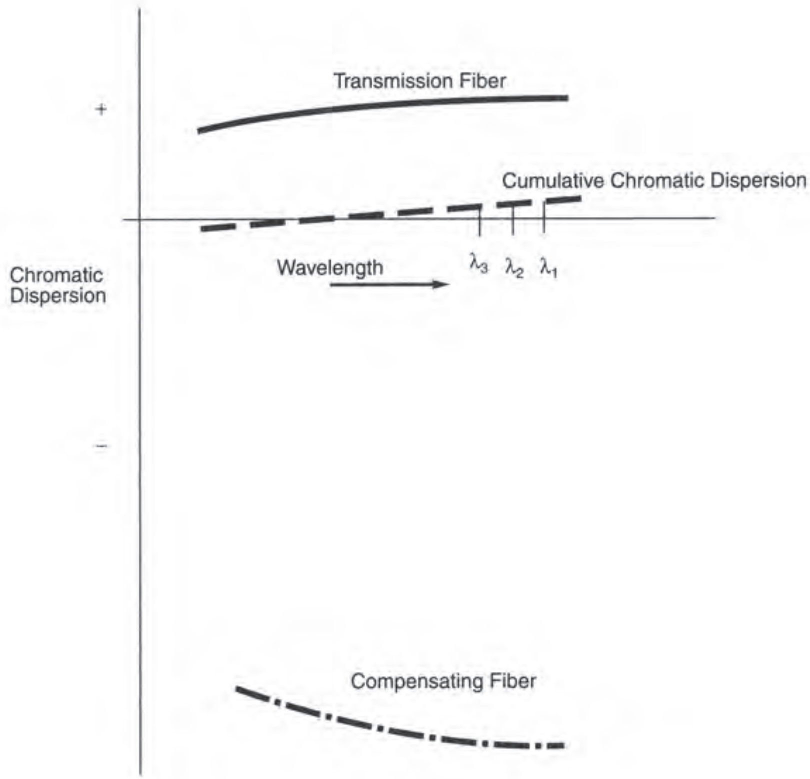

FIGURE 22.7
Dispersion compensation in WDM system, showing difference in cumulative dispersion.

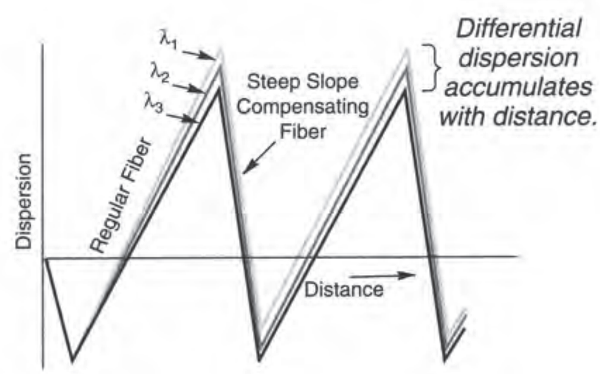

FIGURE 22.8
Cumulative dispersion plotted along a WDM system, which uses short lengths of dispersion-compensating fiber. Total dispersion at the three wavelengths diverges along the length of the fiber because the compensation inevitably is imperfect.

Nonlinear Effects in WDM Systems

Nonlinear effects are proportional to total optical power density and to the distance the light travels through the fiber at high power levels. Long-haul WDM systems are particularly vulnerable because the more optical channels they carry, the higher the total optical power. A single-channel system normally has no trouble transmitting 3 mW, but if a WDM system tries to transmit 80 channels at that power level, the total power reaches 240 mW, which can produce nonlinear effects and interactions among the transmitted channels.

High total power makes WDM systems vulnerable to nonlinear effects.

Nonlinear effects are relatively weak in glass fibers, but their total impact on the signal is proportional to the distance the signal travels in the fiber. In practice, they are not significant in metro systems that run tens of kilometers, but can pose problems in long-haul systems running thousands of kilometers.

Fiber attenuation complicates the picture by reducing the optical power as a signal travels through the fiber, causing nonlinear effects to decline with distance from the light source. As a result, nonlinear effects accumulate over a *maximum effective length*, after an amplifier or transmitter which depends on the fiber attenuation. For a typical single-mode fiber with 0.22 dB/km attenuation at 1550 nm, this is about 20 km per span between optical source and receiver. The value is smaller for fibers with higher attenuation. In long-haul systems, the maximum effective length is multiplied by the number of spans between amplifiers.

Attenuation limits nonlinear effects to a maximum effective length.

Four-wave mixing poses particular problems in DWDM systems. As you learned in Chapter 5, it occurs when signals at three input frequencies combine to generate a mixed signal at a fourth frequency:

$$\nu_1 + \nu_2 - \nu_3 = \nu_4$$

The three input signals need not all be at different frequencies; two of them could be on the same optical channel. The equal spacing of WDM channels means that the new frequency is likely to fall on another optical channel, producing noise and crosstalk.

Four-wave mixing increases if the three input waves remain in phase as they pass through the fiber. That occurs when there is no chromatic dispersion to spread them out along the fiber. If the fiber has some minimum chromatic dispersion in the transmission window—typically at least 1 ps/nm-km—the optical channels do not remain in phase over long distances, and the four-wave mixing signal is reduced, as shown in Figure 22.9. This level of dispersion is too low to limit transmission bandwidth on the individual channels.

Dispersion compensation does not enhance four-wave mixing because it uses lengths of fibers with positive and negative dispersion that combine to produce low dispersion over the entire fiber span. In such systems, the local dispersion is nonzero, so the signal pulses do not stay in phase over long distances. In short, near-zero *cumulative* dispersion over a fiber span is fine, but zero *local* dispersion in a single fiber can keep signals in phase and enhance four-wave mixing.

The dependence of nonlinear effects on power density makes them sensitive to the choice of fiber. Fibers with small effective areas concentrate signal power, making them more vulnerable to four-wave mixing and other nonlinearities. Typically, the effective area is large in step-index single-mode fibers, somewhat smaller in nonzero dispersion-shifted fibers, and smallest in dispersion-compensating fibers. Other trade-offs occur between effective area and reduced dispersion slope; increasing effective area tends to increase dispersion slope, and conversely, reducing the dispersion slope tends to decrease the effective area.

Fiber type affects nonlinear effects.

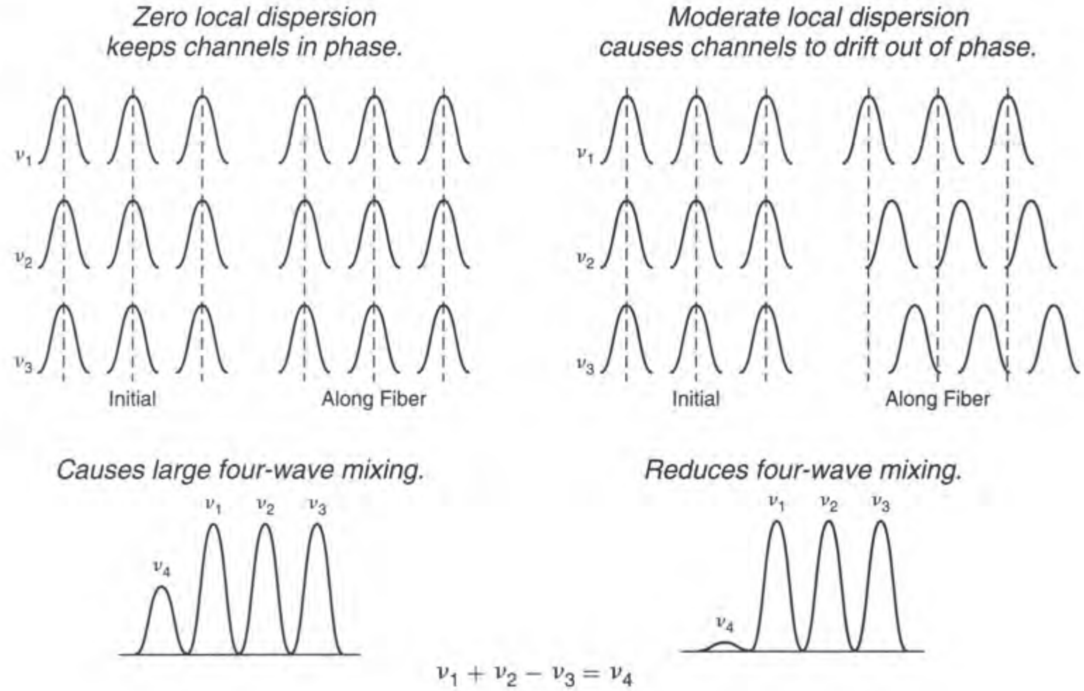

FIGURE 22.9 *Four-wave mixing is high at the zero-dispersion point, where signals stay in phase over long distances. Some local dispersion causes the wavelengths to drift out of phase, reducing four-wave mixing.*

Designers can take advantage of these differences in properties when they select fibers for dispersion compensation. Nonlinear effects can be reduced by placing a large-effective-area fiber close to transmitters or amplifiers, where the optical power is highest. Fibers more vulnerable to nonlinear effects can be placed in parts of the fiber span where power is lower.

Optical Amplification and WDM Design

The role of optical amplifiers is more complex in WDM systems than in single-channel systems. In single-channel systems, the variation of amplifier gain with wavelength limits the usable transmission band. In WDM systems, this variation adds requirements to balance gain across the transmission spectrum. The available optical power must be shared among the transmitted channels, and gain can be saturated when total input power reaches a high level, although power per channel remains modest. Amplified spontaneous emission is a concern because it adds to background noise across the gain spectrum.

These considerations apply to all types of optical amplifiers, but this section will concentrate on erbium-doped fiber amplifiers because they are the most common type in use.

Amplifier Power Levels

Amplifier output is divided among many WDM channels.

Erbium-fiber amplifiers have a maximum output power, typically 17 to 24 dBm for C-band amplifiers. The limit comes from energy transfer, both in exciting the erbium atoms and in stimulating emission, which causes the saturation effects described in Chapter 12. This peak output is concentrated on a single wavelength in a single-channel system, but in a WDM system it is divided among all populated optical channels, limiting the power per channel. As a result, the power per channel in a WDM system decreases with the number of channels. For example, an amplifier operated at its maximum output of 80 mW could deliver 20 mW on each of 4 optical channels, 10 mW on each of 8 channels, or 1 mW on each of 80 channels.

Multichannel operation can limit gain on individual channels to much lower values than are possible with a single channel. If an amplifier saturates at a power level of 15 dBm (30 mW), it can amplify a single −15 dBm input channel by 30 dB. However, if the input were 30 input channels each at −15 dBm (30 μW), the total output power still would be limited to 30 mW, or 1 mW per channel, a gain of only 15 dB. This has important system consequences because it means the fiber span between amplifiers can include only 15 dB of loss—equivalent to 75 km of fiber with 0.2 dB/km attenuation. Adding optical channels without upgrading the amplifiers can downgrade transmission capability. Depending on operational details, doubling the number of optical channels would reduce the power level on each channel by 3 dB, equivalent to 15 km of 0.2 dB/km fiber.

Amplifier power also affects maximum data rate per optical channel. As you learned in Chapter 11, pulse detection requires a minimum number of photons, so average powers must be increased to deliver those photons in shorter pulses. Thus dividing power among more optical channels, which can transmit more data in parallel, reduces the power available to transmit signals at a higher data rate per channel. In short, when operating near system margins, increasing the number of optical channels trades off directly with increasing the data rate per channel.

Gain Flatness and Channel Equalization

Filters can equalize amplifier gain on optical channels.

Erbium amplifier gain typically is flat to within 1 to 3 dB across its operating range. This sounds good until you start cascading amplifiers. That variation is acceptable for a single amplifier, but not in a cascade. A difference of 2 dB between channels becomes 10 dB after five amplifiers, which could lead to the loss of weaker channels. Gain-equalizing filters can compensate for this by reducing the power on the strongest optical channels, as shown in Figure 22.10. The filter's extra attenuation offsets the stronger gain at certain wavelengths, making amplifier gain flat across the spectrum. Some imperfections inevitably remain, but careful engineering can limit channel-to-channel power differences to no more than a few decibels over a chain of more than 100 amplifiers.

Another approach to gain equalization is by adding another amplifier to add power to the system rather than subtract it. Raman amplifiers are an attractive choice because their gain is higher at longer wavelengths than at shorter ones; they complement the spectrum of erbium amplifiers, which peak at shorter wavelengths. *Hybrid amplifiers* combine a Raman amplifier with an erbium-doped fiber amplifier to give more uniform gain across a range of wavelengths. The Raman amplification stage may be in the transmission fiber, but it requires a strong pump beam.

Switching and Optical Networking

Optical networking manages signals as optical channels.

A central concept of optical networking is managing signals by the optical channel, sometimes called a *lambda* after the Greek letter λ used as an abbreviation for wavelength. The idea is to transmit signals in many bit streams at separate wavelengths, such as 40 channels

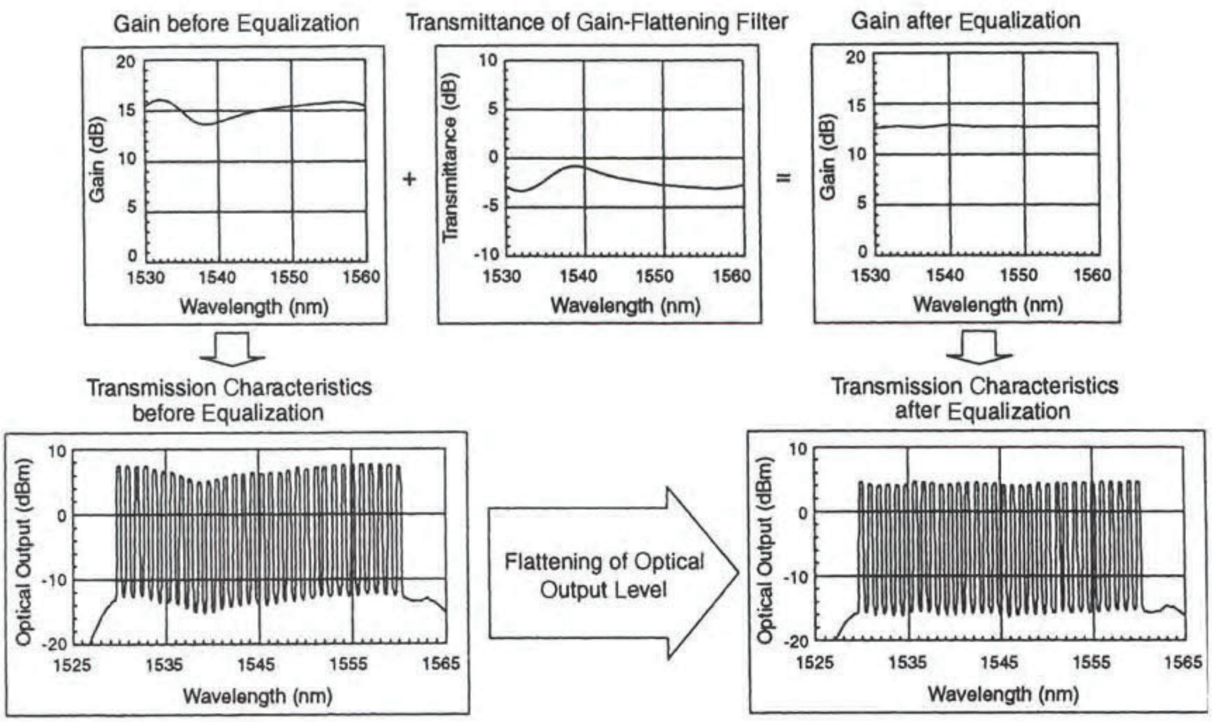

FIGURE 22.10
Effect of gain-flattening filter on an optical amplifier. (Courtesy of Furukawa Ltd.)

at 2.5 Gbit/s rather than a single 100-Gbit/s data stream. Multiple channels offer more *granularity* because each data stream can be detected and processed as an optical channel without disturbing any of the others.

Figure 22.11 shows how a simple optical network switch in Chicago can distribute eight optical channels from Omaha. First the input signal is demultiplexed to its eight component channels. Two wavelengths go to both Detroit and Indianapolis. One wavelength each goes to Minneapolis, Milwaukee, Chicago suburbs, and St. Louis. In this example, the input signals are organized in Omaha, and the Chicago switch directs them on their way without any further processing. The same fiber could carry many other wavelengths going to other destinations, but you couldn't make sense of a drawing with 40 optical channels.

This example looks simple because we've omitted crucial details. The optical switch is just a box; we don't know how it processes the light. The figure doesn't show the conversion of signals from one wavelength to another, although this may sometimes be necessary. To understand how optical networks are designed, let's take a closer look at some key concepts.

Transparent, All-Optical, and Opaque Systems

Both optical switches and networks can be classed into three categories that overlap to some extent.

- *Transparent* systems transmit optical signals without changing their format, as if the light were shining directly through them. They can be amplified, but their wavelength remains unchanged, and the optical signal is never converted into electronic form.

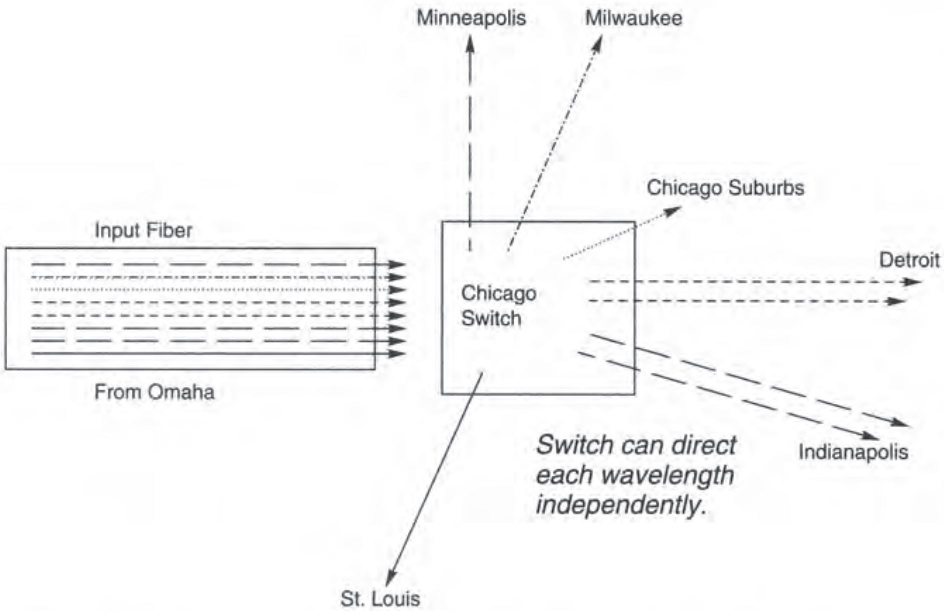

FIGURE 22.11 *Granularity of optical channels.*

- *All-optical* systems transmit and manipulate only optical signals, which are never converted to electronic form. However, the optical signals may be converted to different wavelengths. All transparent networks are all-optical, but some all-optical systems are not transparent.
- *Opaque* systems convert the input optical signal into electronic form for switching or processing, then convert that signal back into optical form. The output may be at the same wavelength or a different one.

