPERSION EXAMPLE

In Example A, considered earlier in this chapter, an 850-nm LED sends 100 Mbit/s through 200 m of 50/125-μm fiber. Modal bandwidth of a typical commercial 50/125-μm fiber is 400 MHz, which is equivalent to a modal dispersion of 2.5 ns/km. For a 200-m length, that corresponds to modal dispersion of 0.5 ns.

To that, you add the chromatic dispersion, calculated from the formula:

$$\Delta t_{chromatic} = D_{chromatic} \times L \times \Delta\lambda$$

A typical value of chromatic dispersion is 100 ps/nm · km at 850 nm, which combined with a linewidth of 50 nm for a typical 850-nm LED gives a chromatic dispersion of 1.0 ns for a 200-m length of fiber. This means that the chromatic dispersion is higher than modal dispersion because of the large LED spectral linewidth.

Adding modal and chromatic dispersions together according to the sum-of-squares formula indicates total dispersion is 1.1 ns. That leaves plenty of room for transmitting 100 Mbit/s, assuming the transmitter and receiver are fast enough.

If you wanted to transmit 1 Gbit/s, you could try using a VCSEL source with a 1-nm linewidth. In that case, total dispersion is

$$\Delta t = \sqrt{(0.5\ \text{ns})^2_{\text{modal}} + (100\ \text{ps/nm} \cdot \text{km} \times 0.2\ \text{km} \times 1\ \text{nm})^2} = 0.50\ \text{ns}$$

This is essentially equal to the modal dispersion of 0.5 ns and is adequate for gigabit transmission over 200 m.

SINGLE-MODE TRANSMISSION EXAMPLE

In Example B, we considered transmitting a 2.5-Gbit/s signal a total of 300 km through a single-mode fiber at 1550 nm. Assume that chromatic dispersion is specified at 3 ps/nm · km, a relatively low value at 1550 nm. Let's consider two cases: a Fabry-Perot laser with linewidth of 1 nm and an externally modulated distributed-feedback laser with linewidth of 0.1 nm. For a first approximation, ignore polarization-mode dispersion.

For the 1-nm laser, chromatic dispersion is

$$\Delta t_{\text{chromatic}} = 3\ \text{ps/nm} \cdot \text{km} \times 300\ \text{km} \times 1\ \text{nm} = 900\ \text{ps}$$

This value is much too high for a 2.5-Gbit/s system.

With the distributed-feedback laser, chromatic dispersion is

$$\Delta t_{\text{chromatic}} = 3\ \text{ps/nm} \cdot \text{km} \times 300\ \text{km} \times 0.1\ \text{nm} = 90\ \text{ps}$$

Thus using fiber with low chromatic dispersion near 1550 nm leaves plenty of margin. However, you could not get away with using step-index single-mode fiber with dispersion around 20 ps/nm-km at 1550 nm.

$$\Delta t_{\text{chromatic}} = 20\ \text{ps/nm-km} \times 300\ \text{km} \times 0.1\ \text{m} = 600\ \text{ps}$$

We also need to consider polarization-mode dispersion. A typical value for new nonzero dispersion-shifted fiber is about 0.5 ps/root-km, so for 300 km of single-mode fiber, the average polarization-mode dispersion is

$$\Delta t_{\text{PMD}} = 0.5\ \text{ps/}\sqrt{\text{km}} \times \sqrt{300\ \text{km}} = 8.7\ \text{ps}$$

This is only a tenth as large as chromatic dispersion, so it makes a negligible contribution to total dispersion, and does not affect transmission at 2.5 Gbit/s.

To fully assess performance of the system, you need to consider transmitter and receiver rise time as well. Suppose both have rise times of 100 ps, reasonable for products designed to work at 2.5 Gbit/s. Then the total pulse spreading is

$$\Delta t = \sqrt{100^2 + 100^2 + 90^2} = 168\ \text{ps}$$

Our earlier formula shows that this is adequate for transmitting 4 Gbit/s.

LONG-DISTANCE SINGLE-MODE EXAMPLE

Long-distance systems must span longer distances of several hundred to several thousand kilometers. Dispersion poses more problems in these systems because the total pulse

Dispersion poses more problems at long distances.

spreading increases with the length of the system. The seriousness of dispersion effects also increases with data rate.

Suppose you are using nonzero dispersion-shifted fiber with a chromatic dispersion of 3 ps/km-nm and PMD of 0.5 ps/root-km to span a distance of 1000 km. As in the last example, you are using an externally modulated laser with spectral width of 0.1 nm. Plug in the numbers and you find:

$$\Delta t_{chromatic} = 3 \text{ ps/nm} \cdot \text{km} \times 1000 \text{ km} \times 0.1 \text{ nm} = 300 \text{ ps}$$
$$\Delta t_{PMD} = 0.5 \text{ ps/}\sqrt{\text{km}} \times \sqrt{1000 \text{ km}} = 15.8 \text{ ps}$$
$$\Delta t_{total} = \sqrt{300^2 + 15.8^2} = 300.4 \text{ ps}$$

This response time is too long to allow good transmission at 2.5 Gbit/s, according to the criteria described earlier. This means that some form of dispersion compensation is needed to reduce average chromatic dispersion along the length of the fiber so it can carry 2.5 Gbit/s.

Things get even more difficult at 10 Gbit/s. If you flip the formula for maximum NRZ data rate to show the maximum allowable pulse spreading for a given data rate, you get

$$\Delta t_{maximum} = \frac{0.7}{\text{Data rate}}$$

For a 10 Gbit/s signal, this gives a maximum pulse spreading of 70 ps, which requires considerably more dispersion compensation. Let's assume that the transmitter bandwidth remains 0.1 nm, although high-speed modulation can increase this number. Suppose we can reduce the average chromatic dispersion by a factor of 5, to 0.6 ps/nm·km. Using the same assumptions above gives us:

$$\Delta t_{chromatic} = 0.6 \text{ ps/nm} \cdot \text{km} \times 1000 \text{ km} \times 0.1 \text{ nm} = 60 \text{ ps}$$
$$\Delta t_{PMD} = 0.5 \text{ ps/}\sqrt{\text{km}} \times \sqrt{1000 \text{ km}} = 15.8 \text{ ps}$$
$$\Delta t_{total} = \sqrt{60^2 + 15.8^2} = 62 \text{ ps}$$

That figure is adequate if you only consider fiber response, but you also have to consider transmitter and receiver rise times. If they both equal 25 ps—very fast devices—that would bring the total pulse spreading to the 70 ps limit. You might want average chromatic dispersion reduced even further, to about 0.3 ps/nm · km.