A network does not have to be entirely of one type. An optical network may include "islands" of transparency that transmit signals purely optically, separated by opaque opto-electronic components such as receiver-transmitter pairs or electronic switches.

An optical network may have islands of transparency separated by opaque elements.

Present systems use both all-optical and electro-optical switches. The big switches at network nodes that direct signals among many input and output ports are electronic, converting input optical signals into electronic form for switching, then converting them back into optical form for transmission. Electronic switches take advantage of highly-developed optical techniques for processing signals and directing them among large numbers of possible output ports. All-optical switches generally are simpler devices used to transfer optical signals between fibers without additional processing. One family redirects all signals in the fiber to another fiber; they are used for protection switching. Another selects one or more wavelengths in a WDM signal and redirects only those wavelengths; they can be used as add/drop switches. The signal-distribution switch shown in Figure 22.11 could be implemented with this type of optical switch.

The advantages of transparent and all-optical networks have been much praised, but in practice some electronic components are much more mature than their optical counterparts. For example, wavelength conversion is simplest to implement by converting optical signals into electronic form and using the electronic signal to drive a laser transmitter at a different wavelength. As you learned in Chapter 12, it's also possible to convert wavelengths by purely optical means, but that technology is still in development.

Wavelength Conversion and Routing

Wavelength conversion is needed to manage optical channels.

Wavelength conversion is an important management tool in optical networks. From the network standpoint, the important part of a signal is the information it carries, not the wavelength of the carrier signal. The wavelength is like a lane on the highway, a channel for carrying information. When signals switch from one fiber to another, the wavelength used on the first fiber may not be available on the second, so the signal may have to be converted to a different wavelength. Ideally, the output wavelength should be tunable, so the signal can be switched to any desired wavelength. This ideal isn't easy to implement, but developers are working on it.

Another application for wavelength conversion is in conjunction with a device called a *wavelength router*, which directs input signals to different ports depending on their wavelength. A wavelength router can be viewed as a fixed wavelength-division demultiplexer, which always directs input signals out a particular port according to their wavelength. In the example of Figure 22.11, an input signal at 1542.14 nm might go to Milwaukee, one at 1542.94 nm might go to Minneapolis, and signals at 1543.73 and 1544.53 nm would go to Detroit. By changing the input wavelength, you can change the output port to which the signal is directed.

System Interfaces and Regeneration

The optical network terminates at electronic transmission centers.

The optical network interfaces with electronic signals at input and output ends. The input electronics typically organize the input electronic signals in some way, such as time-division multiplexing to combine signals for high-speed transmission, or regrouping signals from other inputs. Then the signals are converted to optical form, and optical interfaces combine signals at different wavelengths at the input of WDM systems.

Interfaces generally are at transmission centers, nodes, or hubs, where many transmission lines come together. Typically local signals feed into these nodes, which combine them with signals arriving from other points and transmit them on outgoing lines. Big electronic switches regenerate optical signals at major network nodes spaced several hundred kilometers apart.

Optical regeneration has been demonstrated in the laboratory, but has yet to find practical applications. System requirements for optical regeneration are not yet clear.

Design Examples

Designing WDM systems is a complex task that requires considering performance at many wavelengths and the interactions of the multiple signals passing through the system. To demonstrate how it works, we will concentrate on very simple examples that illustrate a few of the considerations.

Amplifier Gain and Power

Suppose we need to transmit 10 optical channels between two switching nodes 200 km apart. We have a single fiber, and space for an optical amplifier at the midpoint. The attenuation of each span is 25 dB, and the transmitter has 3-mW output on each of the 10 channels after the multiplexer. The company warehouse has optical amplifiers with performance similar to that shown in Figure 12.5. Small-signal gain is 30 dB and peak output is 30 mW. The demultiplexer has 5 dB loss per channel, and we haven't picked the receivers yet.

At first glance it looks like there's plenty of margin for operation on a single channel.

Single-channel transmission

Transmitter signal	5 dBm
Attenuation	–25 dB
Amplifier input	–20 dBm
Amplifier gain	27 dB
Amplifier output	7 dBm

Remember, however, that with nine other channels the total amplifier input is –10 dBm, which reduces the gain.

Multichannel transmission (per channel)

Transmitter signal	5 dBm
Attenuation	–25 dB
Amplifier input	–20 dBm
Amplifier gain	22 dB
Amplifier output	2 dBm
Attenuation (span)	–25 dB
Attenuation (demux)	–5 dB
Receiver power	–28 dBm

Receiver sensitivity of –35 dBm is required to achieve a 7-dB system margin.

Adding Optical Channels

Now let's consider what happens if we add another 10 optical channels to the system using the same amplifier and demultiplexer. The amplifier already produces 16 mW (10 channels at 1.6 mW each), so assume the extra input power reduces gain 2 dB.

20-channel transmission (per channel)

Transmitter signal	5 dBm
Attenuation	–25 dB
Amplifier input	–20 dBm
Amplifier gain	20 dB
Amplifier output	0 dBm
Attenuation (span)	–25 dB
Attenuation (demux)	–5 dB
Receiver power	–30 dBm

If we use the same receivers, the system margin is a rather thin 5 dB. This small margin is one reason why adding channels to WDM systems may require more extensive upgrades of optical amplifiers, transmitters, and receivers.

Gain Equalization

Suppose we have a chain of five amplifiers in a WDM system that transmits signals over a span of 600 km on a dozen optical channels. For proper performance, the receivers require signals between –30 dBm and –28 dBm. How much gain variation can be tolerated in each amplification stage, assuming that the variation is the same for each amplifier?

The allowable variation is 2 dB divided by 5 amplifiers, or 0.4 dB for each amplifier. Gain equalization is needed to keep the variation within these limits.

Gain equalization normally uses filters to block the additional power. If the amplifiers normally have 3 dB of gain variation across the spectrum, the filters will reduce the amplifier gain by 3 dB, so the gain budget will have to be increased. This is an example of the many trade-offs involved in WDM design.

What Have You Learned?

1. WDM is a fundamentally different way to organize signals than time-division multiplexing.
2. Granularity is the subdivision of signals in a network.
3. Total transmission capacity is the amount of information that can be transmitted by a network. It depends on data rate per optical channel, the number of channels, and channel spacing. It also depends on the range of wavelengths transmitted.
4. Optical amplifiers are not available for all fiber transmission wavelengths, limiting the bands usable for long-distance transmission. Fiber attenuation and dispersion limit the bands usable for shorter distances.
5. Dense-WDM is used for long-distance transmission. Coarse-WDM is used for tens of kilometers, and is less expensive to implement.
6. Dispersion limits fiber transmission distance at high data rates. WDM provides higher total transmission capacity on multiple channels with lower data rates.
7. Spectral efficiency measures how tightly signals can be packed, dividing the total data rate by the optical bandwidth used. The best rates achieved without elaborate polarization schemes are about 0.8 bit/s/Hz.
8. Network operators do not populate all available channel slots.
9. The ITU G694.2 standard established a grid of 18 CWDM slots with center wavelengths from 1270 to 1610 nm. This grid is used mainly in metro telecommunications networks.
10. Optical amplifiers have gain bandwidths of tens of nanometers; this limits the wavelength ranges available for long-distance DWDM systems.

11. Raman gain is offset from the pump wavelength by an amount that depends on the material; the value is 13 THz in silica. The Raman wavelength is longer (lower in frequency) than the pump wavelength.
12. Fiber loss varies significantly in the 1250 to 1650 nm; it is 0.35 dB/km at 1310 nm and under 0.2 dB/km at 1570 nm. Variation in the erbium amplifier band is much smaller.
13. Dispersion slope measures variation of chromatic dispersion with wavelength.
14. It's impossible to compensate for chromatic dispersion perfectly across a range of wavelengths because the dispersion curves of different fibers don't cancel each other out perfectly. These differences increase over long distances.
15. Nonlinear effects increase with total power and with the distance the light travels at high power. Attenuation reduces nonlinear effects far from the transmitter.
16. Saturation limits total power available from an optical amplifier. Saturation depends on total power, not just power at a single wavelength, so WDM systems can saturate.
17. Opaque systems convert optical signals into electronic form at some point, typically for switching or wavelength conversion.
18. Wavelength conversion is needed to manage optical channels. Wavelength routers direct signals according to their wavelength.

What's Next?

In Chapter 23 you will learn about the structure of the global telecommunications network.

Further Reading

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

Rajiv Ramaswami and Kumar N. Sivarajan, *Optical Networks: A Practical Perspective,* 2nd ed. (Morgan Kaufmann, San Francisco, 2002)

Questions to Think About

1. Your design packs 2.5-Gbit/s optical channels only 50 GHz apart, but 10-Gbit/s signals require 100 GHz spacing to give adequate performance margins. Which spacing transmits more data if you populate every available slot in the erbium-amplifier C-band? How much is the difference?
2. Wavelengths of 1460 to 1625 nm are available in an unamplified metro WDM system. How many CWDM optical channels can this network transmit with ITU standard 20-nm spacing?
3. You have two types of fiber available for dispersion compensation, step-index single-mode fiber with dispersion of +17 ps/nm-km at 1550 nm, and a reduced-slope non-zero dispersion-shifted fiber with –2 ps/nm-km at the same wavelength.

How much of each do you need for a 95-km span with zero cumulative dispersion at 1550 nm? Which type should you use closer to the transmitter?

4. An erbium-fiber amplifier that has peak output of 20 dBm transmits 40 optical channels. What is the output power on each channel? What happens to the output power if the number of channels is doubled?

5. You need to send 20 optical channels through a series of 50 erbium-doped fiber amplifiers, and the output signals must differ by no more than 6 dB at the end of the amplifier chain. How should you equalize the gain assuming all that difference is due to unequal gain across the range of optical channels?

6. The input to an all-optical switch is a cable containing 48 fibers. Each fiber can transmit up to 40 optical channels. How many optical channels can the switch direct if it has 128 input and 128 output ports?

Chapter Quiz

1. One 40-Gbit/s TDM channel is equivalent to how many 2.5-Gbit/s optical channels?
   a. 1
   b. 4
   c. 16
   d. 40
   e. 100

2. A WDM system has 40 optical channels spaced 100 GHz apart in the 1550-nm region. What is the approximate total spectrum used by the system?
   a. 40 GHz
   b. 100 GHz
   c. 400 GHz
   d. 4 THz
   e. 193.1 THz

3. How is the spacing required between optical channels related to the data rate transmitted on each channel?
   a. The spacing increases with the data rate; the required spacing in gigahertz is larger than the data rate in gigabits.
   b. The spacing equals the signal speed in gigahertz.
   c. The spacing in nanometers is equal to the data rate in gigabits per second.
   d. The spacing decreases with the data rate; it is proportional to the length of the data pulses.
   e. The relationship is impossible to state because data rate is digital and frequency spacing is analog.

4. Which of the following systems can accommodate the most optical channels?

a. coarse-WDM in a system that includes several erbium-doped fiber amplifiers

b. dense-WDM in a system that includes several erbium-doped fiber amplifiers

c. coarse-WDM in a system without optical amplifiers

d. dense-WDM in a system without optical amplifiers

e. not enough information to tell

5. You need to compensate dispersion for a 60-kilometer length of fiber with chromatic dispersion of −3 ps/nm-km. How much fiber with chromatic dispersion of +15 ps/nm-km at the same wavelength (1550 nm) do you need?

a. 5 km

b. 12 km

c. 15 km

d. 25 km

e. 60 km

6. Your fiber has a dispersion slope of 0.08 ps/nm^2-km in the C-band of erbium-fiber amplifiers. Assuming that slope is a straight line, how much does the dispersion change over the width of the C-band?

a. 0.08 ps/nm-km

b. 0.8 ps/nm-km

c. 2.8 ps/nm-km

d. 8 ps/nm-km

e. 28 ps/nm-km

7. A span between two fiber amplifiers includes three types of fiber. Which type should be closest to the output of the first amplifier to reduce four-wave mixing?

a. step-index single-mode fiber with effective area 80 μm^2 and dispersion +15 ps/nm-km

b. reduced-slope single-mode fiber with effective area 50 μm^2 and dispersion –2.5 ps/nm-km

c. large-effective-area single-mode fiber with effective area 68 μm^2 and dispersion –20 ps/nm-km

d. either of the low-dispersion fibers

e. any of the fibers

8. In which fiber is four-wave mixing the largest?

a. step-index single-mode fiber with effective area 80 μm^2 and dispersion +15 ps/nm-km

b. reduced-slope single-mode fiber with effective area 50 μm^2 and dispersion –2.5 ps/nm-km

c. large-effective-area single-mode fiber with effective area 68 μm^2 and dispersion –2.0 ps/nm-km

d. zero dispersion-shifted single-mode fiber with effective area 75 μm^2 and dispersion 0 ps/nm-km

e. 50/125 graded-index multimode fiber

9. An optical amplifier generates 2 mW on each of 40 optical channels. An optical switch downstream diverts half of the optical channels to another cable. If everything else remains constant, what is the power on the remaining channels?

a. 2 mW

b. 3 mW

c. 4 mW

d. 5 mW

e. The amplifier will not work.

10. What is the difference between a transparent and an opaque optical switch?

a. A transparent switch allows light signals to pass through in optical form; an opaque switch converts them to electronic form.

b. A transparent switch is made of clear glass; an opaque switch is made of metal.

c. A transparent switch does not change the data rate of the signal; an opaque switch changes the data rate.

d. A transparent switch transmits a signal; an opaque switch blocks the signal, turning it off.

11. You add eight new optical channels to an optical fiber that already was carrying eight optical channels. The signals must pass through an optical amplifier nearing saturation. All the signals arrive at the amplifier with the same power. What happens to the original eight channels?

a. Output power on all channels after the addition is the same as it was on the original channels because the power per channel is well under saturation.

b. Output power on the original channels is unchanged, but the new channels get less power.

c. Output power per channel drops to a lower level.

d. Output power on the original channels doubles.

e. All channels fail to transmit a signal because they overload the amplifier.

12. Which of the following varies least in the 1550-nm transmission window?

a. chromatic dispersion

b. fiber attenuation

c. erbium-amplifier gain

d. raman amplifier gain

CHAPTER

23

Global Telecommunications Applications

About This Chapter

Now that you have learned about fiber-optic hardware, standards, and system design, it's time to look at how fiber-optic systems are used. Changing technology and regulations are eroding traditional divisions, but it is still useful to separate telecommunications into a few sectors. The largest in scale is the global telecommunications network, the backbone of international telecommunications, including long-distance transmission under the oceans and on land. Other sectors are regional or metro networks and distribution networks for voice, video, and computer data services. You learned a little about these ideas in Chapter 3; now it's time to take a closer look.

In this chapter, I will describe the long-distance fiber-optic transmission systems that carry data, voice, video, and other signals around the world. They include intercontinental submarine cables as well as national and international systems on land. They are the world's biggest telecommunications "pipelines," and they are designed to maximize both transmission speed and distance. Fiber-optic technology has dominated these systems for over a decade, first with single-channel single-mode transmission, and now with high-speed DWDM systems.

These long-distance networks feed into regional or metro networks, which in turn connect switching offices and distribution networks, which deliver services to individual homes and offices. Later chapters will look at those networks.

Defining Telecommunications

The term *telecommunications* is deliberately broad. It dates back to the era when communication specialists were trying to group telephones and telegraphs under one heading. As the telegraph industry faded away, telephony become dominant, but the new word had caught on. It was useful because new communication services were emerging. Radio and television networks spread, broadcasting voice, music, and pictures. Telex relayed printed messages around the world. Facsimile systems transmitted images of documents. Computer data communications grew rapidly. Wireless telephones and pagers spread. They all fell under the broad heading of telecommunications.

Telecommunications encompasses voice, data, facsimile, video, and other forms of communications.

Different types of telecommunications had different origins, but most are *converging* toward a common network that delivers many different services. The main reason is that it costs less to build one versatile network than many specialized ones. Convergence will never be complete because some services inherently differ from others. A mobile phone small enough to fit into your pocket can display brief text messages, but not complex Web pages or large-screen broadcasts of a football game.

Convergence is strongest in the global backbone network that carries digitized signals around the world. Most signals go through two somewhat distinct fiber-optic networks: the traditional circuit-switched system that evolved from the telephone network and the Internet. Note that these are only *somewhat* distinct because they overlap and signals are often converted between the two forms.

Switching is the most obvious difference between the two networks. As you learned earlier, the Internet is packet-switched and traditional telephone networks are circuit-switched. This distinction is not a rigid one. The telephone networks can encode Internet Protocol packets into SONET frames and transmit them on circuit-switched fiber. As viewed from the OSI layered model, the system transmits the signal in IP format on the interchange layer and as SONET on the physical layer. Likewise, the Internet can transmit telephone calls as packet streams using Voice over Internet Protocol (VoIP).

Satellites transmit video signals, and some voice and data services.

Communication satellites provide other long-distance transmission. They relay digital and analog video signals to television broadcasters, cable television companies, and customers of direct-broadcast services. Satellites also provide paging services as well as some data transmission and telephone links, with the last two services used largely in remote areas. You should know that these services exist, but I won't cover them because they aren't part of the global fiber network.

The telecommunications industry is changing at the same time as the technology. The once-monolithic telephone industry has become a collection of many companies that form shifting alliances and tend to merge or divorce at the whims of Wall Street. Competition is real in the long-distance market, where many carriers have separate networks spanning part of—or in a few cases, most of—the world. Competition is more dubious on the local level. Mergers seem to change the name of your local telephone or cable-television company more often than most people change local carriers, but cable and telephone companies do compete with each other to offer data and—sometimes—telephone services. Everything has to interconnect for the system to function, and most of the time most of the network does.

Let's look briefly at the elements of the global telecommunications network before focusing on long-distance transmission.

The Telephone Network

The telephone network spread around the globe in the twentieth century, becoming the backbone of the international telecommunications network. The phone network was intended to carry conversations between any two telephones, so it has connections extended to individual home and office phones around the world.

The telephone network evolved to become the global telecommunications backbone.

Local and regional telephone systems interconnect with each other and with long-distance and international carriers to offer service around the block and around the planet. Telephone numbers provide the information needed for switching signals. In much of North America, you can direct calls within your area code by dialing seven digits. Long-distance calls within the United States, Canada, and parts of the Caribbean require dialing a long-distance code (1), a three-digit area code (XXX), and a seven-digit local number. (You must always dial the area code in places where two or more area codes are overlaid in the same area.) To make overseas calls, you dial an international code (011 from the United States), a 1- to 3-digit country code (e.g., 44 for Britain or 81 for Japan), usually a city code or other regional code (e.g., 207 for inner London or 3 for Tokyo), then a local number (usually 6 to 8 digits). Thanks to this system, you can call most of the phones in the world from your home or office, although you may regret it when you get the bill.

Each traditional twisted-wire-pair telephone line carries only a modest amount of information. A standard analog phone line carries sound frequencies of 300 to more than 3000 Hz, which the industry calls POTS, for Plain Old Telephone Service. (Phone lines can carry frequencies to 4000 Hz, but the upper frequencies are used for control signals.) That is enough for intelligible conversations, but it is far short of the ideal 20- to 20,000-Hz range of the human ear. Pulse-code modulation converts the analog signal to digital format, with one voice channel equal to 64,000 bit/s.