So far we have assumed polarization-mode dispersion is fairly benign by using specifications for new fibers fresh from the factory. However, older fibers or fibers installed under less than ideal conditions may have higher polarization-mode dispersion. Suppose we want to use older fibers with a realistic value of 2 ps/root-km for polarization-mode dispersion. Those calculations give us:

$$\Delta t_{chromatic} = 0.6 \text{ ps/nm} \cdot \text{km} \times 1000 \text{ km} \times 0.1 \text{ nm} = 60 \text{ ps}$$
$$\Delta t_{PMD} = 2 \text{ ps/}\sqrt{\text{km}} \times \sqrt{1000 \text{ km}} = 63.2 \text{ ps}$$
$$\Delta t_{total} = \sqrt{60^2 + 63.2^2} = 87 \text{ ps}$$

This produces polarization-mode dispersion slightly larger than the chromatic dispersion, and yields a total pulse spreading too high for 10 Gbit/s transmission, even ignoring transmitter and receiver rise times.
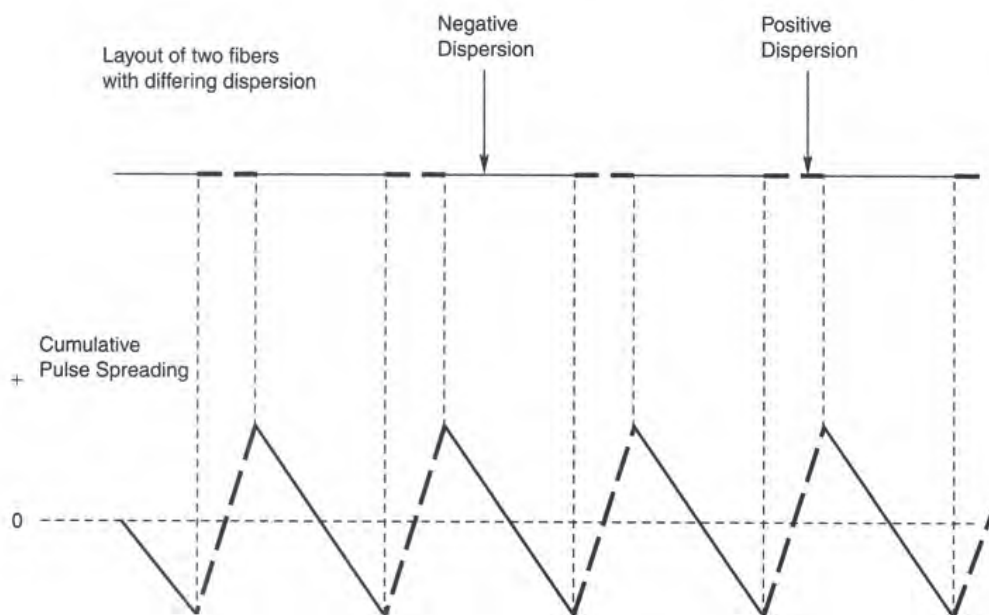
Dispersion Compensation

As these examples show, dispersion compensation becomes essential as transmission speeds and distances increase.

Compensation is easiest for chromatic dispersion. As you learned in Chapter 5, chromatic dispersion has a characteristic sign that indicates whether the shorter or longer wavelengths have gone farther in the fiber. To compensate for chromatic dispersion, you add a length of fiber (or some other optical component) with dispersion of the opposite sign. If the shorter wavelengths are falling behind, adding a length of fiber that makes the longer wavelengths fall behind serves to compress the length of the output pulse.

One approach to chromatic dispersion compensation is by alternating segments of fiber with a different dispersion sign, as shown in Figure 21.7. In this example, one fiber has a positive dispersion and the other a negative dispersion that is smaller in magnitude. This mixture is possible by using two types of nonzero dispersion-shifted fibers, one with zero dispersion at wavelengths shorter than the 1550-nm band, the other with zero dispersion at longer wavelengths. Optical signals pass through alternating segments of the two fibers, so the cumulative chromatic dispersion is first negative, then shifts positive, and so on. This is called *distributed compensation*, because the fibers that compensate for the pulse spreading are distributed along the fiber segment. In this example, the dispersions of the two fibers do not differ greatly in magnitude, but do differ in sign.

An alternative is to insert dispersion-compensation modules at selected points along the length of the system. Each module has chromatic dispersion opposite in sign to that of a certain length of the transmission fiber, so one module might compensate for 20 km of

> Dispersion compensation becomes important at high speeds or over long distances.

FIGURE 21.7

Dispersion compensation distributed along length of system.

dispersion. Modules can be built with fiber Bragg gratings that provide dispersion compensation. However, they more typically consist of easy-to-install coiled fibers designed specifically for dispersion compensation, which have chromatic dispersion higher in magnitude but opposite in sign to that of the transmission fiber. Chromatic dispersion in compensating fibers typically is several times larger than it is in transmission fibers, so shorter lengths are required. That's an advantage because compensating fibers generally have poorer transmission.

It is relatively simple to compensate chromatic dispersion at one wavelength to give a net chromatic dispersion of zero along a fiber segment, as shown in Figure 21.7. It is virtually impossible to compensate for chromatic dispersion across a wide range of wavelengths, as you will learn in Chapter 22.

Compensation for polarization-mode dispersion is very difficult. Unlike chromatic dispersion, PMD is a dynamic effect, which means its degree varies randomly over time. Thus PMD compensation also must be dynamic, with automatic adjustment over time. Both optical and electronic techniques have been demonstrated in the lab, but they are not yet widely accepted or in practical use.

Transmitter and Receiver Response Times

Transmitters and receivers must be matched to fiber characteristics.

So far I have concentrated on fiber dispersion, but transmitter and receiver response times also play critical roles in system bandwidth budgets. Just because a transmitter and receiver are rated to operate at 10 Gbit/s in some situations does not mean they can transmit at that speed in *all* systems. For example, a transmitter and receiver, both having 40 ps rise times, can transmit 10 Gbit/s signals through fiber with up to 40 ps of cumulative pulse spreading. However, they couldn't be used in the example I showed earlier, where total pulse spreading in the fiber was 62 ps. That would require a faster transmitter and receiver, with response times around 25 ps.

For this reason manufacturers sell different models of transmitters and receivers for transmission at the same data rates through different types of fibers. Transmitter–receiver pairs intended for short-distance use do not have to meet the same stringent speed requirements as those used for long-distance systems. In short, you have to match transmitters and receivers to the fiber system used to assure you get the desired transmission speed.

To show you how time-response calculations work, I have chosen simple examples. In reality there are a few other complications that come from the nature of transmitters and receivers. In particular, the range of wavelengths from transmitters is broadened by a couple of distinct effects.

Transmitter Spectral Broadening

As you learned earlier, directly modulating a semiconductor laser changes its refractive index as the density of current carriers changes. This effect, called *chirp*, causes the resonant wavelength to shift during a pulse, effectively spreading the range of output wavelengths. A laser that emits a bandwidth of 0.001 nm in a continuous beam spans a much larger range when directly modulated. External modulation can avoid chirp.

FIGURE 21.8

Modulation broadening is caused by side bands generated by the modulating signal.

External modulation cannot prevent a second type of spectral broadening that arises from any amplitude modulation of a pure carrier signal. The same effect occurs with radio. Modulation generates new frequency components by adding (and subtracting) the frequency of the modulating signal to the carrier frequency, as shown in Figure 21.8. This process generates *side bands* at frequencies both above and below the carrier.

The two side bands contain identical information, so radio transmitters generally suppress one side band to conserve scarce frequency space. Side-band suppression is difficult at optical wavelengths, so normally both side bands are present. The result is an effective broadening of the transmitter spectrum that increases directly with the bandwidth. If modulation produces a 10-GHz range of frequencies, it generates 10-GHz side bands on each side for a total spectral range of 20 GHz. This affects both chromatic dispersion in the fiber and the channel spacing possible in DWDM systems.

● Modulation broadens the wavelengths in an optical carrier.

Cost/Performance Trade-offs

So far I have only mentioned in passing a critical concern in real-world system design—cost. Minimizing cost is an implicit goal in all system design. I list some simple guidelines below, but it is impossible to give hard-and-fast rules for the tough job of making trade-offs between cost and performance. Ultimately it is your judgement as a system buyer or designer whether pushing error rate from 10^{-9} to 10^{-12} is worth an extra $1000. The best I can give you are some ideas to apply in working situations.

● Users must make cost/performance trade-offs.

Choice of Fiber Type

The choice of fiber type has a tremendous impact on the cost and performance of any system. A fundamental choice is between single- and multimode fiber, but several variations on both types are available. Installation is expensive, so you want to allow room for your system to expand. Bandwidth requirements inevitably increase, just as each new generation of software demands more computer memory and hard-disk space.

You need single-mode fiber if your system spans more than a couple of kilometers. Premium types, preferred for amplified WDM systems, are nonzero dispersion-shifted fibers, with low dispersion through the 1550-nm erbium amplifier window. Some single-mode fibers are optimized for "metro" applications over distances to a few hundred kilometers; others are optimized for terrestrial long-haul or submarine systems. Step-index single-mode fiber, with zero dispersion near 1310 nm, is still common, but normally is used at shorter distances because of its high dispersion at 1550 nm. Low-water fibers have a broader transmission window where amplification is not critical.

Single-mode fiber sometimes is used at distances shorter than a couple of kilometers, especially for high-speed transmission, such as 10-Gigabit Ethernet.

Graded-index multimode fibers are normally preferred for networks where transmission distances are up to a couple of kilometers, speeds are moderate, and many connections are required. Their big advantage is easier coupling to each other and to low-cost light sources. Bandwidth is significantly higher for 50/125-μm fiber than for 62.5/125-μm fiber. Remember that transmission speeds are sure to increase, so plan for future capacity expansion.

Plastic fibers and large-core step-index multimode fibers have very limited distance ranges. Plastic fibers have high attenuation, and large-core multimode fibers have low bandwidth. Their applications remain limited, but they can be valuable in certain situations, such as short data links inside equipment.

Other Guidelines

It's impossible to give a comprehensive set of guidelines for fiber system design, but I can give a set of rough-and-ready suggestions, starting with a few common-sense rules:

● Your time is valuable. If you spend an entire day trying to save $5 on hardware, the result will be a net loss unless you are mass-producing the design.

● Installation, assembly, operation, and support are not free. For a surprising number of fiber-optic systems, installation and maintenance cost more than the hardware. You may save money in the long term by paying extra for hardware that is easier to install and service.

● It can cost less to pay an expert to do it than to learn how yourself. Unless you need to practice installing connectors, it's much easier to buy connectorized cables or hire a fiber-optic contractor for your first fiber-optic system.

● You can save money by using standard mass-produced components rather than developing special-purpose components optimized for a particular application.

Some basic cost trade-offs are common in designing fiber-optic systems.

● The performance of low-loss fiber, high-sensitivity detectors, and powerful transmitters must be balanced against price advantages of lower-performance devices.

● Low-loss, high-bandwidth fibers generally accept less light than higher-loss, lower-bandwidth fibers. Over short distances, you can save money and overall attenuation by using a higher-loss, more costly cable that collects light more efficiently from lower-cost LEDs. (Because of the economics of production and material requirements, large-core multimode fibers are considerably more expensive than step-index single-mode fibers.)

● The marginal costs of adding extra fibers to a cable are modest and much cheaper than installing a second parallel cable. However, if reliability is important, the extra cost of a second cable on a different route is a worthwhile insurance premium.

● LEDs are much cheaper and require less environmental protection than lasers, but they produce much less power and are harder to couple to small-core fibers. Their broad range of wavelengths and their limited modulation speed limit system bandwidth.

● Fiber attenuation contributes less to losses of short systems than do losses in transferring light into and between fibers.

● Topology of multiterminal networks can have a large impact on system requirements and cost because of their differences in component requirements. Coupler losses may severely restrict options in some designs.

● Light sources and detectors for 1300 and 1550 nm cost more than those for the 650- or 800- to 900-nm windows, although fiber and cable for the longer wavelength may be less expensive.

● 1550-nm light sources cost more than 1300-nm sources.

● Fiber and cable become a larger fraction of total cost—and have more impact on performance—the longer the system.

● Balance the advantages of eliminating extra components with the higher costs of the components needed to eliminate them. For example, it's hard to justify the expense of two-way transmission through a single fiber over short distances.

● Optical amplifiers or high-power laser transmitters make sense in systems distributing signals to many terminals.

● Compare costs of high-speed TDM on single channels or lower-speed TDM at multiple wavelengths.

● *Dark fibers*—extra fibers installed in the original cable that were never hooked up to light sources—are often available in existing cables.