The land-line telephone system is so pervasive that all fax machines, mobile phones, pagers, and dial-up computer modems work with it and take advantage of its infrastructure. These devices are assigned standard telephone numbers, and their signals must be compatible with telephone lines. All share the same dialing system, as shown in Figure 23.1, so mobile phones can reach wire-line phones and vice versa. However, this also means that somebody on the opposite side of the world can send a fax to your home voice line in the middle of the night. If you pick up the receiver, it's like getting someone who doesn't speak English—both parties on the line are speaking in the audio spectrum, but in different languages.

Phones, faxes, and modems share a common dialing system.

The Internet and Computer Networks

The Internet began as a network linking the computer networks at major research universities and laboratories but has since expanded to link many of the world's computers. It generally carries digital data more efficiently than phone lines. As with the telephone network, you can divide the Internet into a long-distance backbone system and regional and local transmission systems.

The Internet carries digital data more efficiently than phone lines.

FIGURE 23.1
Voice, fax, and other signals share the telephone network.

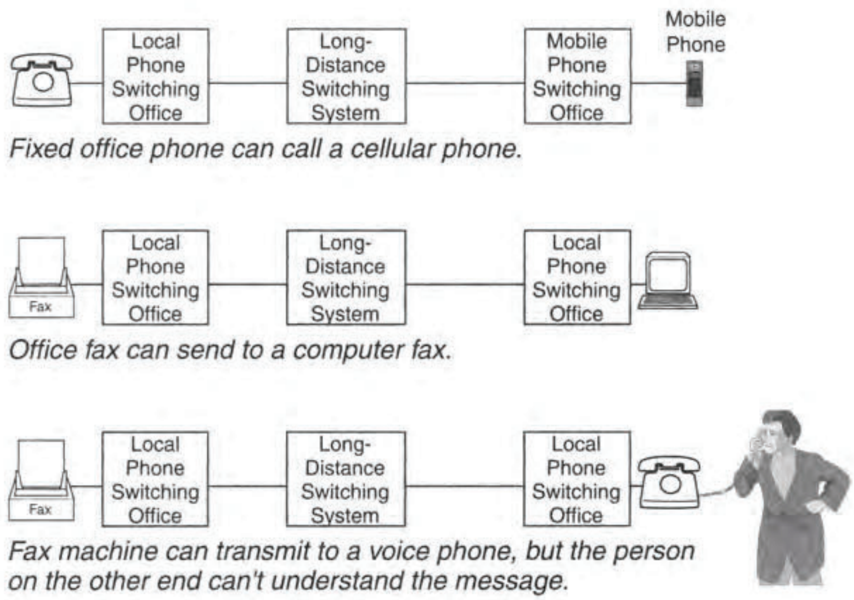

In practice, there is considerable overlap between computer networks and telecommunications networks. You can connect a personal computer to the following:

- A dial-up modem linked to an analog telephone line, which sends data to an Internet service provider over phone wires.
- A wireless modem linked to the cellular telephone network. A new generation of mobile phones will include facilities for data transmission.
- A cable modem connected to the cable-television network, which the cable company links to the Internet.
- A Digital Subscriber Line (DSL) connection over copper phone lines, which transmits digital data at frequencies higher than the voice band. The same wires may simultaneously transmit analog voice signals, as described in Chapter 25.
- A home, office, or university network that links to the Internet.
- The Internet, usually via high-speed digital lines linked to a regional Internet node, if you have a very powerful computer.

All these services connect you to the Internet, and direct your signals through the switches and routers that process Internet traffic. You will learn more about computer networking in Chapter 26, but the details of Internet transmission are beyond the scope of this book.

Cable-television systems receive signals from distant sources for local distribution.

Video Communications

Video signals are broadcast through the air by ground or satellite transmitters, or are transmitted through cables. There are three main types of video distribution—long-distance distribution from central sources to individual broadcasters and cable companies, direct broadcast from satellites to home receivers, and local distribution from cable companies and broadcast stations to individual homes. You will learn more about these in Chapter 27.

Traditionally, video distribution has been separate from the global telecommunications network, but this is changing. Much distribution of programs to broadcast stations or cable head-ends is via satellite, but some video signals are distributed over fiber. Cable companies link to the Internet to provide cable modem service and to the telecommunications network if they compete with the local telephone company to offer telephone service.

Television is changing from analog to digital technology, which means that video transmission technology is also changing. Broadcasters, cable companies, and satellite television companies are all trying to adapt their infrastructure for digital transmission.

Television distribution is affected by the change to digital technology.

The Internet carries streaming video and video clips, as well as video conferences and calls, but usually the bandwidth is limited and the signal does not match television standards. For example, NASA Television can show you scientists talking about the latest Hubble Space Telescope photos, but the streaming video version on the Internet won't show you the Hubble photos in their full glory. Video clips on the Internet may show more detail, but they take longer to download than they do to play.

You'll learn more about video communications in Chapter 27.

Special-Purpose Networks

Some communication systems don't fit into the major categories we've described so far. The U.S. Department of Defense's high-speed network is designed to survive hostile attack. Electric utilities and railroads use dedicated systems along their rights of way to monitor operations. Often these networks connect with the global telephone network and the Internet. You can call many military phones from the civilian phone network, and many utilities lease surplus fiber capacity to Internet or phone carriers.

Special-purpose networks are part of the convergence trend, but convergence does have its limits. Like "one size fits all" clothing, it may not meet special needs. The U.S. military doesn't want to rely entirely on commercial cables because they are far more vulnerable to potential enemies than are their own customized systems. Likewise, commercial cables are unlikely to run alongside the powerlines that electric grid operators need to monitor.

The Global Telecommunications Network

The global telecommunications network is built around the pieces described so far, and reaches into communities around the world. The actual network is far too complex to depict, but Figure 23.2 shows a greatly simplified view. The diagram ignores the existence of competing local wire-line and mobile telephone companies, long-distance competition, and satellite and broadcast video. It also ignores local computer networks in homes, coffee shops, businesses, and universities. All of these components are interconnected, however, so you can use your cell phone to call the wire-line phone at the local pizza shop, then call your uncle in Australia.

Diverse communication systems are interconnected to form a global network.

The global network operates on many levels and through many media. It starts with local connections, which link to regional networks, then to national and international networks. When you make a call, your phone transmits the signals to a local switching

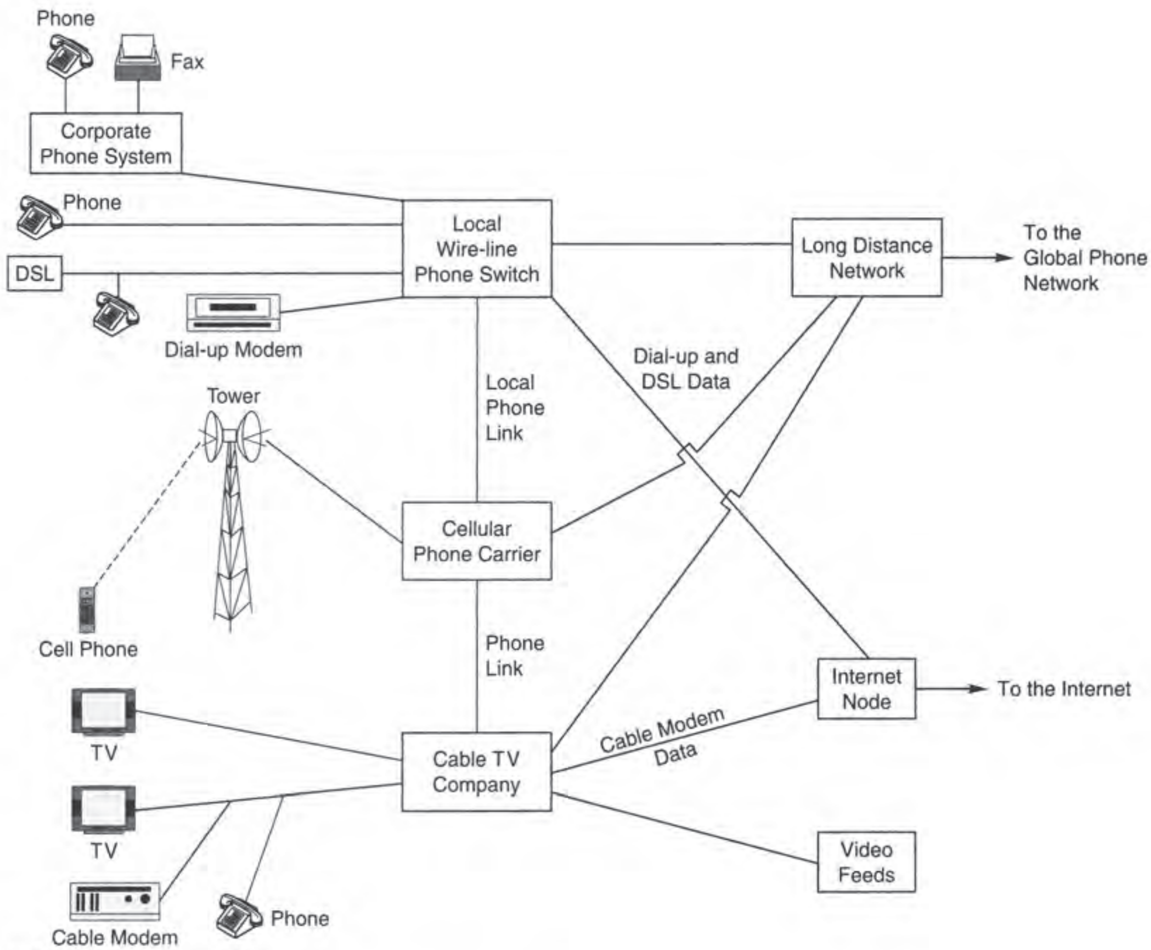

FIGURE 23.2
Interconnection of telecommunication networks.

center, which may send it to another customer of that switching center, to the local cell-phone network, or to the regional network. The regional network may drop your call off at a nearby town, or route it to the national long-distance network. The long-distance network may carry your call anywhere in the country, or to international carriers that make connections around the world. Then the network repeats those steps in the opposite direction to make the final connection.

The links connecting to an individual phone have low capacity, enough to carry one call. Connections between neighboring phone networks have much higher capacity, and can carry hundreds or thousands of calls. Regional networks have even higher capacity, and national and international networks have the highest.

The actual transmission capacities and layout of the transmission grid depend on the demand for services, the distribution of people and industry, economics, and politics. High-capacity transmission lines run between major population centers such as New York and London, or Chicago and San Francisco. Less capacity is needed to serve smaller cities like Syracuse, New York, and Worcester, Massachusetts. Satellites make connections to remote points impractical to reach by cable, like the Falkland Islands. Telecommunication companies exist to make money, so they provide better connections to rich countries than to poor ones. The United States is well connected to Mexico and Canada, for example, but not to Cuba, where connections remain minimal after decades of economic sanctions. (Ironically, the first international submarine telephone cable laid from the United States connected Key West to Cuba in 1950.)

The telecommunications network is like the circulation system of the human body. Blood flows from tiny capillaries to larger veins, which in turn feed larger veins that carry more blood. After the blood passes through your heart and lungs, it is divided into smaller and smaller arteries and ultimately reaches the tiny capillaries. Individual phone and Internet links are the capillaries of the telecommunications network.

Signals are multiplexed for high-capacity transmission.

Individual input signals are combined or multiplexed together at various stages in the telecommunications network. Telephone calls are converted into strings of 64,000 bits per second, and strings of bits from separate calls are shuffled together to make higher-speed signals. Those faster signals, in turn, are shuffled together to carry even more calls. Telephone networks have a hierarchy of data transmission rates; the Internet uses a similar hierarchy, although it combines input signals in a different way.

Both hierarchies are well-established and accommodate a wide range of equipment. This is especially true in the case of the telephone network, which has long remained compatible with older hardware. With wire cutters and a screwdriver, you can attach a massive black 1950s-vintage dial phone to the same set of phone wires that can carry DSL broadband Internet service. The Internet is not that flexible.

Let's look at the important hierarchies, starting with the one that connects to your phone line.

Time-Division Multiplexed Telephone Hierarchy

The digital telephone hierarchy use *time-division multiplexing*, with rates based on interleaving input signals into a faster data stream at an exact multiple of the input data rate. Appendix C lists important standards and the data rates they carry. Figure 23.3 shows how time-division multiplexing works in the North American *plesiochronous digital hierarchy*, used at the lower levels of the telephone network. This is a decades-old AT&T standard, but it does the job well.

The digital telephone hierarchy time-division multiplexes input data streams.

The analog voice signals from your phone are first converted to digital form at a rate of 64,000 bits per second to produce a single *voice channel*, sometimes called a *DS0* signal, that is input to the multiplexing stage. Twenty-four DS0 voice channels are interleaved to produce one *DS1* signal at 1.5 Mbit/s, which is transmitted over a *T1* line. ("DS" means digital signal, and the number is the step in the hierarchy; "T" means transmitted signal, with the same numerical coding.) Four DS1 signals are interleaved to give one 6.3 Mbit/s DS2 signal, and seven DS2 signals are interleaved to give one 45 Mbit/s DS3, transmitted over a T3 line.

In practice, not all the slots for lower-speed inputs are filled. The number of input voice channels depends on the calling volume, which peaks during the business day but drops to near zero around 3 A.M. The 1.5 Mbit/s T1 rate and the 45-Mbit/s T3 rate are the most

FIGURE 23.3
Time-division multiplexing in the North American digital telephone hierarchy.

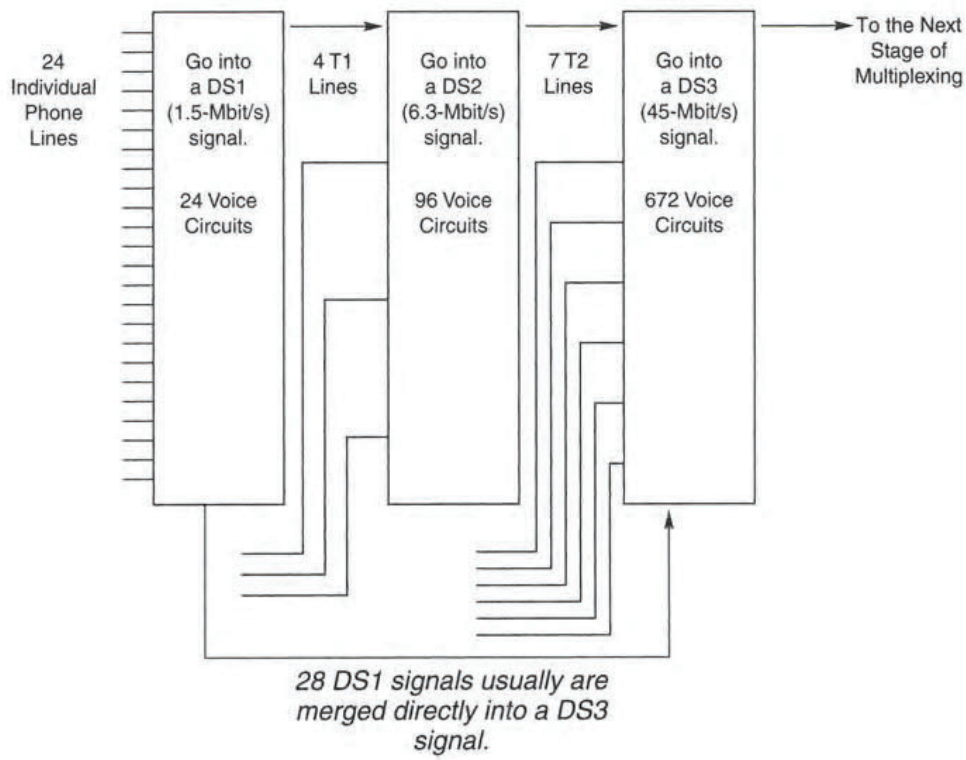

important line rates in use; most phone companies now skip directly from the DS1 to the DS3 rate, without bothering with the intermediate DS2 rate.

The original specifications called for a 274-Mbit/s T4 rate, but it never found wide application and is now ignored. During the 1980s, a variety of other time-division multiplexing rates were introduced in North America. The major family was based on a nominal data rate of 405 Mbit/s (equivalent to 6048 voice channels, or nine T3 carriers) with additional rates of 810 Mbit/s, and 1.7 and 2.4 Gbit/s that were multiples of the 405 Mbit/s rate. These systems were not formally accepted as standards, and have since been supplanted by SONET and related systems.

European/ITU Telephone Hierarchy

Europe uses a different digital hierarchy.

The digital telephone hierarchy developed by the European Conference of Postal and Telecommunications Administrations combines 64,000 bit/s voice circuits differently than in North America. It first combines 30 digital voice circuits to give a 2.048 Mbit/s *Level 1* channel. Next it combines four Level 1 signals to make an 8.448 Mbit/s *Level 2* signal (120 voice channels). Additional multiplexing steps combine four lower-level channels, giving *Level 3* (480 voice channels) at 34.305 Mbit/s, *Level 4* (1920 voice channels) at 139.264 Mbit/s, and *Level 5* (7680 voice channels) at 565.148 Mbit/s. These rates are sometimes called *E rates* because they originated in Europe. The standards were formally adopted by the International Telecommunications Union, and are often called ITU or CCITT standards, from a French translation of the name of an ITU commission. All are listed in Appendix C.

In practice, Level 1 and Level 3 are the most widely used rates, similar to the T1 and T3 rates used in North America, and Level 4. As in North America the highest data rate, the 565 Mbit/s Level 5, was never widely used and is now obsolete.

Japan also developed a system of its own. The first two levels are the same as the North American T1 and T2 line rates, but the next step is a factor of 5 (rather than the 7 used in North America) to 32.064 Mbit/s or 480 voice channels, and the final step is a factor of 3 to 97.728 Mbit/s or 1440 voice channels.

SONET/Synchronous Digital Hierarchy

In the 1980s, the telecommunications industry wrote a new set of time-division multiplexing standards for higher-speed transmission of mixed voice and data. These became the *Synchronous Optical Network* (SONET) and *Synchronous Digital Hierarchy* (SDH) standards covered in Chapter 20. The two standards are quite similar, but differ in details important for compatibility with the older North American and ITU digital telephone hierarchies. As you learned in Chapter 20, they package data into frames, which allocate specific slots for input channels, making them function as circuit-switched systems. SONET and SDH can also transmit packets, but not as efficiently as a true packet-switched system such as the Internet Protocol.

SONET packages input signals into frames to guarantee capacity.

The base of the SONET standard is the 51.84-Mbit/s OC-1 (Optical Carrier-1); an electronic signal at that rate is called STS-1. The OC-1 rate was designed to handle the T3 input signals common in North America, but it can also handle input signals in ATM or IP formats. The next step up the SONET/SDH ladder is OC-3, 155.52 Mbit/s, which in SDH is designated as STM-1. This format was designed to match Level 4 of the ITU digital hierarchy, as shown in Figure 23.4.

Overhead consumes about 10% of SONET/SDH capacity.

If you compare data rates, you will notice that the SONET/SDH rates are about 10% higher than the rates quoted for the input signals from the digital telephone hierarchies.