● Remember that human actions—not defective equipment—cause most fiber-optic failures. Consider ring or mesh topologies that can survive a single break. Take the extra time and spend the extra money to make important systems less vulnerable to damage. This means labeling and documenting the system carefully, as well as not leaving cables where people can trip over them.

Always leave room for future upgrades. Bandwidth requirements are sure to increase, and it costs less overall to install higher-capacity fiber now than to install a cheap one that you have to replace with a more expensive one later.

- Install extra fibers in your cables; they're a lot cheaper than going back later to install more cables.
- Leave margin for repair, such as slack in cables. It costs much less than complete replacement later.
- Leave room for expansion by adding WDM to your system.
- Watch for potential bottlenecks that might prevent future expansion.

As you grow more familiar with fiber optics, you will develop your own guidelines based on your own experience.

What Have You Learned?

1. Design of fiber-optic systems requires balancing sometimes-conflicting performance goals as well as costs.

2. The system loss budget is calculated by subtracting all system losses from the transmitter output power plus the gain of any optical amplifiers. The resulting output power should equal the input power required by the receiver plus system margin.

3. Significant losses can occur in coupling light from sources into fibers. You need multimode fibers to collect light from LED sources and single-mode fibers to collect light from edge-emitting laser sources. Large-core fibers are more efficient for large-area LEDs.

4. Total fiber loss equals attenuation (dB/km) multiplied by transmission distance. Multimode fibers may suffer transient losses in the first 100 to 200 m.

5. Total loss from connectors, couplers, and splices is their characteristic loss multiplied by the number of each in the system. You calculate the most likely loss using average loss and the worst case using maximum specified loss.

6. System margin is a safety factor that allows for repairs and aging of components. Typical values are 5 to 10 dB.

7. Optical amplifiers boost signal strength, but because of their high cost they are best used in long, high-speed systems or systems that distribute signals from one source to many receivers.

8. Transmission capacity budgets calculate bandwidth or bit rate; they depend only on source, fiber, and receiver characteristics. You can estimate capacity by calculating response time.

9. Response time of a system is the square root of the sum of the squares of component response times. Calculations must include transmitter and receiver response times as well as fiber dispersion.

10. Modal dispersion and chromatic dispersion combine to limit capacity of multimode fibers. Because of these capacity limits, multimode fibers are rarely used over more than a couple of kilometers.

11. Chromatic dispersion and polarization-mode dispersion limit capacity of single-mode fibers, which usually transmit over a kilometer or more. Chromatic dispersion depends on source spectral width as well as fiber dispersion.

12. Dispersion compensation becomes important at high speeds or over long distances. The compensating elements can be fibers distributed through the length of the system, or lumped at certain points.

13. Transmitters and receivers must be matched to fiber characteristics.

14. Modulation broadens the range of wavelengths in the optical carrier. The largest effect is chirp for directly modulated lasers, but external modulation also broadens transmitter spectral range.

15. Installation can cost much more than hardware. With demand for transmission capacity rising steadily, you should keep your upgrade paths open.

What's Next?

Chapter 22 covers the fundamentals of optical networking design, including wavelength-division multiplexing.

Further Reading

Vivek Alwayn, *Optical Network Design and Implementation* (Cisco Press, 2004)

Gerd Keiser, *Optical Fiber Communications*, 3rd ed. (McGraw-Hill, 2000)

Rajiv Ramaswami and Kumar N. Sivarajan, *Optical Networks: A Practical Perspective* (Morgan Kaufmann, 2002)

Questions to Think About

1. Why is the relative light-collection efficiency of fibers in decibels proportional to 20 times, rather than 10 times, the log of the ratio of their diameters?

2. Suppose the amplifiers in a transatlantic cable are limited to 12 dB of gain to limit noise. How far apart can they be spaced if the fiber attenuation averages 0.24 dB? You can neglect splices, and the system contains no connectors.

3. You are installing a fiber-optic data link between a laboratory and a remote data-collection center 5 km away. You want to use a VCSEL transmitter with −5 dBm output and a fiber with loss of 2.5 dB/km at 850 nm. The system includes patch panels on each end, with two connector pairs at each patch panel, and additional connectors on the terminal equipment. If connector loss is 0.5 dB, and you want a 10 dB margin, how sensitive must your receiver be?

4. You have to design a system with 1 Gbit/s data rate using return-to-zero (RZ) digital coding. What is the 3-dB analog bandwidth of this system? What NRZ data rate could this system transmit?

5. You want to transmit 1 Gbit/s through 100 km of single-mode fiber, using a transmitter and receiver that each have response times of 0.4 ns. The transmitter has line width of 0.1 nm. Neglecting polarization mode dispersion, what is the maximum chromatic dispersion allowable in the fiber?

6. Can you get away with using a 1550 nm VCSEL in the system of Problem 5 if the VCSEL has linewidth of 0.5 nm, output of −5 dBm, and the fiber loss is 0.25 dBm? (This assumes you can find such a VCSEL.) Check both the pulse spreading and the power level, assuming receiver sensitivity of −30 dBm.

Chapter Quiz

1. A large-area LED transfers 10 μW (10 dBμ) into an optical fiber with core diameter of 100 μm and numerical aperture of 0.30. What power should it couple into a fiber with 50 μm core and NA of 0.2?

    a. 10 dBμ

    b. 9.5 dBμ

    c. 3 dBμ

    d. 1.0 dBμ

    e. 0.4 dBμ

2. A connector is specified as having loss of 0.6 dB ± 0.2 dB. What is the maximum connector loss in a system containing five such connector pairs?

    a. 0.6 dB

    b. 3.0 dB

    c. 4.0 dB

    d. 5.0 dB

    e. none of the above

3. A 100-Mbit/s signal must be sent through a 100-m length of fiber with eight connector pairs to a receiver with sensitivity of −30 dBm. The fiber loss is 4 dB/km, and the average connector loss is 1.0 dB. If the system margin is 5 dB, what is the minimum power that the light source must couple into the fiber?

    a. −13.0 dBm

    b. −13.4 dBm

    c. −16.0 dBm

    d. −16.6 dBm

    e. −20.0 dBm

4. A system is designed to transmit 622 Mbit/s through 50 km of cable with attenuation of 0.4 dB/km. The system contains two connector pairs with 1.5 dB

loss, a laser source that couples 0 dBm into the fiber, and a receiver with sensitivity of −34 dBm. How many splices with average loss of 0.15 dB can the system contain if the system margin must be at least 8 dB?