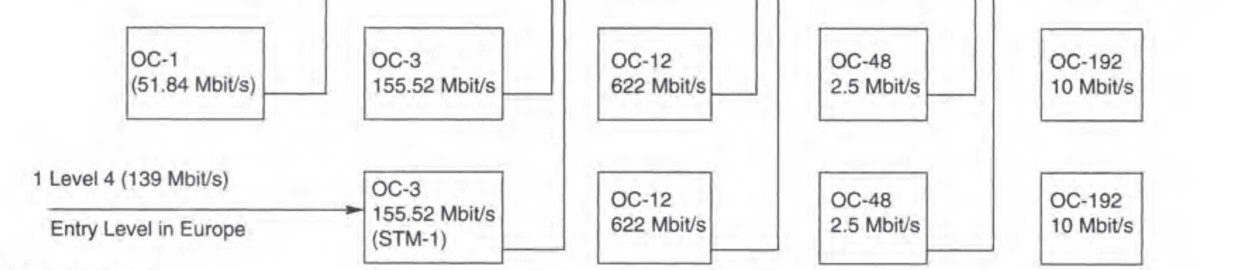

FIGURE 23.4

SONET/SDH time-division multiplexing hierarchy.

This reflects additional "overhead" data needed to manage and monitor signal transmission. The SONET standard includes capabilities to spot failures and redirect signals to maintain service. The SONET and SDH standards also specify network structures that make recovery from failure easier.

From 155.52 Mbit/s, both SONET and SDH step up by factors of 4, to 622 Mbit/s, then to roughly 2.5, 10, and 40-Gbit/s. So far 10 Gbit/s is the upper limit in most practical systems, although developers have begun to offer 40-Gbit/s equipment.

SONET and SDH are time-division multiplexing formats, which cover the speed and sequence of bits transmitted in a single data stream. This means they are designed for use on a single optical channel. You can stack multiple SONET or SDH signals (or mixtures of the two) onto a fiber using wavelength-division multiplexing.

Internet Transmission

The Internet uses telephone-standard data rates.

Internet data transmission rates are the same as those used for digital telephone lines. However, Internet data streams are generated differently from telephone data, and are not necessarily transmitted in the same formats.

The Internet relies on packet switching to generate data streams, not the time-division multiplexing used in telephone networks. Internet routers put data packets into a queue and transmit them in the sequence they arrive. If no other data packets are waiting, incoming packets are transmitted right away. If other data packets are waiting, the new packets wait their turn.

In practice, the output data rate is matched to one of the standard TDM line rates. The Internet data stream may be packaged into SONET frames for transmission through fibers, or transmitted as a string of bits at the same rate but without the SONET information and overhead.

Optical Channels

An optical channel carries one SONET/SDH signal.

SONET/SDH signals modulate a single optical channel. As you learned earlier, wavelength-division multiplexing allows a single fiber to carry many optical channels. Because the optical channels are entirely independent, WDM also allows a single fiber to carry different kinds of traffic at different data rates.

ITU standards specify the spacing of optical channels in both DWDM and CWDM systems. The 50- or 100-GHz channel spacing of DWDM allows potentially huge transmission capacity. The erbium-amplifier C- and L-bands each can transmit about 80 optical channels with 50-GHz spacing, or 40 with 100-GHz spacing. At 10 Gbit/s per channel, that adds to a total capacity of 1.6 Tbit/s on a single fiber. Single fibers have carried higher data rates—above 10 Tbit/s—in the laboratory, but there's no demand for that much capacity. In fact, most WDM systems carry only a tiny fraction of their maximum potential capacity.

Actual Transmission Capacity

The telecommunications bubble led to a tremendous overexpansion of fiber transmission capacity, particularly in long-haul terrestrial and submarine systems. Carriers believed that

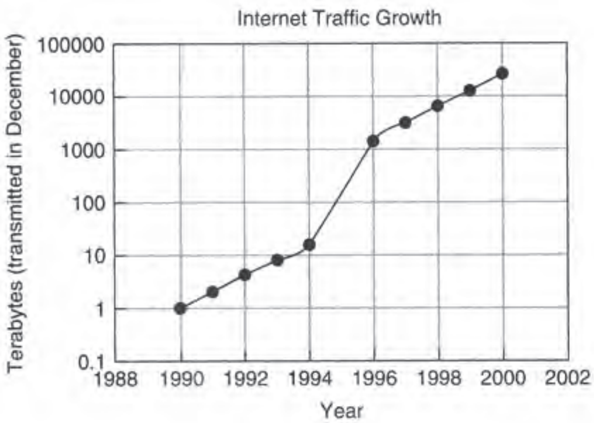

FIGURE 23.5
Growth of Internet Traffic based on data by Andrew Odlyzko, University of Minnesota.

Internet transmission capacity was doubling every three months, or by a factor of 16 each year. Later these numbers began to look dubious. The doubling claim seems to have come from a February 1997 press release from WorldCom, a company whose accounting later became highly suspect. Andrew Odlyzko of the University of Minnesota tracks Internet traffic and says traffic might have grown that fast briefly as the World Wide Web became popular in 1995 and 1996. However, his figures show that traffic since then has roughly doubled each year, as shown in Figure 23.5. That's respectable growth, but by a factor of eight less than had been claimed.

Nobody actually counts transmitted bits, but Odlyzko estimates that the volume of U.S. Internet backbone traffic averaged about 300 Gbit/s during one month in late 2002. That's less than half the capacity of a single fiber carrying 80 DWDM channels at 10 Gbit/s each, although that comparison isn't really fair. Internet traffic varies during the day, slumping in the early morning. The traffic also is going in different directions all across the country, carried by many different networks running many different places. Only a few links between major cities need even a single 10-Gbit/s channel.

A few Internet links between major cities carry 10 Gbit/s.

By late 2002, potential fiber capacity was much larger. Typically the cables that telecommunications carriers installed along major transmission routes contained lots of extra fibers to allow for growth. Buying rights of way and installation costs ran as high as 95% of the total cost, so the extra fibers seemed a bargain compared to having to go back and lay another cable. Each of those fibers can potentially carry many optical channels, but the carriers still don't use them all. They lit fibers, and channels on the fibers, only when they needed them. Most of the fibers were dark, and even the lit fibers had most of their optical channels available. Even the lit channels weren't full. The market research firm RHK Inc. estimated that Internet backbone traffic in late 2002 amounted to only 5% to 15% of fully-provisioned capacity. At the same time, Washington consulting firm TeleGeography estimated that only 10% of installed fibers carried any traffic, and that in those fibers only 10% of the available channels had been lit.

Fiber capacity greatly exceeds the demand for long-distance links.

So a glut of fiber spread through the network, particularly for long-distance transmission. The cost of buying and installing that fiber in anticipation of traffic that never came helped drive carriers such as WorldCom and Global Crossing into bankruptcy.

THINGS TO THINK ABOUT

The Fiber Glut

The irony of the fiber glut is that the fiber-optic industry succeeded much too well in meeting the demand for transmission capacity. In the 1990s, the capacity of a single optical fiber soared from a single channel to several dozen. The key breakthrough was the invention of the *erbium-doped fiber amplifier.* Wavelength-division multiplexing had been tried before, but it had been impractical without a way to amplify the signals in optical form.

In many ways, fiber optics was the right technology at the right time. The Internet was hungry for transmission capacity, and fiber could provide it. Excited investors poured money into telecommunications technology. In the early stages of the bubble, some people became very rich. But the industry climbed the peak of demand so fast that it never noticed that it had run over the cliff until it looked down, and fell with a big splat.

Now we're left with more fiber than we know what to do with. The companies that laid the cable have little to show for it now, but the fiber itself is a potential resource that we're only starting to tap. For example, five British radio telescopes have been linked with the 76-meter Jodrell Bank dish to form a 217-kilometer array called e-MERLIN that will be able to see the universe in radio waves as sharply as the Hubble Space Telescope can see it at visible wavelengths. e-Merlin requires a steady stream of data at 30 Gbit/s from each of the five dishes—but that's only three optical channels on a single fiber. We can only wonder what other opportunities the surplus fiber will present.

Fiber Transmission Business

Cable transmission capacity may be shared.

The economics of operating fiber-optic networks can be almost as complex as the technology. The sharing of transmission capacity is one subtle complexity that arises from the economics of laying cables.

Fibers are cheap compared to laying cable, so it makes economic sense to add plenty of extra fibers. Some companies made a business of laying cables and leasing fibers to other companies that provided telecommunication services. Carriers quickly realized that leasing fibers on somebody else's cable could cost much less than laying their own cable. Carriers also began swapping capacity in each other's cables so, for example, a long-distance and a regional carrier wouldn't both need to lay cables to all the towns they served. Swapping capacity sometimes extended to optical channels on a single fiber, or even to the multiplexed components that went into a single TDM data stream.

Carriers want redundant transmission paths to guarantee service.

These transactions could get quite complex, involving multiple companies, and creating some unexpected consequences. For example, big carriers like to provide redundancy in their networks by dividing their traffic to a particular city between two or more cables. That way, a single cable break won't knock out service totally. However, when a fire in a Baltimore railroad tunnel destroyed a fiber-optic cable, some carriers discovered that their "redundant" optical channels were routed through that one cable. Further investigation showed that wasn't the only cable that carried several "different" fiber routes.

Competitive carriers often lease capacity on fibers running through the same cable. They don't interfere with each other's transmission because the signals remain separate. However, users who have contracts with both companies to ensure service during a cable outage could have an unpleasant surprise.

Submarine Cables

The largest links in the global telecommunication network are submarine fiber-optic cables, which cross oceans to link continents. They have been vital links since the age of the telegraph, shrinking the world. In the rest of this chapter, we will look at them and their long-haul terrestrial cousins that form the backbone of telecommunications on land. Later chapters will look at regional and local systems.

Submarine fiber cables are the backbones of intercontinental telecomm-unications.

Submarine cables come in many types. Some cross the few kilometers of seawater separating an island from the mainland; one of the first to use fiber was an 8-km cable from Portsmouth, England, to the Isle of Wight off the English coast. Many cross tens or hundreds of kilometers of sea; the Mediterranean and Caribbean seas are crisscrossed with submarine cables. Some span thousands of kilometers of ocean; the longest run across the bottom of the Pacific and from Europe to Japan.

Submarine cables must meet extremely tough requirements. Their transmission capacity should be as high as possible because the cables are costly to make, lay, and operate. The cable, and any optical amplifiers or repeaters, must withstand harsh conditions on the bottom of the ocean for a design life of 25 years. Components must be extremely reliable because it is very expensive to recover the cable from the sea floor and haul it to the surface for repairs. The cable should transmit digital signals cleanly to be compatible with modern equipment. These specifications veritably call out "fiber optics," and since the 1980s fibers have been standard for submarine cables. Figure 23.6 shows how they have spread around the world.

The Impact of Submarine Cables

Submarine cables date back to the days of the electrical telegraph, and for well over a century they have played a vital role in binding the world together. Undersea telephone cables came long after telegraph cables, and the first transatlantic fiber-optic cable was not laid until 1988. However, since then the technology has grown at amazing speed.

The first submarine telegraph cable was laid in 1850.

The first submarine telegraph cable was laid in the English Channel in 1850, as Europe expanded its electrical telegraph system. It carried only a few messages between England and France before a fisherman snared it and hauled a piece to the surface. He thought it was a strange type of seaweed. That experience taught submarine cable engineers an important lesson—waterproof isn't enough. Fishing trawlers and ship anchors remain the biggest threat to cables in shallow water, so modern cables are buried a meter below the sea floor except in ocean depths below a few hundred meters.

The first attempt to lay a transatlantic telegraph cable failed in 1857, and it was not until 1866 that a reliable cable began operating under the Atlantic. Very long telegraph cables were possible in the nineteenth century because mechanical relays could amplify their dots and dashes. Long-distance telephone transmission proved much harder because it required electronic amplifiers. Vacuum tubes relayed signals on land, but transatlantic telephone

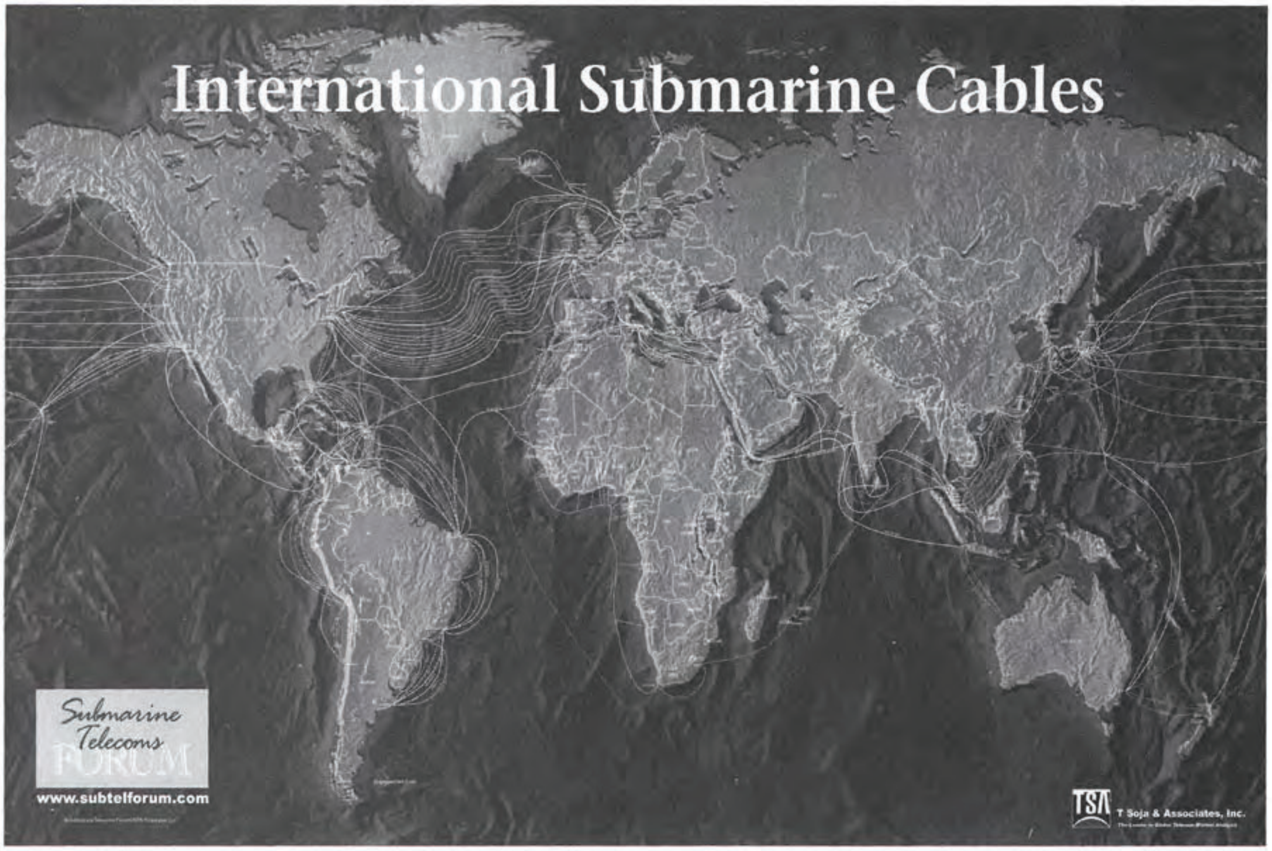

FIGURE 23.6
Global submarine fiber-optic cable systems. (Courtesy of Submarine Telecoms Forum and WFN Strategies, LLC)

calls had to rely on short-wave radio links until 1956, when the first transatlantic telephone cable, TransAtlantic-1 (TAT-1), began operation between Britain and Canada. It included vacuum-tube amplifiers sealed into special cylinders built to withstand the tremendous pressure at the ocean bottom.

TAT-1 was made of copper coaxial cable, which offers the highest bandwidth of any standard metal cable. However, coax attenuation increases with the square root of transmission frequency ν and can only be reduced by increasing the inside diameter D of its outer conductor.

$$\text{Attenuation} = \frac{\text{constant} \times \sqrt{\nu}}{D}$$

Switching to fiber increased cable capacity by a factor of 8 over coax.

This means that to transmit higher frequencies, coaxial cables must be made fatter, or have repeaters spaced closely together. Later engineers were able to squeeze up to 4200 telephone circuits onto a single coaxial cable, but to span the Atlantic it had to be an unwieldy 5.3 cm in diameter and required one repeater every 9.5 km—a total of 664.

By the mid 1970s, it was clear that coax had reached the end of the line, and satellites looked like they would eventually dominate intercontinental communications. However, Bell Labs turned to fiber, and in 1980 announced a design for the first transatlantic fiber-optic cable, TAT-8. By using single-mode fiber transmitting at 1300 nm, they calculated they could transmit 280 Mbit/s on each of two fiber pairs, with a third fiber pair kept in reserve. With digital data compression, that was equivalent to a total of 35,000 voice channels, over eight times more than TAT-7. After eight years of hard work, TAT-8 began service at the end of 1988 with repeaters averaging more than 50 km apart.

That began over a dozen years of explosive submarine fiber-optic growth. Repeaters gave way to optical amplifiers in the mid-1990s, and wavelength-division multiplexing allowed each fiber to carry multiple optical channels. The Apollo transatlantic cable, which began service in February 2003, started operation with 16 channels carrying 10 Gbit/s on each of four fibers, a total of 640 Gbit/s. That's more than a thousand times the ground-breaking capacity TAT-8 offered just 15 years earlier. Plans call for upgrading Apollo's capacity to 80 channels per fiber, a maximum of 3.2 Tbit/s. With the current fiber glut, that's likely to take a while.

When it began service in 1988, TAT-8 carried a thousand times more traffic than TAT-1 had when it was laid 32 years earlier. By then TAT-1 had been retired, shut down in 1978 after 22 years of service. Ironically, new generations of fiber-optic technology made TAT-8 itself obsolete much faster. Dwarfed by the huge capacity of WDM submarine cables, TAT-8 was quietly retired in 2002 after a mere 14 years of service because its capacity was too small to justify the cost of maintaining it.

TAT-8 was retired as obsolete in 2002 after 14 years of service.

Submarine Cable Basics

The design of submarine cables is shaped by the submarine environment, which is very different than the environment for terrestrial cables. The underwater environment is very stable, very extreme, and very hard to reach. Repairs of undersea cables are difficult and expensive, so systems are designed to operate without service for long periods. Specifications usually call for no more than two underwater repairs in a cable's nominal 25-year lifetime, and the target is no repairs.

Submarine cables are designed for long life, high pressure, and high speed.

Electrical power must be transmitted from the cable termination points on land, so power is at a premium. Early systems used repeaters, but since the mid-1990s all submarine cables have used only optical amplifiers underwater. Full "three-R" regeneration—reshaping and retiming pulses as well as reamplifying them—is done only at the cable termination points on shore. For intercontinental cable systems spanning thousands of kilometers, this imposes very stringent requirements on the levels of noise, dispersion, and nonlinear effects in the transmitting fiber and the optical amplifiers.

As you learned in Chapter 8, the fiber-optic cables used in submarine systems are highly specialized, with fibers embedded deep in the core of a pressure-resistant structure. The outer layers of deep-sea cables are medium-density polyethylene. Heavy metallic armor covers the polyethylene in shallow-water cables, which are buried to protect them from fishing trawlers and ships' anchors—the undersea counterparts of backhoes. If undisturbed, the cable structure should withstand intense pressures and exclude salt water for decades.

Submerged amplifiers are housed in pressure-resistant cases.

Optical amplifiers are mounted inside pressure-resistant cases originally developed to house repeaters. They are built into the cable but are larger in diameter, so at first glance they resemble a rabbit swallowed by a python. Submarine cable developers still call these cases "repeaters," but don't be fooled—repeaters have not been used on submarine cables for several years (although the repeaters on old cables have not been replaced).

Submarine cables fall into two broad classes, *unrepeatered* and *repeatered*. In the world of submarine cables, these terms define whether or not the system includes optical amplifiers with their pump lasers in the same underwater housing as the optical amplifier. Underwater pump lasers mark a key dividing point because they are electronic components subject to failure, and because they require electrical power to be transmitted through the cable. These two types can be further subdivided according to the distance they span and their configuration, but we will concentrate on the basic categories. As you will see later in this chapter and in Chapter 24, these systems have counterparts on dry land.

Unrepeatered Undersea Cables

Unrepeatered cables can span up to 400 km.

The length of unrepeatered cables is limited by the need to sustain sufficient optical power to avoid repeaters. In simplest form, this generally means distances less than a couple hundred kilometers from end to end, so the transmitter output is not attenuated below the receiver threshold. This distance can be stretched to as much as 350 or even 400 km by adding a postamplifier after the transmitter, a preamplifier before the receiver, and remotely pumped amplifiers, as shown in Figure 23.7.

Transmission distance in a repeaterless system normally is limited by the system power budget, which depends on the transmitter output power and the receiver sensitivity. For

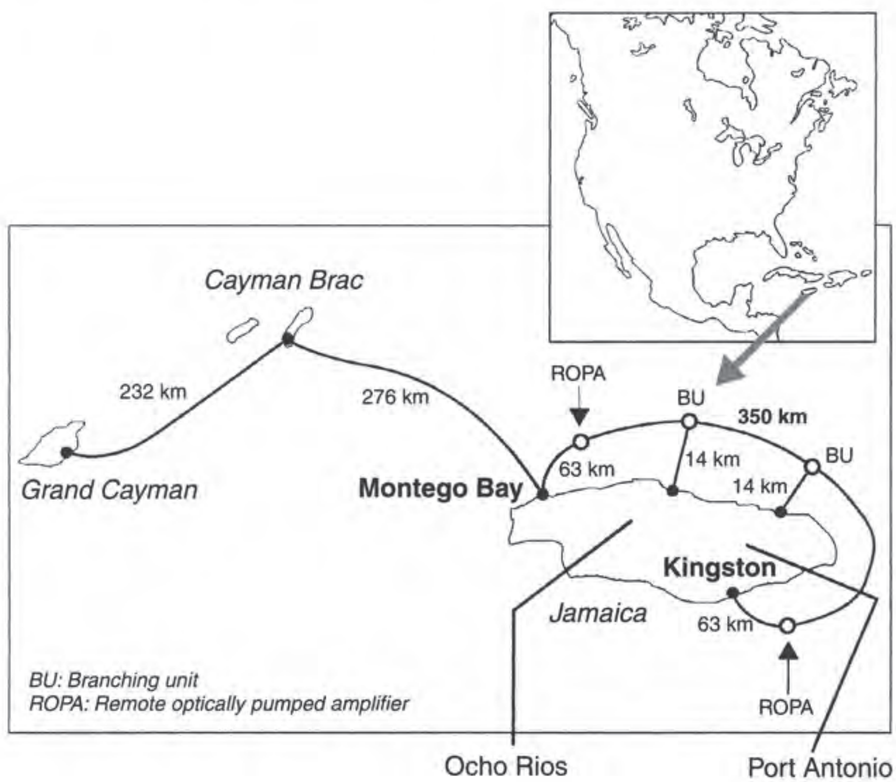

FIGURE 23.7
Unrepeatered systems can hop a series of islands to span long distances. Use of optical postamplifiers, optical preamplifiers, and remotely pumped optical amplifiers can stretch single spans. (Copyright Alcatel)

example, a system in which the transmitter output is 28 dB above the receiver sensitivity normally can span about 130 km without amplification. Longer distances are possible by increasing transmitter power and/or receiver sensitivity.

Onshore or remotely powered amplifiers can stretch unrepeatered transmission distances.

As you learned earlier, a postamplifier can increase the output power from a transmitter. Likewise, a preamplifier can increase receiver sensitivity. It's possible to stretch "repeaterless" transmission further by siting *remote optically pumped amplifiers* offshore. In these systems, the cable from shore to the offshore amplifier includes a separate fiber that carries light from an onshore pump laser, which then does not require submerged pump electronics. Fiber attenuation makes it impossible for a remote pump laser to deliver as much pump power as a pump laser could if it was in the submerged housing, so the amount of amplification is less than with a conventional optical amplifier. The design also requires dedicating one fiber slot in the cable to a pump fiber for each remote optical amplifier. However, post- or preamplification with optically pumped amplifiers can stretch the power budget to as much as 88 dB for single-channel transmission at 2.5 Gbit/s. In Figure 23.7, the use of two remote optically pumped amplifiers allows the cable around Jamaica to stretch a total of 350 km. (Transmitter powers can be higher than on longer cables because nonlinear effects do not build up over the short distance.)

Using remote amplifiers limits the number of fiber pairs that can carry traffic because repeater housings can contain only a limited number of amplifiers. Thus a 100-km system with no amplifiers can carry signals on as many fibers as can fit in the cable (a few dozen in current designs), but a 300-km system would be limited to fewer fibers.

Unrepeatered submarine cables are widely used to connect the mainland with offshore islands, link islands with each other, or loop along the coast of a continent or large island. Most run only a few kilometers to tens of kilometers between islands or from the mainland to an offshore island. Examples are across the English Channel, or between islands in Hawaii, Denmark, Japan, or Indonesia.

Repeaterless systems can span longer distances by island-hopping, as shown in Figure 23.7. Another approach is to run a series of unrepeatered cables between coastal cities in a *festoon system*, such as the one around Italy shown in Figure 23.8. Laying offshore festoon cables to link coastal cities may cost less than burying cables on land—particularly where cities are along the coast and the terrain is mountainous, as in Italy.

Repeatered Submarine Cables

Transmitting signals over spans longer than about 300 to 400 km requires submerged repeaters, which put important constraints on repeatered systems. The cable itself must carry electrical power from the termination points on shore. This electrical power transmission capability is limited, restricting the total number of amplifiers in the chain, and thus both the total distance spanned and the number of usable fiber pairs. Typical repeater housings can hold only 8 to 16 fiber amplifiers, also limiting the number of fiber pairs.

Repeatered cables can fit fewer active fiber pairs.

All repeatered submarine systems are effectively long-distance systems, although the total distances range from hundreds of kilometers to a total of more than 20,000 km in the longest systems. They are designed in various configurations. Some run thousands of kilometers between two points on opposite sides of an ocean, such as the east coast of the United States and the west coast of Europe. Others are loops or rings. Many large systems

FIGURE 23.8
Submarine fiber cables link coastal cities in Italy. They are part of a network that includes land lines. (Courtesy of Corning Inc.)

run along coasts, landing at a number of points, such as the SEA-ME-WE-3 (Southeast Asia-Middle East-Western Europe) system shown in Figure 23.9. They may include underwater branching points, where optical channels or whole fibers worth of signals are added and dropped. Many long systems like SEA-ME-WE-3 have intermediate landing points where signals can be regenerated if needed.

Submarine cables are designed from end to end as a unit.

The terminal points of submarine cables link with long-distance terrestrial cables and often with other submarine cables as well. One important difference between submarine and terrestrial long-distance systems is that the submarine systems are designed from end to end to function as units, while terrestrial links are part of larger networks built up of many cable spans. This means that submarine designers can count on complete control of

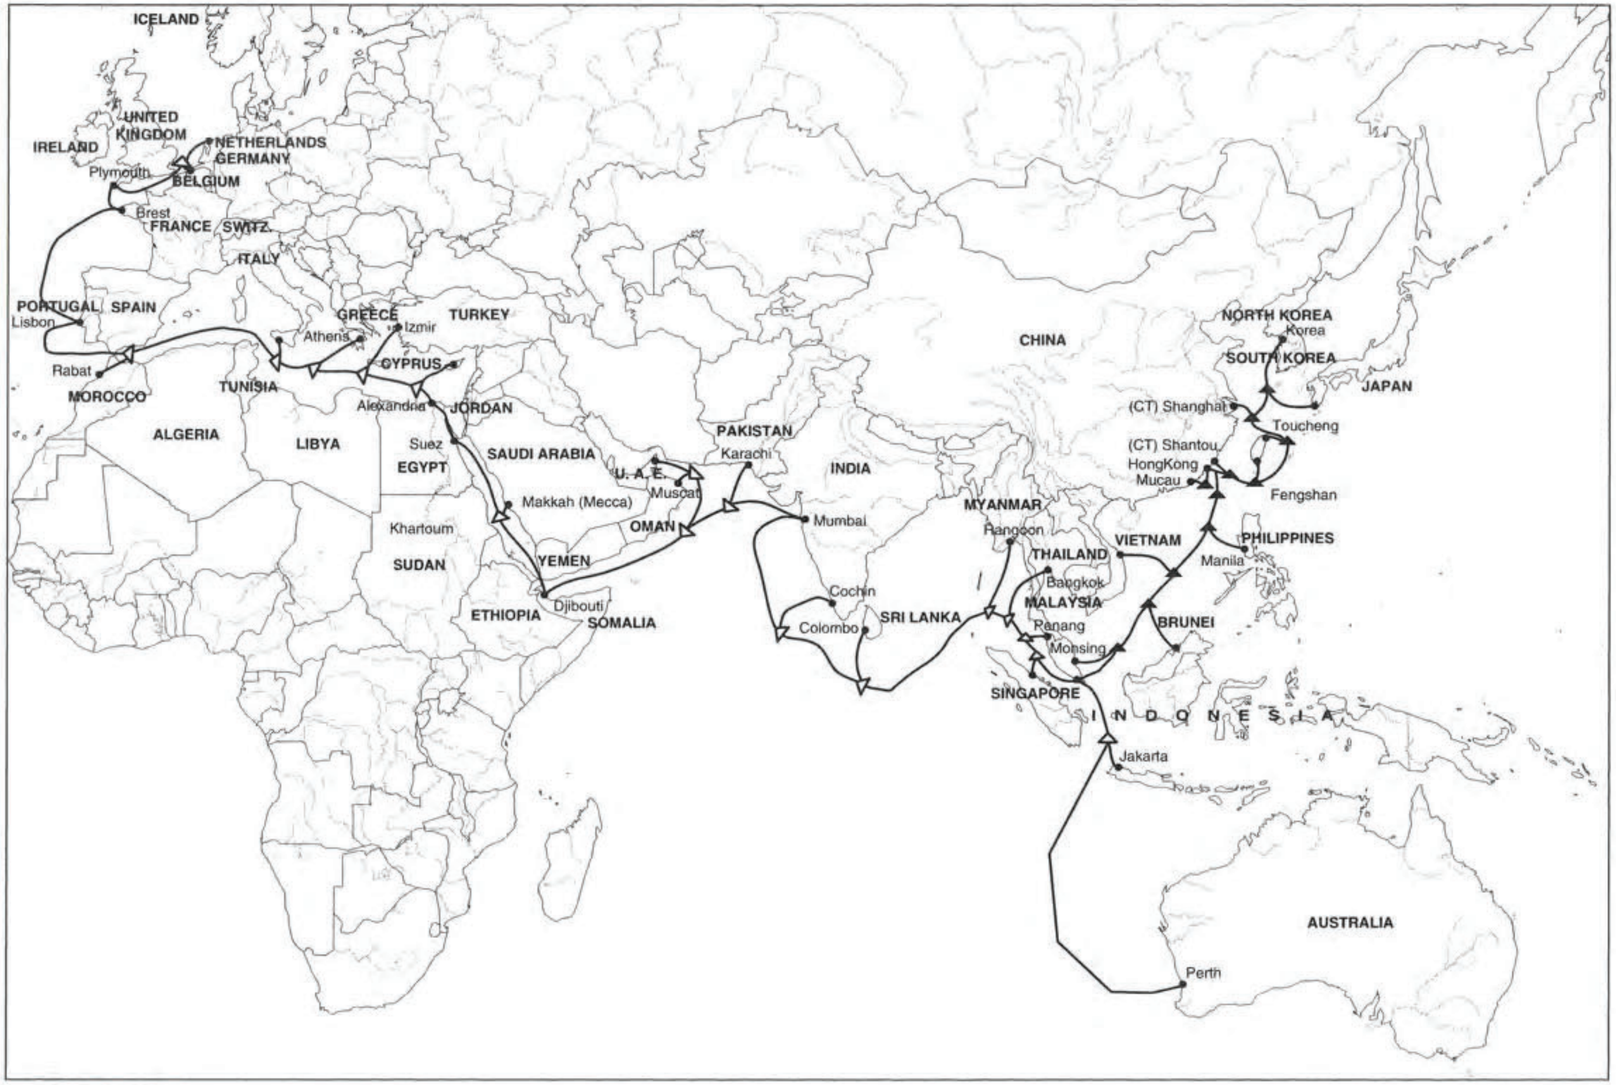

FIGURE 23.9

SEA-ME-WE-3 submarine cable system lands at many points from western Europe to Southeast Asia, spanning a total length of 27,000 km.

their entire cable system, and deploy different fibers along its length to optimize performance. As you learned in Chapter 22, this makes it possible to install fiber with large effective area near the transmitter to control nonlinear effects, while using other fibers elsewhere along a span. Designers of most terrestrial cables cannot count on this flexibility.

Another important difference is that submarine cables link points on islands and continents, but not in the ocean. All drops go to land. Terrestrial cables go through many sparsely populated regions, but farmers, ranchers, and rural villages need much more bandwidth than do fish. Networks of terrestrial cables link many points throughout the regions they serve; submarine cables link points on the edges of the oceans.

Design constraints are tight for high-speed repeatered submarine cables.

Repeatered submarine cables spanning a few thousand kilometers face stringent design constraints to maintain signal quality. Gains of optical amplifiers in submarine cables are kept low—typically around 10 dB—to control noise from amplified spontaneous emission and nonlinear effects, which accumulate along the length of the cable. This limits optical amplifier spacing to about 50 km, roughly half the distance in terrestrial long-haul systems. Using higher input power and lower gain also helps to equalize power across the spectral range of the optical amplifier, which is critical in systems that may contain 100 or more amplifiers in series.

Precise dispersion management is crucial. In state-of-the-art submarine systems, a cable run between amplifiers includes three types of fiber. Large-effective-area fiber is used for the first part of the run after the optical amplifier to minimize nonlinear effects arising from the high-power levels. Typically a length of nonzero dispersion-shifted fiber designed for submarine use follows, with zero dispersion shifted to a wavelength longer than the erbium-fiber amplifier band. Then comes a length of standard step-index single-mode fiber. The dispersion in each fiber segment is large enough to limit four-wave mixing, but the overall dispersion is low enough to allow high-speed transmission across the erbium-fiber amplifier band. Raman amplification in the final fiber segment both preamplifies the signal for the optical amplifier and equalizes gain across the WDM spectrum.

Potential capacity of some cables exceeds 1 Tbit/s.

The capacity of repeatered submarine cables has increased steadily, as shown in Table 23.1. The first system, TAT-8, transmitted at 1300 nm. Single-channel 1550-nm transmission began in the early 1990s to stretch repeater spacing. WDM systems in the erbium amplifier window followed in the late 1990s, with the first carrying 4 or 8 channels per fiber at 2.5 Gbit/s, and total capacity to 40 Gbit/s. The latest systems have slots for dozens of optical channels, but transmitters and receivers are not installed on all wavelength slots when the cable is laid. The Apollo transatlantic cable had 16 channels at 10 Gbit/s each on four fiber pairs when it began service in 2003, with total data rate of 640 Gbit/s. Each fiber can transmit up to 80 channels, for total capacity of 3.2 Tbit/s. The use of WDM with extra channel slots leaves a window for upgrades to prevent early obsolescence like TAT-8, but those won't be needed until the fiber glut is eased.

Table 23.1 samples the longest and highest-capacity repeatered submarine cables. Many shorter repeatered systems have been installed on routes in the Mediterranean and Caribbean that don't require the biggest possible capacity. For example, the 1300-km Black Sea Fiber Optic Cable System links Bulgaria, Ukraine, and Russia with two fiber pairs carrying a single 2.5-Gbit/s channel, with provisions for adding WDM. The shorter lengths of these cables and their 2.5-Gbit/s data rates ease design requirements and dispersion limitations.

Table 23.1 Initial capacities of some major undersea fiber cables.

System	TAT-8	TAT-10	TAT-12/13	SEA-ME-WE-3	Pan American Crossing	Apollo
Operational	Dec. 1988	1992	1996	1998	2001	2003
Location	US–UK and France	US–Germany	US–UK–France–US loop	Germany to Singapore	California–Mexico–Panama, Venezuela, St. Croix	US–UK–France
Initial data rate per fiber pair	278 Mbit/s	565 Mbit/s	5 Gbit/s	2.5 Gbit/s per optical channel	10 Gbit/s per optical channel	10 Gbit/s per optical channel
Working pairs	2	2	2 each half of loop	2	3	4
Fiber	Single-mode	Single-mode	Dispersion-shifted to 1550 nm	Zero dispersion at 1580 nm	Nonzero dispersion-shifted	Dispersion-compensated mix
Repeater spacing	Over 50 km	Over 100 km	None	None	None	None
Wavelength	1300 nm	1550 nm	1550 nm	Up to 8 near 1550	Up to 32 near 1550 nm	16 initially, eventually 80 near 1550 nm
Optical amplifiers	None	None	Yes	Yes	Yes	Yes
Total cable capacity	560 Mbit/s	1130 Mbit/s	10 Gbit/s	Up to 40 Gbit/s	960 Gbit/s	3.2 Tbit/s
Notes				Optical add-drop capability	Not all channels installed initially	Not all channels installed initially

Branch Points and Landings

As you can see in Figures 23.6 through 23.9, submarine cables have a variety of configurations. These illustrations show two distinct approaches to making connections to terrestrial networks—branch points and landings. These are highlighted in Figure 23.10.

Submarine cables may branch underwater.

Branch points are undersea cable junctions where fibers or optical channels are divided between two or more destinations. The first transatlantic fiber-optic cable, TAT-8,

FIGURE 23.10
Undersea cable branch connection is an offshore add-drop multiplexer.