a. none

b. 10

c. 20

d. 30

e. 40

f. none of the above

5. A 2.5-Gbit/system must span a distance of 2000 km, with optical amplifiers every 80 km. If the fiber loss is 0.3 dB/km at 1550 nm and there is one 0.1 dB splice every 16 km, what must the amplifier gain be if the system is not to gain or lose signal strength across its entire length?

a. 20 dB

b. 24.4 dB

c. 26.4 dB

d. 30 dB

e. 34.4 dB

6. You need to transmit identical 1-Gbit/s signals to 200 homes using a 1310-nm laser source. The homes are 1 to 4 km from your transmitter and use receivers sensitive to 30 dBm. What transmitter power do you need to achieve a 5-dB system margin if your fiber has a 0.4-dB/km loss at 1310 nm, each signal path from transmitter to home includes 6 connectors with a 0.5-dB average loss, and you split the signal in a 1 × 200 tree coupler with no excess loss?

a. 4.6 dBm

b. 9.6 dBm

c. 2.6 dBm

d. 0.0 dBm

e. −0.4 dBm

7. What is the duration of a single-bit interval in a 1.7-Gbit/s signal?

a. 1.7 ns

b. 1 ns

c. 0.588 ns

d. 0.294 ns

f. 0.170 ns

8. What is the response time of a system with transmitter response of 2 ns, receiver response of 1 ns, and 100 m of multimode fiber with dispersion of 20 ns/km (including both modal and chromatic dispersions)?

a. 2 ns

b. 2.236 ns

c. 2.646 ns

d. 2.828 ns

e. 3 ns

9. What is the total dispersion of 10 km of graded-index fiber with modal dispersion of 2.5 ns/km and chromatic dispersion of 100 ps/nm · km when it is used with an 850-nm LED having a 50-nm spectral width?

a. 5 ns

b. 25 ns

c. 50 ns

d. 55.9 ns

e. 75 ns

10. What is the total dispersion of 10 km of single-mode fiber with chromatic dispersion of 17 ps/nm · km and average polarization-mode dispersion of 0.5 ps/$\sqrt{km}$ at 1550 nm when used with a laser source with spectral width of 1 nm?

a. 1.58 ps

b. 10 ps

c. 17 ps

d. 170 ps

e. 172 ps

11. You are designing a 1-Gbit/s system using NRZ-coded signals with a transmitter with 0.3 ns rise time and a receiver that also has 0.3 ns rise time. What is the maximum total dispersion allowable through the entire length of the fiber?

a. 0.1 ns

b. 0.3 ns

c. 0.44 ns

d. 0.56 ns

e. 0.7 ns

12. You generate a 2.5-Gbit/s NRZ signal with rise time of 0.15 ns and spectral width of 0.1 nm. You have to send it through 80 km of nonzero dispersion-shifted fiber with chromatic dispersion of 6 ps/km·nm at the transmitter wavelength. Average polarization-mode dispersion is 0.5 ps/root-km. What is the total pulse dispersion in the fiber?

a. 4.5 ps

b. 45 ps

c. 48 ps

d. 52.5 ps

e. 80 ps

Optical Networking System Design

About This Chapter

Optical networking adds another dimension to the concepts of system design you learned in Chapter 21 for the simple case of a single optical channel per fiber. This chapter covers systems that transmit two or more optical channels per fiber, where signals are managed by wavelength or optical channel. The extra optical channels increase the complexity of system design, and the number of factors that must be considered.

This chapter opens with a review of optical networking concepts. It then explains how optical channels are packed together, contrasting wavelength-division and time-division multiplexing, and dense and coarse channel spacing. Then it covers the properties of optical fibers and optical amplifiers that affect optical networking design. Finally it covers optical switching and channel management, including the importance of wavelength conversion. Optical network design is still a young field, so we will not cover it in as much detail as we did single-channel design in Chapter 21.

Optical Networking Concepts

Optical networking organizes signals by wavelength as well as by the time sequence of digital data. Wavelength-division multiplexing packs a number of optical channels into a transmitting fiber, with each optical channel transmitted at a different wavelength. In an ideal optical network, signals can readily be converted to different wavelengths to rearrange or redistribute them, as cars shift between lanes on a highway. WDM began as a way to squeeze more data through an optical fiber, but is evolving into a new way of managing data by wavelength as well as by digital coding.

There are several key concepts in optical networking.

● *Signal granularity* measures how signals are subdivided within an optical network. The greater the granularity, the more potential ways there are to organize the signals. In general granularity is a good thing because data transmitted over most networks is assembled from many small data streams, not massive flows between two points. Think of the traffic as being like many automobiles, not a few 200-car freight trains.

Granularity is the subdivision of signals in a network.

● *Total transmission capacity* of a system measures how much information a fiber can transmit. It equals the sum of the data rates on all the individual optical channels carried by the system. To maximize transmission capacity, you pack channels as closely together as possible and transmit the highest possible data rate on each channel. Raising the data rate increases the bandwidth required for each optical channel, so trade-offs are inevitable. Cost-performance trade-offs also are inevitable because not all fiber-optic routes require the greatest possible bandwidth, and packing high-speed channels tightly together can be very expensive.

● *Fiber transmission capacity* depends on attenuation and dispersion, which are functions of wavelength. Variations in attenuation and dispersion limit the transmission capacity of certain parts of the spectrum more than others. The degree of limitation depends on overall transmission distance.

● *Amplification capacity* also varies across the spectrum. Good amplifiers are not available at all wavelengths, and amplifiers do not have uniform gain across their operating ranges. Amplification is a must for long-distance optical networks, but may not be required in other types.

● *Switching capacity* for optical networking depends on the ability to manipulate signal wavelength. Wavelength conversion technology is still in development.

Two broadly different families of optical networking technologies have been developed for different types of applications, and further variations may emerge in coming years.

DWDM is used for long-haul systems; CWDM is used for shorter links.

● *Dense-WDM* packs as much transmission capacity as possible into as few fibers as possible. Its main applications are in spanning distances of hundreds of kilometers or more, where infrastructure costs are relatively high. Sharing the capacity over more channels reduces overall costs significantly, even if the cost per channel is high.

● *Coarse-WDM* reduces equipment costs for WDM systems that span moderate distances—tens of kilometers or less. Infrastructure costs over these distances are relatively low, so it's important to limit the cost per channel.

These concepts are the cornerstones of optical network design. Let's look at how they are used.