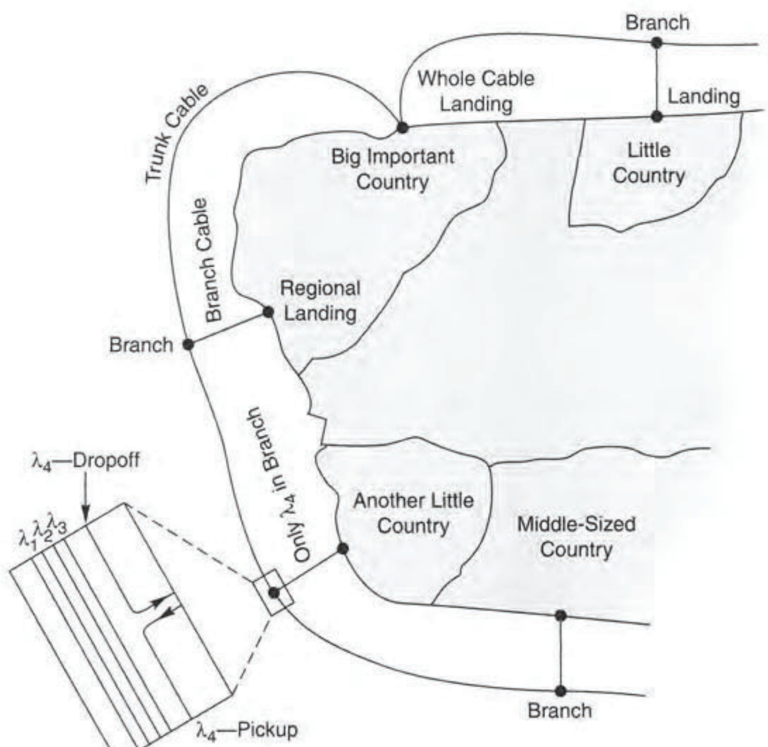

divided off the French coast, with signals split between Britain and France. Cables such as SEA-ME-WE-3 have offshore branching points that essentially drop signals at a coastal city. Figure 23.10 shows how a fiber drops one wavelength at "Another Little Country," while the remaining wavelengths continue on the submerged cable offshore.

Landings link submarine cables to the terrestrial network.

Landings are points where the entire submarine cable comes on land and connects with the terrestrial network. Transatlantic cables like TAT-8 normally only land at their end points, where they collect signals from and distribute signals to terrestrial networks. However, cables that run around continents, like SEA-ME-WE-3, land at many points, where the signals may be regenerated as well as linked to local terrestrial networks. Festooned cables generally don't require regeneration at every landing, and some of the jumps may be short.

Many transoceanic cables form loops, with a pair of parallel cables running between continents and making connections on land. The loops function like SONET rings, providing backup transmission capacity in case one of the two cables breaks. The TAT-12/13 and Apollo cables listed in Table 23.1 are examples.

Long-Haul Terrestrial Systems

Terrestrial cables are not as long as the longest submarine cables.

Long-haul terrestrial telecommunication systems, like long-distance submarine cables, carry high-speed signals and serve as backbones of the global telecommunications network. The same principles underlie the operation of submarine and terrestrial systems. The main

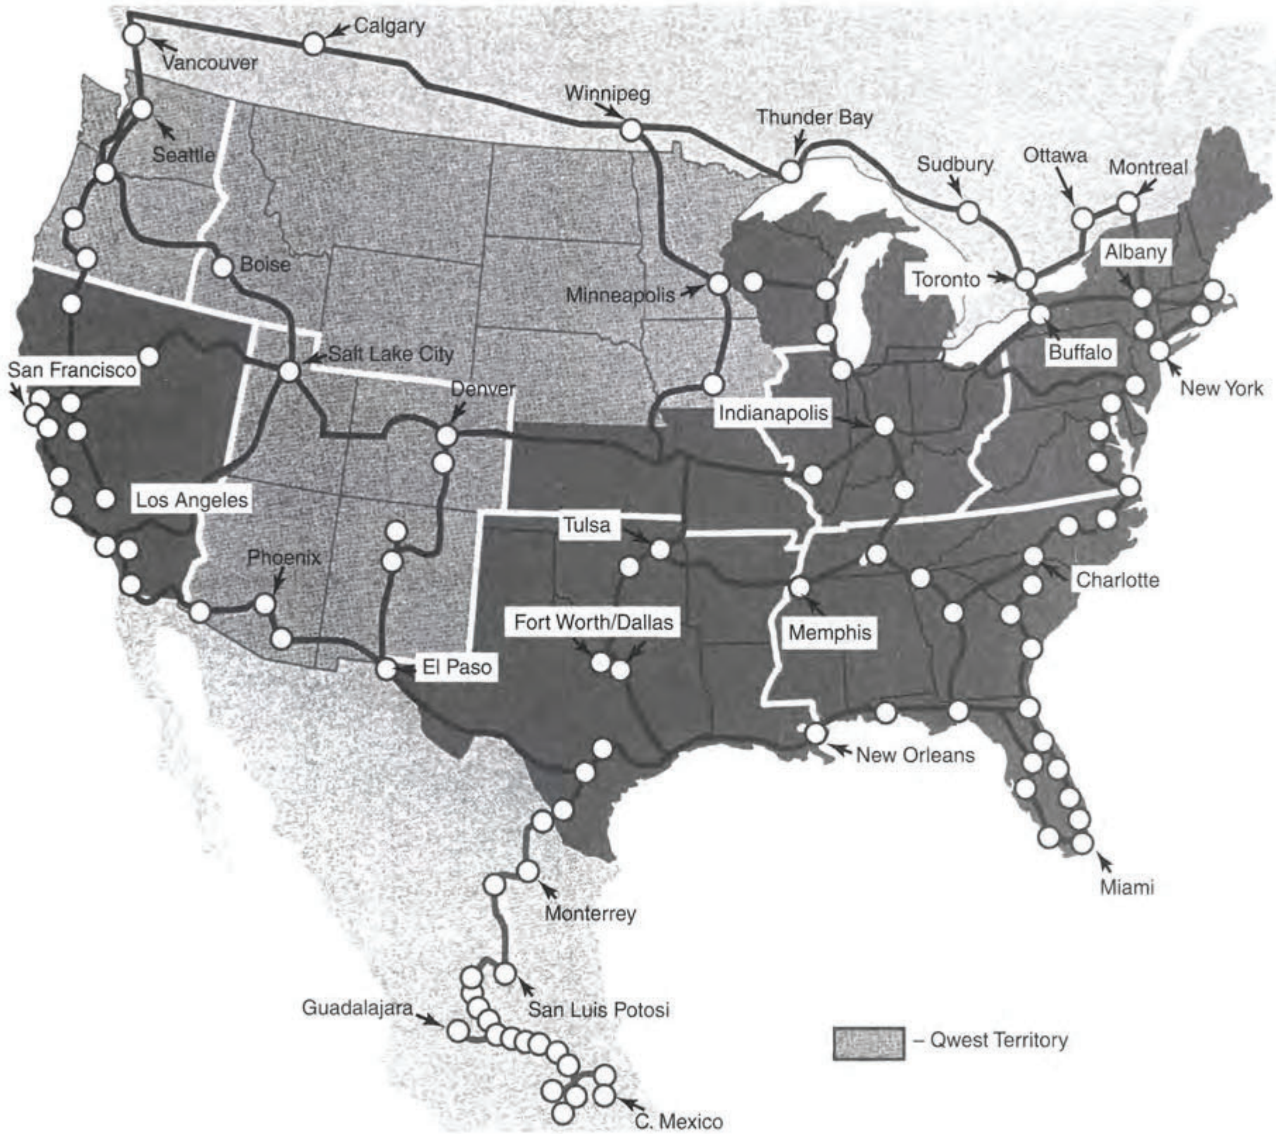

FIGURE 23.11

Qwest's North American fiber-optic backbone system. Nodes are farther apart in the company's local service area in the western United States. (Courtesy of Qwest)

differences are in the details. Terrestrial systems generally are part of a network mesh connecting major urban centers or telecommunication transport nodes, which are scattered across continents. This means that most long-haul terrestrial systems do not have to span the intercontinental distances of long-haul submarine cables. It also means that terrestrial networks often are installed piece by piece rather than as entire systems. Terrestrial environments are much more accessible, making repairs and powering amplifiers much easier.

This gives long-haul terrestrial telecommunications systems a distinct topology, as shown in Figure 23.11 for a representative network operated by Qwest. Like a national

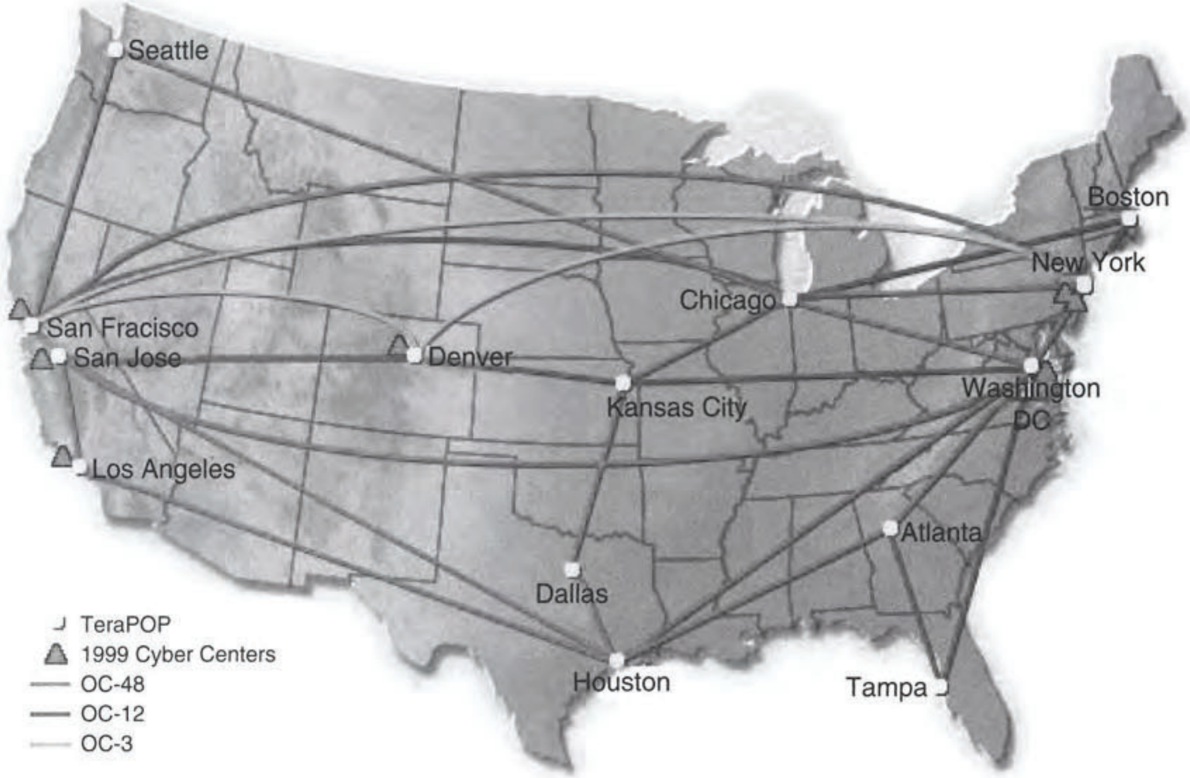

FIGURE 23.12
Qwest IP network in the United States. (Courtesy of Qwest)

railroad or interstate highway map, the main routes connect major population centers. In fact, many long-distance telecommunication lines run along railroad or highway rights of way. The nodes indicated are major interfaces with regional and local telecommunications networks.

It's worth comparing the long-distance backbone network with Qwest's Internet backbone, shown in Figure 23.12. The Internet backbone looks different because it represents Internet Protocol connections on a higher layer (layer 3) than the physical connections made through fiber in the long-haul network. The long loops between points such as Los Angeles and Dallas show that a connection is made between the two points, but don't try to show the physical route the signal takes. You can also see that IP signals may need to make two or more hops even between major hub cities. For example, messages routed from New York to Seattle would have to go through Chicago or San Francisco.

Transmission Rates and Distances

As you can see from Figures 23.11 and 23.12, long-haul terrestrial systems generally don't go as far between drop points or nodes as do transoceanic submarine cables.

Typical terrestrial backbone systems run a few hundred to a thousand kilometers between hubs or nodes where the signals are regenerated and redistributed. The Qwest long-distance network in Figure 23.11 shows two types of configuration. In the light gray area, where Qwest provides regional as well as long-distance telephone service, the cables run long distances between major switching nodes. In the darker gray area, where Qwest's main business is long-distance service, the cable has many add-drop points as well as major urban switching nodes.

These shorter cable runs relax many requirements on optical amplifiers, because noise, pulse dispersion, and differences in gain as a function of wavelength do not accumulate as much over such distances. Terrestrial cables may be able to use 25 to 30 dB of optical amplifier gain instead of the 10 dB limit in transoceanic submarine cables. Thus instead of one amplifier every 50 km, you could have one every 100 km or more. The shorter distances reduce the accumulation of nonlinear effects. The shorter distances also ease the requirements on gain uniformity, so terrestrial systems can use a wider range of wavelengths in WDM systems and thus transmit more optical channels.

The shorter spans of land cables relax requirements on optical amplifiers.

Terrestrial systems also can accommodate more fibers because they do not have the same stringent limits on electrical power and number of optical amplifiers as submarine cables. Terrestrial systems can obtain power locally and house amplifiers in buildings. These advantages allow terrestrial cables to carry signals on many more fibers than submarine cables, so terrestrial cables have much larger total transmission capacity.

Terrestrial cables have more fibers than do long-haul submarine cables.

The data rates on individual optical channels in long-haul terrestrial systems, like those in submarine cables, have increased steadily since the first single-channel long-distance fiber networks were installed in the early 1980s. The first systems transmitted 400 Mbit/s at 1300 nm. By the early 1990s, the state of the art in commercial systems was 2.5 Gbit/s at a single wavelength. Today, the state of the art in long-distance terrestrial systems is multiple optical channels per fiber, each transmitting 2.5 or 10 Gbit/s. Those data rates are achievable over several hundred kilometers using nonzero dispersion-shifted fibers for best WDM performance. Extensive dispersion compensation is used to avoid the need for regeneration at longer distances.

Long-haul transmission at 40 Gbit/s would be much more difficult because the dispersion limit on distance increases with the square of the data rate. Laboratory experiments have demonstrated long-haul transmission at 40 Gbit/s, but commercial interest has largely evaporated since the telecommunications bubble deflated. Virtually no one needs that much new capacity when plenty of dark fibers are available.

Peak transmission rates possible through today's fiber networks normally are reached only on the busiest routes. Figure 23.12 shows that only a fraction of Internet data traffic goes at the maximum 10-Gbit/s (OC-192); many lines are 155 Mbit/s, 622 Mbit/s, or 2.5 Gbit/s. The map shows Internet Protocol capacity in 2001, close to the peak of the bubble, but none of the routes required more than one 10-Gbit/s optical channel. Internet traffic has increased since then, but needs have been met by adding optical channels.

Only a few very busy routes carry peak data rates.

In principle it's possible to adapt transoceanic submarine cable technology for terrestrial systems that run across broad continents like North America. Boston is about as far from London as it is from Los Angeles. However, the current market is not willing to pay for that technology.

Long-Haul Network Connections

Long-haul land cables link to international and regional networks.

Long-haul terrestrial networks have two distinct types of connections. One is with other long-distance networks, such as international submarine cables. To make calls from Chicago to Berlin, for example, you need a terrestrial connection from Chicago to the landing point of a transatlantic cable, a submarine connection across the Atlantic, and a terrestrial connection from the European landing point to Berlin. These connections typically are on the east and west coasts and on southern borders, depending on where the traffic is going.

The other type of connections are to regional and metro telecommunication networks, which you will learn about in Chapter 24. These regional companies include not only the dominant local telephone companies, but also competitive local carriers, cell phone networks, cable companies, and Internet Service Providers. They may also include divisions of the long-distance carriers that are licensed to provide local service.

If you look closely at Figure 23.11, you will note that every major node in the United States is on a ring of cable. These are SONET-type rings, which provide insurance in case of equipment failures or cable breaks. For example, if a flash flood east of San Diego washed out the main cable from Los Angeles to Phoenix, traffic between the cities could be redirected through Salt Lake City, Denver, and El Paso.

Add-Drop Multiplexing and Wavelength Conversion

Long-distance terrestrial cables may add and drop signals at intermediate points.

Unlike long-haul submarine cables, long-haul terrestrial cables may need to make connections at intermediate points along their routes. Typically these are cities large enough to generate significant traffic, but not large enough to be hubs. This is done with add-drop multiplexers, which you learned about earlier.

Add-drop multiplexers can take various forms. They may be static, always directing signals in the same ways, or dynamic, able to switch signals in different directions. They also may split off the contents of an entire fiber, or individual optical channels in a fiber carrying WDM traffic. The choice depends on the type of system and the amount of traffic.

Normally signals are dropped at the intermediate location, and others added in their place. In WDM systems, this may require converting the wavelengths of some signals to wavelengths that are available in the through cable.

Wavelength conversion also may be necessary at hubs, where signals are switched in different directions and reorganized. For example, signals from both San Francisco and Seattle may reach Salt Lake City at 1540 nm, but only one 1540-nm channel may be available to Denver. One of those signals must be converted to a different wavelength.

Types of Long-Distance Services

So far I have concentrated on the technology of piping high-speed data over long distances. You may be wondering about the structure of the industry that handles the job. That structure has been changing thanks to new regulations that have broken up old monopolies, and the growth of many new companies. A few years back, you could separate carriers into local, long-distance, and international. Now this is no longer possible. Some international carriers also own regional telephone companies and provide long-distance service in the United States. Companies have been sold and merged at a dizzying rate.

This isn't a book about the telecommunications business, so I won't go into detail, but you should recognize a few distinct services:

- The *public switched telecommunications network*, which has grown from the telephone network to provide service on a call-by-call basis. You use it to make long-distance phone calls and to send faxes over phone lines. Generally long-distance calls pass through two or more companies, and you shouldn't notice the difference.
- The Internet, which transfers data packets among users around the world. Most of its long-distance traffic goes over a set of fibers dedicated to Internet transmission.
- Private leased lines, which are transmission capacity that businesses lease on fibers from carriers whose business is providing that capacity. This can get complicated because some carriers actually lease lines to provide part or all of their capacity. For example, long-distance calls from South Dakota to Minneapolis might go through a line that a long-distance carrier had leased from another company, which laid cable along the right of way of a gas pipeline. Sometimes carriers will even lease lines from each other to avoid the costs of building a pair of separate parallel transmission lines.

All these types of services also exist in regional telecommunications networks, covered in Chapter 24.

What Have You Learned?

1. Telecommunications encompasses voice, data, facsimile, video, and other forms of communications, which are carried by a global network that includes fiber-optic systems, satellites, and other media.
2. The global telecommunications network evolved from the telephone network. It connects local, regional, and long-distance networks so you can dial phones around the world. Voice phones, faxes, dial-up modems, and pagers share a common dialing system based on circuit switching.
3. The Internet was developed to connect computer networks. It transmits bursty digital data more efficiently than circuit-switched telephone lines. It uses packet switching.
4. Cable-television systems distribute video signals locally, and also carry Internet and telephone signals. Video signals also are distributed by direct broadcast satellites.