Optical Channel Density

WDM divides a block of spectrum among optical channels.

A WDM optical network divides a block of spectrum among multiple optical channels. The space required for an optical channel depends on its data rate; the higher the data rate, the broader the bandwidth the channel requires. For example, the signal produced

by modulating a single-line light source at 10 Gbit/s occupies four times more bandwidth than does an identical light source modulated at 2.5 Gbit/s. Allocating the spectrum among optical channels is a fundamental step in optical network design.

WDM Compared to High-Speed TDM

Before we slice up the optical spectrum, we should compare WDM to high-speed time-division multiplexing. In principle, a single 160-Gbit/s data stream could carry 16 10-Gbit/s signals or 64 2.5-Gbit/s signals. Why not time-division multiplex the slower signals to higher speeds and avoid the need for so many separate transmitters at different wavelengths?

Part of the answer is that high data rates impose limitations on fiber transmission. Higher-speed TDM signals run into transmission distance limits; increase the data rate by a factor of four, and the maximum transmission distance drops by a factor of 16 because of chromatic dispersion. The pulse durations are four times shorter, and modulation effects broaden the transmitter spectrum by a factor of four, multiplying pulse spreading caused by chromatic dispersion by a second factor of four. A signal that can travel through 1600 km of fiber at 2.5 Gbit/s can travel only 100 km at 10 Gbit/s, and a mere 6.25 km at 40 Gbit/s. Combine 16 2.5-Gbit/s signals into one signal at 40 Gbit/s and dispersion limits it to 6.25 km; but transmit 16 separate 2.5-Gbit/s signals at different wavelengths and they can all travel 1600 km. Current technology performs well at 10 Gbit/s, but transmission at 40 Gbit/s is difficult, and transmission at 160 Gbit/s is extremely difficult, even in the laboratory.

Another part of the answer is granularity. There's more need for separate 2.5- or 10-Gbit/s channels than for 40-Gbit/s channels. Internet traffic maps mostly show 2.5-Gbit/s links, with 10-Gbit/s transmission only on the highest-traffic routes between the biggest cities. The companies that provide transmission capacity usually subdivide the capacity of their fibers into sizes that their customers want. WDM enables them to split fiber capacity into 2.5- and 10-Gbit/s slices, providing the granularity needed for today's traffic. Customers can transmit independent traffic on each wavelength, and often can choose the transmission format.

Overall, WDM offers better transmission distance and granularity than high-speed TDM for today's technology and transmission demands.

## Spectral Range and Optical Channels

Only a limited spectral range is usable in any fiber system, depending on factors such as attenuation, dispersion, and amplification that will be described later. The available spectrum is divided among multiple optical channel slots, which normally are of an identical width set by international standards. For example, a range of 3200 GHz, corresponding to 1562.42 to 1535.82 nm (192.0 to 195.2 THz), can be divided into 32 100-GHz channels, 64 50-GHz channels, or 16 200-GHz channels.

The space that each channel requires depends on the modulation rate. As shown in Figure 22.1, a 10-Gbit/s signal spreads the modulation spectrum of a carrier signal across a broader range than does a 2.5-Gbit/s signal, so it can't fit in as narrow an optical channel. The degree of separation possible depends on the optics as well as the modulation bandwidth.

Fiber transmission distance is limited at high data rates.

FIGURE 22.1

Faster signals spread across more spectrum and require wider slots.

Channel Spacing

10 Gbit/s

High-speed channels have high crosstalk if spaced too closely.

2.5 Gbit/s

● 10-Gbit/s signals usually go in 100-GHz channels in practical systems.

Advanced technology can reduce the modulation bandwidth and sharpen the selectivity of the optics to improve channel spacing. However, limits on total transmission capacity are unavoidable. So far the best results obtained in the laboratory without using elaborate polarization schemes have squeezed about 0.8 bit per second into 1 Hz of bandwidth. This figure of merit, called *spectral efficiency*, corresponds to fitting a 40-Gbit/s signal into a 50-GHz optical channel. Spectral efficiencies of practical commercial systems are much lower; they typically assign 10-Gbit/s signals to 100-GHz channels and 2.5-Gbit/s signals to 50-GHz slots.

Populating Channels

● Not all wavelength slots are populated.

One confusing dichotomy in optical networking is the often vast difference between the actual operating load of a fiber-optic system and its stated capacity. The difference arises because WDM systems are designed with slots to accommodate a certain number of optical channels, but carriers do not immediately *populate* all these channels with operating transmitters and receivers. Also optical systems often are modular, so only part of the hardware is installed initially, with the remaining optics added as transmitters and receivers are installed on new channels.

Figure 22.2 shows an example of this dichotomy, a WDM system that has potential slots for 40 optical channels. Its *potential capacity* is the number of available slots times the

FIGURE 22.2

Partial provisioning of a 40-channel DWDM system.

Only three receivers and transmitters actually installed.

maximum data rate per channel. The carrier operating the system, however, has installed transmitters and receivers on only three of the optical channels; all the others, shown in dashed lines, are reserved for future expansion. In this example, the carrier has bought optics that separate input signals into five groups of optical channels, plus optics that divide one of those five groups of optical channels into eight separate channels. Only three of those eight channels are populated, and the transmitters and receivers may not be operating at the maximum rate. Thus this system could be carrying only 7.5 Gbit/s (2.5 Gbit/s on each of the three populated channels), but have a potential capacity of 400 Gbit/s (10 Gbit/s on each of 40 channels).

This example is typical of the incremental approach that carriers take to installing transmission capacity. Operating companies do not need all the potential capacity immediately, but they do want to have room for future expansion. Transmitters and receivers currently are expensive, but their prices are coming down, so the carrier populates only the channel slots that are needed immediately. If more capacity is needed later, the carrier can populate more slots with cheaper transmitters and receivers, so the carrier buys only the capacity it needs today.

In fact, transmission loads vary, and during the telecommunications bubble many carriers vastly overestimated the maximum transmission capacity they would need. The result was that the actual peak load of a system like the one in Figure 22.2 would fall far short of its 7.5-Gbit/s capacity.

Dense- and Coarse-WDM

So far we have focused on the practice of packing optical channels closely together to squeeze as much bandwidth as possible through a single optical fiber. This is called *dense-WDM* or *DWDM*. In practice, DWDM systems have channel spacing of 200 GHz or less. They

Coarse-WDM can reduce transmission costs.

provide huge transmission capacity over long distances, and during the telecommunications bubble they filled the tremendous demand for bandwidth. Today most of these systems are underutilized.

No one seriously disputes the advantages of transmitting signals through one fiber on many separate optical channels. However, they do question the advantages of DWDM, which was developed specifically to provide high-speed, long-distance transmission. Long-haul applications could justify the high costs of precision optics to split the spectrum into narrow slices and of cooled high-speed laser transmitters to provide precise wavelengths that fit into the narrow channel slots. However, those high costs prevented DWDM technology from being used in other applications that could benefit from wavelength-division multiplexing.

An alternative approach, called *coarse-WDM* or *CWDM*, was developed to avoid the high costs associated with precise wavelength control. It divides the spectrum into fewer slices, wide enough to accommodate the wide wavelength variations of much cheaper uncooled diode lasers. Each CWDM channel can carry many DWDM channels, as shown in Figure 22.3. Two variations of CWDM have been developed: one for dividing a single high-speed signal into four slower data streams, and the other for sending up to 18 separate signals through a single fiber without amplification.

● 10-Gigabit Ethernet uses four-channel CWDM.

The 10-Gigabit Ethernet standard has an option for splitting the data stream into four separate signals transmitted at 2.5 Gbit/s, for a total of 10 Gbit/s. This allows the use of graded-index multimode fibers over longer distances than those possible at 10 Gbit/s and the use of lower-cost 2.5-Gbit/s transmitters with single-mode fiber. The coarse-WDM transmitters use uncooled distributed-feedback lasers emitting at center wavelengths of 1275.7, 1300.2, 1324.7, and 1349.2 nm. The broad range of wavelengths is possible because no optical amplifiers are needed for the distances used in 10-Gigabit Ethernet. Loose wavelength tolerances and using uncooled lasers lower costs. The lasers could operate at temperatures of 0 to 70°C with wavelength drifting no more than 5 nm, which would keep the signals in the proper coarse-WDM slots.

FIGURE 22.3

CWDM and DWDM channel widths.

FIGURE 22.4

CWDM channels specified in ITU G694.2 standard.

The ITU G694.2 standard established a grid of 18 CWDM slots with center wavelengths from 1270 to 1610 nm for telecommunications. As shown in Figure 22.4, this range spans the band of lowest attenuation in silica fibers. A few channels lie under the 1.38-μm water peak, but these are easily made available by using low-water fibers. Each of these channels is 20 nm wide. This is broad enough to accommodate the wavelength variations of an uncooled laser transmitter, which drifts about 9 nm when temperature changes by 100°C. The specification allows use of interference filters with central passbands limited to the central 13 nm, which can be made quite inexpensively with insertion loss less than 1 dB. You can see this spectral profile in Figure 22.4.

The major applications for CWDM are expected to be in metropolitan networks, which span no more than tens of kilometers and generally do not require amplifiers. Typically these networks operate at data rates to 2.5 Gbit/s, allowing the use of inexpensive directly modulated DFB laser transmitters. All channels except those in the E-band water peak can be transmitted through existing standard step-index single-mode fiber.

The CWDM standard is designed to be used by itself, but blocks of DWDM channels can be slipped into single 20-nm CWDM channels. Up to 16 DWDM channels can fit into a CWDM channel at 1550 nm, and up to 22 into a CWDM channel near 1310 nm. A variety of alternatives are possible, such as using CWDM channels at 1470, 1490, 1590, and 1610 nm to supplement DWDM channels in the 1530–1565-nm erbium-amplifier band.

Overall, CWDM and DWDM each have their own advantages. CWDM is best suited for transmission over moderate distances, where amplification is not required. DWDM is best suited for longer distances, where amplification is needed, because the closely packed wavelengths all fit in the erbium-amplifier range.

> CWDM has a grid of 18 channels that are each 20 nm wide.

Operating Ranges of WDM Systems

Most WDM technologies can span much of the optical and infrared spectrum. WDM systems have been built using visible LEDs and plastic fibers, as well as the infrared sources and silica fibers described above. However, both the properties of the fiber and the amplifier technology used (if any) limit the operating range of any individual WDM system.

Gain Bandwidths of Optical Amplifiers

Optical amplifier gain limits the usable wavelength range.

WDM systems that require optical amplification are inherently limited by the gain bandwidths of the amplifiers. As you learned in Chapter 12, optical amplification depends on stimulated emission, and any material produces stimulated emission over only a limited range of wavelengths. The properties of optical amplifiers vary considerably and are summarized in Table 22.1. In practice, individual amplifiers are limited to usable gain bandwidths of tens of nanometers because WDM systems require uniform gain across their operating ranges; you don't want gain of 5 dB at one band and 35 dB at another.

Doped-fiber amplifiers are limited by the gain bandwidth of the light-emitting elements, praseodymium, thulium, and erbium. Erbium-doped fiber is the most widely used, and has usable gain over the broadest range of wavelengths, from 1530 to 1610 nm. As you learned in Chapter 12, the gain of erbium amplifiers varies widely over that range. They usually are designed specifically for either the C-band at 1530 to 1565 nm or the L-band at about 1570 to 1610 nm. The gain profile can change somewhat with fiber composition, but in general the bandwidth over which gain is reasonably uniform is tens of nanometers.

Raman gain is offset 13 THz from the pump in silica fibers.

Raman gain is offset from the pump wavelength by an amount that depends on the fiber composition. For standard silica fibers, the peak is 13 THz lower in frequency than the pump and has a bandwidth of about 5 THz, as you saw in Figure 12.14. Raman gain can be used to amplify light anywhere in the silica fiber transmission band by picking a pump band offset by the proper amount from the amplified wavelength. Pumping at multiple wavelengths can produce gain across a wide range, but if the range is too large the pumps overlap signal wavelengths, causing interference.

Semiconductor optical amplifiers have a spectral width of a few tens of nanometers, with the peak gain depending on the composition of the active layer. This means it is possible to make optical amplifiers with different central wavelengths by growing active layers with different compositions, but most of these devices are not standard commercial products.

Table 22.1 Optical amplifier properties

Type of Amplifier	Wavelength Range (nm)	Band Limits
Praseodymium-doped	1290–1320	Praseodymium stimulated emission
Thulium-doped	1450–1500	Thulium stimulated emission
Erbium-doped	1530–1610	Erbium stimulated emission
Raman	Broad, offset from pump band by 13 THz	Tens of nm gain band
Semiconductor	Broad, but composition dependent	Tens of nm from one composition

The gain bandwidths of amplifiers are broad enough to accommodate many DWDM channels, but only one or two CWDM channels. For example, the 35-nm wide gain of a C-band erbium fiber amplifier can hold about 40 channels with 100-GHz spacing or about 80 channels with 50-GHz spacing. However, as Figure 22.4 shows, it holds only one CWDM channel, with a second at the edge of the erbium band at 1530 nm. This illustrates why DWDM is used in long high-capacity systems that require optical amplification.