5. Optical fibers are the high-speed backbone of the global telecommunications system, which carries telephone signals at a hierarchy of standardized data rates that are combined by time-division multiplexing.
6. The lower speeds of the digital telephone hierarchy are based on the separate standards for time-division multiplexing used in North America, Europe, and Japan. Higher speeds are based on the SONET/SDH standards, which package input signals into digital frames that guarantee circuit capacity for telephone calls as well as computer data.
7. The Internet transmits signals at standard rates used in the global telecommunications network, but does not use time-division multiplexing to assemble them.
8. Users usually share transmission capacity on the same cable, the same fiber, or the same optical channel.
9. Long-distance fiber-optic cable capacity was overbuilt during the telecommunications bubble, leaving a glut of capacity.
10. Submarine fiber-optic cables are the backbone of intercontinental telecommunications. The first, using electro-optic repeaters, was laid in 1988 and is already obsolete. Current cables use optical amplifiers and have much higher capacity.
11. The undersea environment shapes the design of submarine fiber-optic cable. The cables must withstand high pressure, be extremely reliable, and have very high transmission capacity.
12. Unrepeatered submarine cables are the simplest types because they avoid submerged amplifiers. Onshore or remotely powered amplifiers can stretch transmission to reach 350 to 400 km. They typically link islands to each other or the mainland.
13. Repeatered submarine cables can span many thousands of kilometers. Limited electrical power and space in repeater housing restrict the number of active fibers. They often link continents.
14. Constraints on design of repeatered submarine cables are tight, limiting repeater spacing to about 50 km and requiring precise dispersion compensation. Their capacity has increased steadily.
15. Submarine cables can include undersea branching points, which divide cable capacity among two or more landing points.
16. Long-haul terrestrial cables resemble long-haul undersea cables, but differ in details because of their environments. Generally, long-haul land cables span shorter distances than submarine cables, and their network topology is different.
17. Terrestrial cables can accommodate more parallel fibers than can submarine cables.
18. Long-haul terrestrial networks link to regional networks and to international submarine cable networks.

What's Next?

In Chapter 24, you will learn about regional and metro telecommunication networks.

Further Reading

Cybergeography: *http://www.cybergeography.org* (map compilations)

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

Jeff Hecht, *City of Light: The Story of Fiber Optics Revised and Expanded Edition* (Oxford University Press, 2004). See the Epilogue.

Jeff Hecht, "Optical Networking: What's Really Out There?" *Laser Focus World 39* 2, 85–88 (February 2003)

International Cable Protection Committee: *http://www.iscpc.org* (tabulation of submarine cables)

Peter K. Runge and Patrick R. Trischitta, eds., *Undersea Lightwave Communications* (IEEE Press 1986) (design of first transatlantic fiber cable)

Patrick R. Trischitta and William C. Marra, "Applying WDM Technology to Undersea Cable Networks," *IEEE Communications Magazine 36*, 2, 62–66 (February 1998)

Questions to Think About

1. The data rate of a SONET OC-12 carrier is 622.08 Mbit/s. If that corresponds to 8064 voice channels, and each voice channel is 64,000 bits per second, how much of that signal is overhead?
2. SONET frames are a fixed length. Internet packets are a variable length, with the header indicating the packet size. How does this difference relate to the difference between circuit and packet switching?
3. TAT-8 transmits 560 Mbit/s and began operation at the end of 1988. Atlantic Crossing 1 began operation in 1998 transmitting 40 Gbit/s on a parallel route. How much did the data rate increase over that decade?
4. According to "Moore's Law" the capacity of integrated circuits doubles every 18 months. Judging from the increase in data rates on transatlantic cables in Question 3, what was the doubling time of fiber-optic capacity from 1988 to 1998?
5. Using the results of Questions 3 and 4, if fiber-optic capacity continues to expand at the same rate, what would the capacity of a transatlantic fiber-optic cable be in 2005? How does that compare with the Apollo cable as it began operation in 2003? How does it compare with the potential capacity if all possible optical channels were activated?

6. In 1983, the peak data transmission rate in a commercial terrestrial fiber-optic cable was 400 Mbit/s on a single fiber. In 2000, manufacturers claimed systems they offered could transmit 160 optical channels at 10 Gbit/s through a single fiber. How much of an increase is that and what is the doubling time?
7. In 2003, the laboratory record for highest data rate through a single optical fiber was 10 Tbit/s. If the state of the art in 1983 was 400 Mbit/s, how much of an increase is that and what is the doubling time?

Chapter Quiz

1. What type(s) of signals travel in the global telecommunications network?
   a. voice telephone
   b. digital data
   c. facsimile
   d. video
   e. all of the above
2. What principle is used to combine telephone signals as they enter the global telecommunications network?
   a. packet switching
   b. frequency-division multiplexing
   c. digital-to-analog conversion
   d. time-division multiplexing
   e. wavelength-division multiplexing
3. How does the Internet relate to the global telecommunications network that evolved from the telephone system?
   a. The two interconnect, and both carry digital data along separate paths.
   b. The telephone network carries only analog signals; the Internet transmits only digital data.
   c. The two are identical.
   d. The Internet has replaced the global telecommunications network.
   e. Only the Internet can carry packet-switched signals.
4. A single voice channel in the North American Digital Hierarchy corresponds to a speed of
   a. 4000 Hz.
   b. 4000 bit/s.
   c. 14,400 bit/s.
   d. 64,000 bit/s
   e. 1.5 Mbit/s.

5. How many T1 signals go into a SONET OC-3 signal?

a. 84
b. 96
c. 672
d. 2016
e. 155 million

6. Which of the following signals can feed a SONET OC-3 system?

a. ATM format
b. packet-switched Internet Protocol
c. T3 from the Digital Telephone Hierarchy
d. multiple T1 circuits
e. all of the above

7. How do unrepeatered submarine cables differ from repeatered cables?

a. Only repeatered cables contain electro-optic repeaters.
b. Unrepeatered cables do not include any optical amplifiers.
c. Unrepeatered cables do not include any optical amplifiers powered by pump lasers under water.
d. Unrepeatered cables can transmit signals farther.
e. Unrepeatered cables include copper wires, which transmit electrical power.

8. Unrepeatered cables *cannot* be used for which of the following?

a. festoon systems along the coast of a country
b. transatlantic cables
c. links across the English Channel
d. island-hopping systems
e. a cable crossing 200 km of ocean

9. Which of the following techniques can stretch transmission distance of an unrepeatered submarine cable?

a. optical postamplifier onshore to boost output power of the transmitter
b. remote optical amplifier powered by light from an onshore pump laser
c. preamplifier onshore to boost input power before the receiver
d. a and c
e. all of the above

10. Which of the following factors limits the number of fiber pairs usable in a repeatered submarine cable?

a. fiber attenuation and dispersion
b. the need for dispersion management on all fibers
c. limited space in repeater housings and limited electrical power for optical amplifiers

d. The core of the cable cannot be larger than a certain size.

e. the number of optical channels used for wavelength-division multiplexing

11. A submarine cable has four fiber pairs, each carrying 2.5 Gbit/s at each of four wavelengths. What is the cable's total data rate?

a. 2.5 Gbit/s

b. 10 Gbit/s

c. 40 Gbit/s

d. 80 Gbit/s

e. 160 Gbit/s

12. Why is the amplifier spacing in a transatlantic fiber-optic cable limited to 50 km?

a. to limit the noise and unequal amplification of optical channels that accumulate over transatlantic distances

b. Electrical power for pump lasers is not available under water.

c. Water pressure reduces the gain that can be achieved.

d. Water pressure increases fiber attenuation.

e. all of the above

13. Why would different types of fiber be used in the same submarine fiber-optic cable?

a. to reduce attenuation

b. to compensate for chromatic dispersion

c. to increase power levels that can be transmitted

d. to balance attenuation at different wavelengths

e. because the factory ran out of the first type partway through the cable

14. All WDM channels in a 5000-km fiber system must have power within 6 dB of each other in order for the system to operate properly. If amplifiers are spaced every 50 km, how uniform must their gain be across the WDM spectrum, assuming differences accumulate uniformly over the cable length?

a. 0.05 dB

b. 0.06 dB

c. 0.12 dB

d. 0.5 dB

e. 0.6 dB

15. How does the terrestrial long-distance network differ from repeatered submarine cables?

a. Terrestrial systems do not require optical amplifiers.

b. Terrestrial systems form a grid interconnecting many points on land.

c. Terrestrial systems connect termination points on the coast.

d. Terrestrial systems do not require electrical power.

e. Only terrestrial systems have branching points.

CHAPTER

24

Regional and Metro Telecommunications

About This Chapter

Regional and metro telecommunication systems bridge the gap between the local networks that serve your community and the long-distance telecommunications that form the backbone of the global network. Their function is to collect signals from points across the region and distribute them within the region and to major nodes of long-distance systems. In general, the operation and design of regional and metro networks fall between those of long-distance and local systems: They typically span shorter distances and carry lower-speed signals than long-distance backbone networks, but span longer distances and carry faster signals than local networks. They also differ in connectivity from long-distance and local networks. In this chapter you'll learn about these differences and how regional and metro networks work.

Defining Regional and Metro Telecommunications

The distinction between the terms "regional" and "metro" can be subtle, and may owe more to history and marketing than to different network architectures. Wire-line telephone charges traditionally were split between long-distance calls, listed individually on the phone bill, and regional calls, billed in increments or at a flat rate. The first breakup of AT&T in 1984 split the company between long-distance service (which remained AT&T) and the original seven Bell Regional Operating Companies, which offered services within their own regions. Long-distance and regional services were differentiated on the basis of the area codes existing at the time of the breakup, which have since been subdivided in most states. Those regional phone companies—and competitive carriers formed later—now operate regional networks in their own operating areas.

Metro networks originated more recently, when other companies decided they wanted to carry signals within regional areas. By that time these signals included digital data and video as well as telephone traffic. Not surprisingly, these networks were concentrated in heavily populated metropolitan areas, where demand for transmission capacity was high. As the telecommunications bubble swelled, metro networks were much hyped. The term has survived partly because it sounds more modern, but also because it emphasizes that these networks serve a highly populated metropolitan area with a lot of traffic heavy with data.

Both regional and metro networks consist mostly of fiber. Functionally their main distinction is in how they connect the locations they serve, but this difference is vanishing as networks evolve. Regional and metro networks use essentially the same hardware, and from this point on we'll cover the two networks together, noting differences explicitly.

Regional Distribution

Regional networks are distribution systems.

Regional and metro networks are distribution systems, not information pipelines like links in the long-distance network. Regional networks connect many points; the long-distance system carries large traffic volumes over long distances.

A regional telecommunication system plays the same role as a grid of major streets, which links neighborhood roads to limited-access expressways. Figure 24.1 shows how a regional network in a rural area might look. The network is a grid of cables connecting telephone switching offices in each town. Low-capacity cables (thin lines) run between small towns, and from small towns to their larger neighbors. Larger-capacity cables (thicker lines) connect the larger towns to the regional hub, a small city, where they connect with long-distance lines.

In a sense, this network is a hub-and-spoke system, with high-capacity links spreading from the regional hub to larger towns, and those larger towns, in turn, connecting to small towns. However, the system also includes cross-links between towns, which form rings like those used in SONET systems. If you cut any one cable in the network of Figure 24.1, the towns on both sides of the break would still have connections to all the other towns in the region.

Regional phone networks connect local phone switches with each other and the long-distance network.

Real networks may have more levels of hubs from which lines branch out to smaller hubs, with regional rings providing a backbone to serve the large hubs. Figure 24.2 shows one such network built in Oregon, which covers both the Portland metropolitan area in the northwest and the rural areas in the east. It delivers service to county seats and major towns, connecting to other networks (not shown) that form grids throughout the state. This is only one of the networks serving the region; a separate regional network links telephone switching offices throughout the state, but it's operated by another organization. (Remember that, as with long-distance networks, fibers carrying signals for two different companies may run through the same cable.)

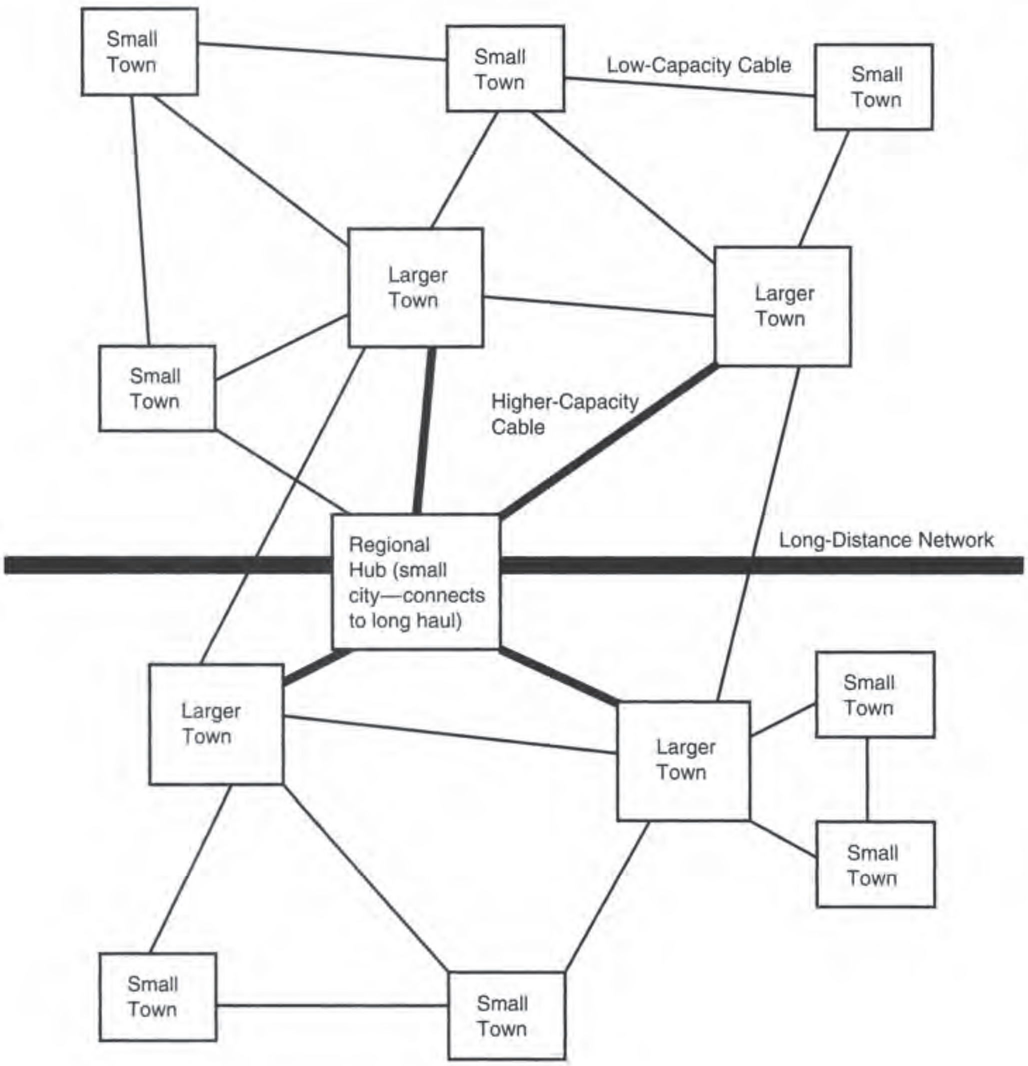

FIGURE 24.1
Regional telecommunication network in rural area.

Regional networks connect to local phone lines at telephone switching offices, which are the hubs for local phone service. As you will see in Chapter 25, you can trace all local phone lines to a switching office or central office, where they make connections with other local phone lines and with regional and long-distance networks. Regional networks also connect with other local services, including cable television companies, mobile phone systems, and Internet Service Providers that rely on telephone lines for connections with the Internet backbone.

Regional network concepts also apply to urban and suburban areas, but in that case the transmission lines link adjacent suburbs, and feed into the urban center, where they connect with long-distance lines. The result is a network that looks like a more elaborate version of Figure 24.1, but with the nodes labeled suburbs and city center. The city center is the hub, with individual suburbs on spokes, and cross-links between adjacent suburbs.

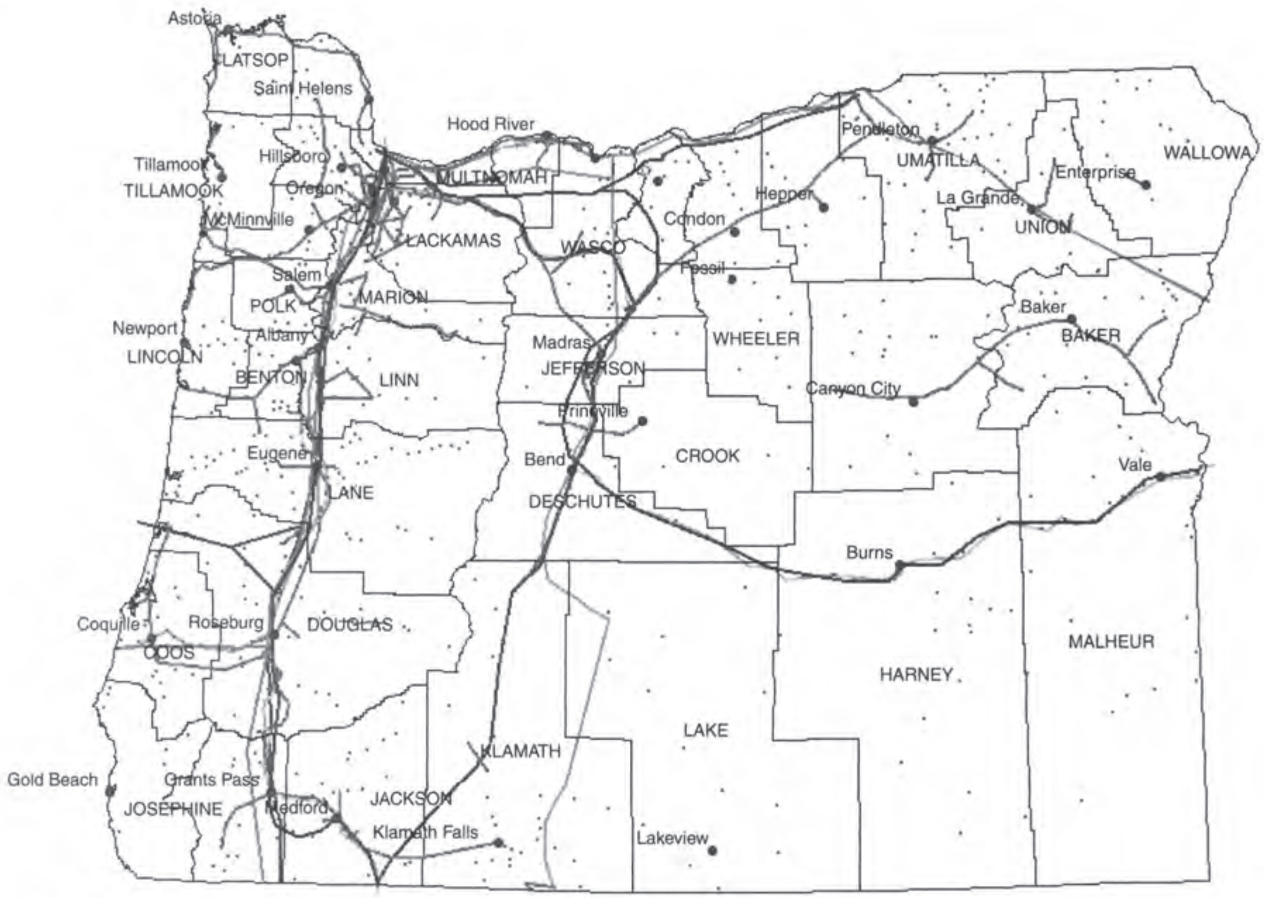

FIGURE 24.2

State of Oregon Enterprise Network provides links throughout the state. (Courtesy of Oregon Economic & Community Development)

Regional phone networks are meshes linking switching offices.