Technically, optical amplifiers are available throughout the 1270 to 1610 nm range of CWDM. However, most of those wavelengths are not widely available from commercial products, and amplification of the whole CWDM range would require several parallel optical amplifiers, a very cumbersome arrangement. In practice, CWDM systems that span the whole possible range are not amplified.

Fiber Bandwidth Limits

Fiber attenuation also can limit the wavelengths usable for optical networking, but the limit is much less severe than that imposed by the gain bandwidth of amplifiers. Fiber loss is below about 0.5 dB/km from about 1250 to 1650 nm, except near the 1380-nm water peak in high-water fibers, as shown in Figure 22.4. That entire band is usable in low-water fibers, which are promoted for CWDM applications because they transmit well at 1350 to 1430 nm. Fibers can carry signals at wavelengths outside this band, but large differences in attenuation accumulate over long distances and the resulting differences in power reaching the receiver or optical amplifier can cause problems.

Another fiber issue that can limit usable bandwidth is differences in chromatic dispersion with wavelength. Such differences discourage the use of widely separated wavelengths in long-distance systems. The limitations are much less important for distances under 100 km, where the band usable for CWDM is nearly 400 nm wide. It's possible in principle to pack hundreds of DWDM channels in the same bandwidth, but there's little demand for that much capacity in unamplified systems.

> Fiber loss is below 0.5 dB/km from 1250 to 1650 nm in low-water fibers.

Factors in WDM Design

You learned in Chapter 21 how system characteristics such as fiber attenuation and dispersion affect transmission at a single wavelength. Matters become more complex in WDM optical networking because signals on different optical channels may experience different levels of attenuation, dispersion, and other properties. That makes the properties of the fiber, the amplifier, and other components important factors in WDM system design. We will concentrate on amplified systems because they experience the most severe constraints.

Fiber Attenuation

Figure 5.2 shows that fiber attenuation varies significantly across the low-loss window from 1250 to 1650 nm. Loss in the 1310-nm region is about 0.35 dB/km compared to a minimum below 0.2 dB/km in the erbium amplifier band. This difference means that 100 km of fiber has 35 dB loss at 1310 nm, but only 20 dB loss at 1550 nm, a 15-dB difference at

FIGURE 22.5

Attenuation of low-water AllWave fiber in erbium-fiber window. (Copyright Lucent Technologies Inc.)

AllWave™ Cabled Fiber

Fiber loss varies little across the erbium-amplifier band.

the receiver assuming that input powers were equal. This difference is large enough to impact the design of CWDM systems that span that range.

The difference is much smaller over the erbium-fiber band used for DWDM, as shown for a low-water fiber in Figure 22.5. For the fiber shown, attenuation ranges between 0.192 and 0.200 dB/km in the 1530- to 1610-nm band, rising to 0.205 dB/km at 1620 nm. In this case, the difference in loss for a 100-km length is small, 19.2 dB at 1570 nm compared to 20.0 dB at 1610 nm. The difference is a mere 0.8 dB, small enough to be smoothed out by techniques used to equalize amplifier gain across the erbium amplifier band.

Dispersion Slope and Compensation

You learned in Chapter 21 that dispersion compensation can reduce the total chromatic dispersion of a fiber link. Dispersion compensation is relatively straightforward at a single wavelength, but recall that chromatic dispersion also varies with wavelength. This makes it necessary to consider the *dispersion slope*—the change in chromatic dispersion as a function of wavelength—in designing WDM systems. As shown in Figure 22.6, the value of the dispersion slope differs between types of fiber. It is largest for large-effective-area fibers, around 0.086 ps/nm^2-km in the C-band erbium-fiber window. Fibers also can be designed with smaller cores and reduced dispersion slopes, about half that value. Conventional nonzero dispersion-shifted fibers have intermediate slopes. This variation of dispersion with wavelength complicates the task of dispersion compensation.

Dispersion compensation can't be perfect across a range of wavelengths.

In a single-wavelength system, chromatic dispersion can be compensated for simply by adding fibers or other components that have dispersion of the same magnitude but the opposite sign, to give net dispersion of zero at the transmission wavelength. However, as Figure 22.7 shows, the curves plotting the dispersions of the transmitting and compensating fibers are not exactly the inverse of each other. At the wavelength λ_2, the dispersion of

FIGURE 22.6

Chromatic dispersion of reduced-slope and large-effective-area fibers shows difference in dispersion slopes.

1 km of compensating fiber completely offsets the dispersion of 5 km of transmission fiber, leaving a net cumulative chromatic dispersion of zero. However, if you add the curves together at shorter wavelengths the net cumulative dispersion is slightly negative and at longer wavelengths it's slightly positive. Thus the dispersion compensation works perfectly for one wavelength, but not for the other wavelengths λ_1 and λ_3.

These differences accumulate with distance. If you plot the cumulative dispersion along the length of the system, you see the sawtooth pattern shown in Figure 22.8. Dispersion increases gradually as light passes through the transmission fiber, then drops sharply as the signal passes through the compensating fiber, which has negative dispersion. The pattern repeats exactly for λ_3, for which the positive and negative dispersion exactly offset each other. However, the peaks gradually rise for the other two wavelengths, λ_2 and λ_1, where there is a net positive dispersion. This spreading sawtooth pattern betrays the imperfection of dispersion compensation across a range of wavelengths.

Similar effects occur when you compensate for chromatic dispersion by combining lengths of two or more types of transmission fiber. Some submarine cables use two or more types of fiber in each span between optical amplifiers. The bulk of the span is non-zero dispersion-shifted fiber with a slightly negative dispersion—about 2 ps/nm-km at 1550 nm—and a zero-dispersion wavelength longer than about 1600 nm. The negative dispersion is offset by the positive dispersion of about 17 ps/nm-km of the step-index single-mode fiber. (Large-effective-area fiber may be used near optical amplifiers to reduce nonlinear effects, as we will see later.) Although great care is taken to assure the best possible dispersion compensation, the best is not perfect, so dispersion plots along the length of the cable show a spreading sawtooth curve similar to the one shown in Figure 22.8.