The topology of a network also depends on its history. Established regional telephone companies built their networks piece by piece over many years. As communities grew and traffic increased, the companies expanded their transmission capacity and added new lines. Their networks are meshes that reach throughout the region, but they make their major connections at local telephone switching offices, where signals from each community are collected and organized for transmission to other communities.

New metro networks built by different companies can have different structures. Figure 24.3 shows a metropolitan network built by one company that offers transmission service to everyone in the area. This system is a ring that provides service to many organizations along its route, making links between different points. The dashed line shows a link that carries video signals from a local television station to a cable company's distribution center. Other links could run between city hall and the county office complex at the bottom, and between the mobile phone switching center and the suburban phone company at top. A metropolitan network makes many connections besides those to traditional telephone companies.

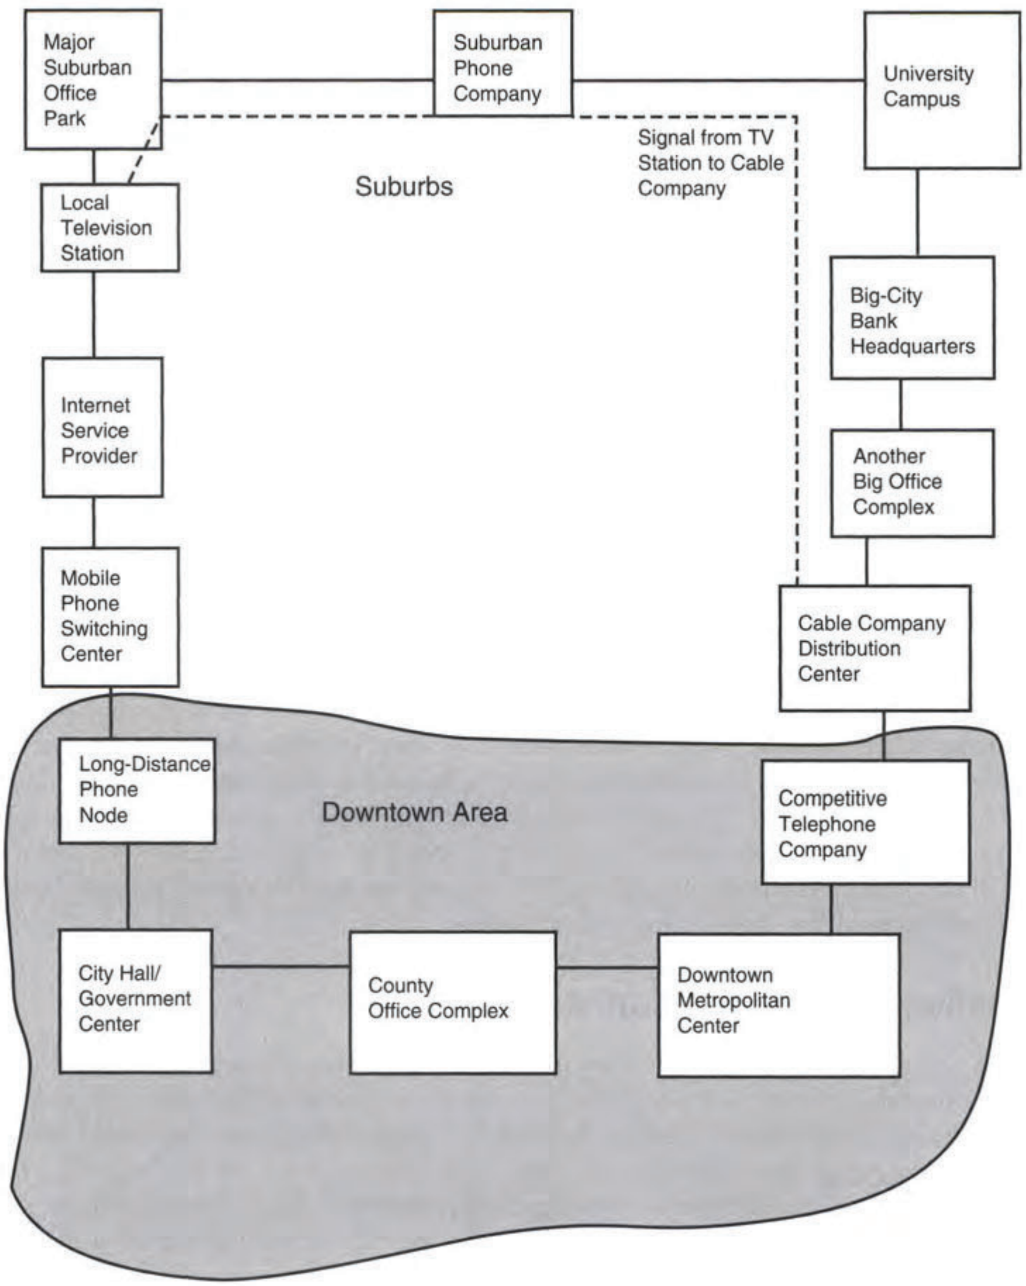

FIGURE 24.3
A metropolitan ring network links many sites.

Other companies may build rings serving other parts of the metropolitan area. These companies can lease capacity to anyone willing to pay for it, including regional phone companies that may decide it's cheaper to rent rather than lay cable along a similar route. The rings and meshes built by all the companies operating in the area combine to form a regional or metropolitan network, carrying a variety of services. However, these networks are not integrated seamlessly, and may operate quite differently, as you will learn later. Let's start by looking at the established regional telecommunications networks.

Regional Telecommunications Networks

Regional phone systems are part of the public switched telephone network.

The well-established regional networks built by telephone companies are part of what is called the *public switched telephone network (PSTN)*, which collects, packages, switches, and distributes telephone signals. This traditional network is a circuit-switched system, which creates voice circuits between individual phones. The switching and packaging are done at local *switching offices*, which digitize voice signals and time-division multiplex the bit streams to create signals at the hierarchy of higher speeds described in Chapter 20. The regional network carries these signals between central offices or from a central office to the long-distance network. You'll learn about connections from the central office to individual phones in Chapter 25.

Recall that this regional phone network is a distribution system with a high level of interconnections that function as a mesh linking many points. It's also a public network, which essentially rents the use of a voice line in time increments to anyone with a telephone connection.

Regional networks also can carry other types of signals. For example, a business may want its own 45-Mbit/s T3 line to carry voice and data signals between its facilities, or may rent a line to an Internet node. The network hardware doesn't care where the signal came from or how it was produced; it just requires a standard data rate along the line.

Historically each regional phone network was a monopoly without competition throughout its region. Now competition is common in urban and suburban areas, although the individual phone user sees it only in the local and long-distance markets. Like long-distance networks, regional networks often buy the right to transmit signals through fibers in cables owned by other companies. Competition to provide local phone service further complicates the picture, although in practice competitive carriers generally lease regional transmission facilities from other carriers.

Regional Telephone Connections

The traditional telephone system was strongly hierarchical, and that legacy remains strong in the circuit-switched phone systems. The hierarchy organized central offices depending on the sizes of the communities they serviced. The central offices that served small towns linked to larger towns, which in turn linked to still larger ones. Only the largest central offices had direct connections to the long-distance network. Today's networks are not as rigidly hierarchical, but that structure is evident in the rural network of Figure 24.1.

Trunk lines connect central offices.

The connections between central offices are called *trunk lines* and today virtually all of them are fiber-optic. In metropolitan or suburban areas, trunks typically run several kilometers between central offices in adjacent communities. Calls made within a region may go through a series of trunk lines. The capacity required depends on the volume of calls on that route.

Legacy and New Networks

Regional networks are upgraded gradually to higher capacity.

Regional networks inevitably contain a mixture of old (*legacy*) and new equipment. The old stuff generally isn't as old as it is in local phone systems, but regional carriers generally do not replace systems in good working condition as long as their transmission capacity is adequate. Normally they replace one segment of the network at a time, both to limit expenses and to avoid disrupting service.

Replacements generally occur where old equipment fails or more capacity is needed. Carriers try to avoid new construction wherever possible, so they may prefer to lay new cables in existing underground ducts, or add wavelength-division multiplexing and new optical channels to existing systems transmitting only one wavelength.

Generally, regional networks require less transmission capacity than long-distance systems. Thus regional networks make much less use of WDM technology than long-haul systems. Upgrades also are not evenly distributed. Although users are steadily increasing the demand for transmission capacity, the growth is not even. New capacity is most likely to be needed in areas where new development has outstripped existing systems.

Other Connections

Regional networks interconnect with other systems.

The regional networks built for fixed ("wire-line") telephone signals now carry a variety of other signals and must connect with other equipment. Mobile telephones have become common, so regional networks must connect with cell-phone towers and switching centers; this means links in each region ("cell") for each mobile phone carrier serving the region.

Competitive telephone companies also require connections. Many of them use the same local wiring as the dominant local carrier, but have separate switches, either at the dominant phone company's central office or another site. Regional networks also have to connect to telephone services provided by local cable-television companies.

Data services also require special connections. Internet Service Providers need to lease lines to make connections to the Internet. So do data-processing and computer centers at large companies and universities. Large companies also may lease lines to connect with other facilities in the area, which don't have to go directly through the switched telephone network.

Transmission Requirements

Regional networks transmit at a wide range of data rates.

Transmission requirements in regional networks cover a much broader range of data rates and distances than in long-distance systems.

The economics of long-haul transmission dictate that signals be packed as tightly as possible. You trade off the cost of the equipment to package signals against the cost of transmission. That means you have to amortize the extra money you spend at the transmitter and receiver against the money you save on fiber, cable, and optical amplifiers. If you have to send signals 5000 km across the ocean, you can afford to spend much more at the transmitter and receiver end than if the signal is only going 50 km.

Individual telephone circuits are grouped together for efficiency, but the degree of grouping depends on the traffic level. Small towns may generate only enough traffic to fill a few 1.5-Mbit/s T1 lines on the busiest days. Larger towns will generate signals at the 45 Mbit/s T3 rate or higher. Regional networks carry signals at data rates from the 1.5 Mbit/s T1 speed up to 2.5 or 10 Gbit/s, depending on customer requirements. For long-haul systems, it usually makes sense to time-division multiplex slow signals, combining multiple inputs to make a single higher-speed data stream. Regional networks usually multiplex to lower rates.

Metro Networks

Metropolitan networks are similar to regional networks, but specifically serve metropolitan areas with large concentrations of people and industry. There are three basic variations on the concept, and the distinctions can sometimes be hazy:

- *Metro telephone/telecommunications networks* developed by expanding the regional networks of telephone carriers in a metropolitan area. They are basically enhanced regional networks that carry services needed in a metropolitan area. They tend to carry heavier traffic than regional networks and have somewhat different structures.
- *Wholesale metro networks*, run by individual companies, sell service directly to other companies rather than operating a public telephone service.
- *Metro-area networks* (MANs) and wide-area networks (*WANs*) typically serve a single organization spread throughout a metropolitan area, such as a city or county government or a large corporation.

The distinctions between the first two types of networks are caused mostly by differences between the type of carriers that operate them. A metro telecommunications network is run by a company that serves many customers in the region and provides retail public switched telephone service. In general, it has many points of connectivity in the region, as shown in the regional network of Figure 24.1. Many of these connections are made through telephone switching offices.

In contrast, a wholesale metro network generally does not provide retail public switched telephone service, although it may sell bulk capacity directly to large companies. This means it has fewer customers, and thus needs fewer connections through the region. Wholesale networks also tend to have fewer connections throughout the region, and may take the form of a simple ring, as in Figure 24.3. They may have connections to telephone switching offices, as well as other types of connections.

A metro telecommunications network has extensive interconnectivity.

Another important distinction is the degree of connectivity. A metro telecommunications network has extensive interconnection capability because it was built to serve telephone-like functions; each node, as in a switched telephone network, must be able to reach every other node. In contrast, wholesale metro networks only make the connections specified by their customers. You can see the difference by comparing the regional network of Figure 24.1 with the metro ring in Figure 24.3. The regional network can link any pair of towns on the map, although the connection may go through several other towns; however, many customers of the metro ring transmit only between a pair of points. The suburban telephone switch and the bank headquarters in Figure 24.3 don't receive the television signals that go past them on the ring.

Let's look at a couple other key features of these networks.

A metro core ring links metro access loops serving parts of a metro area.

Metro Network Topology

Regional networks have moved away from the hierarchical structure of older telephone networks, in which smaller switches directed signals through larger switching offices. A modern metro telecommunications network in a large city looks more like Figure 24.4, with

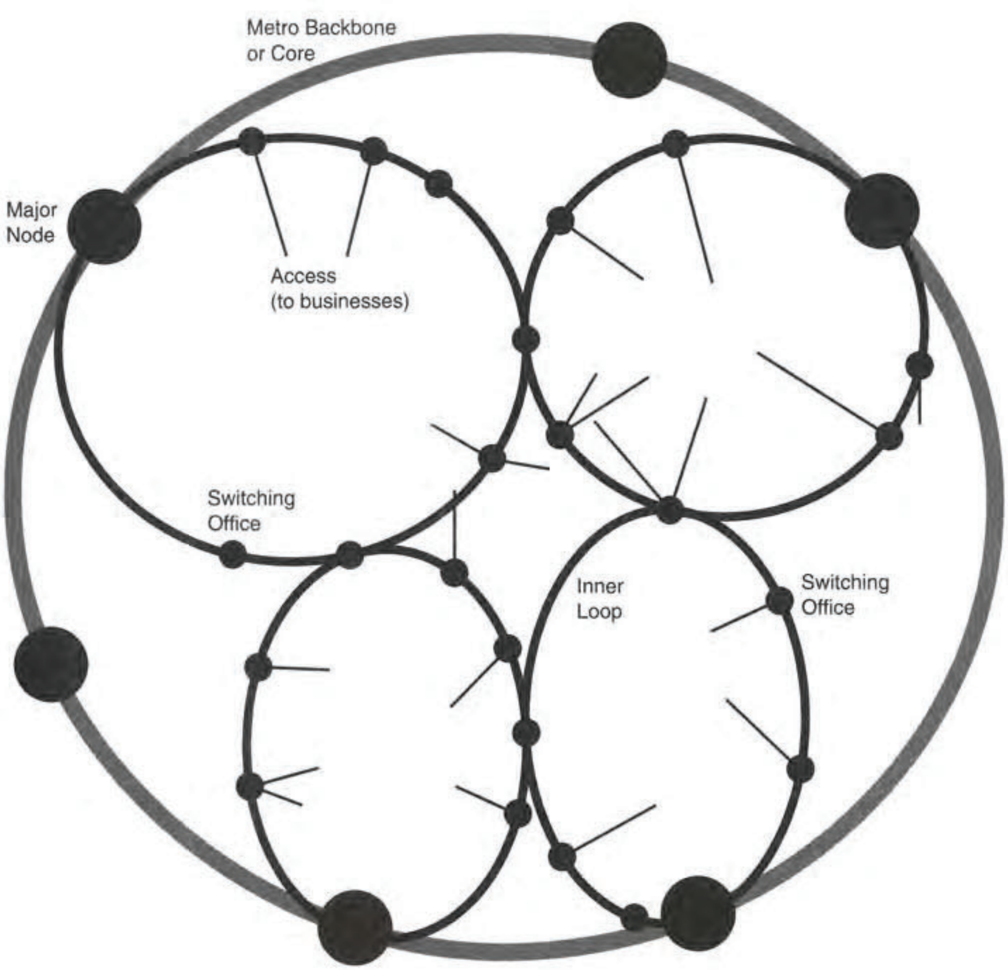

FIGURE 24.4
Metro core and access network in big city.

connections arranged in two sets of rings. The thick outer line is the *metro core* loop, which connects the telephone switches that carry the largest volume of traffic. Nodes on the metro core connect to *metro access* loops, which connect to smaller switching offices that serve parts of the service area. In Figure 24.4 these show as inner and outer loops; they may be arranged differently, but they are always form rings to maintain service in case of a single-point failure.

The same approach could be used in a regional network. In that case, the core connection might be a ring that serves small cities, while the access connections would be the rings connecting the small towns. This highlights an important point—the types of service delivered are as important as the area where they are delivered.

The metro or regional network provides interconnections between two points on the *network edge*. For a telephone network, these points are central offices that make switched connections to individual telephones. For leased lines, they are individual company facilities dedicated to making the same point-to-point connections.

Switching offices are at the network edge.

The network edge or outer part of the metro network is also called the *access network*. In practice, this refers to a connection to a central office or to an organization that buys transmission capacity on the metro network. For example, the suburban telephone office, the big-city bank headquarters, and the university in Figure 24.3 are all part of the access network.

MANs and WANs

MANs and WANs connect organizational LANs.

Both metropolitan telecommunication networks and wholesale metro networks sell transmission capacity to retail or wholesale customers. MANs and WANs do not; they are nominally connections between local area networks (LANs), which organizations use to connect their computers. MANs and LANs often buy or lease capacity on metro networks.

The Internet sometimes is considered a wide-area network, but we won't use that term here because it might muddle the picture. In practice, Internet links run through the same cables as other metropolitan transmission lines but are operated by independent carriers and Internet service providers, which agree to exchange signals with each other. You'll learn more about the Internet and data transmission in Chapter 26.

Regional/Metro Services and Equipment

Regional and metro networks are intermediate between long-haul and local systems in terms of size and traffic volumes. This is evident both in the services they carry and the equipment they use.

Rural regional networks span more territory than do metro networks.

Metro networks usually span no more than about 200 kilometers, and most metro loops are shorter in total length. This means that they require few if any optical amplifiers, except to compensate for lossy components, so they can use wavelengths for which amplifiers are not available. Regional networks can be considerably larger. Functionally a regional network could cover an entire area code, which can span several hundred kilometers in large sparsely populated states, so it may require more amplifiers.

Traffic volumes depend on the population of people and industries. Population density in the United States varies from less than 10 people per square mile in the most sparsely populated states to thousands of people per square mile in major urban areas. Information-intensive industry is concentrated in densely populated areas, so metro networks need much higher capacity than do regional networks in rural areas. Thus unlike long-distance networks, regional and metro networks generally don't have to provide both high transmission capacity and long transmission distances. Traffic volumes are still growing, so metro and regional networks are not as overbuilt as long-distance systems.

Services and Transmission Speeds

As mentioned earlier, metro and regional networks carry both public switched telephone traffic and point-to-point connections between company facilities on leased lines. Thus these networks carry a variety of services at different speeds.

Public networks combine input signals into high-speed data.

Public telecommunications networks are built on the assumption that input from subscribers comes in small chunks, which are combined into larger data flows in a regional or metro network. Public networks combine voice channels using time-division multiplexing to make higher-speed signals. Metro and regional networks see only the combined signals.

Metro and regional networks also receive signals directly from their customers in a variety of formats, depending on customer requirements. Network operators may transmit the signals in the raw input format—a practice called providing "dark fiber" because the customers "illuminate" it with signals in the transmission format they prefer. Alternatively, network operators may merge input signals from multiple access customers into a composite signal at a higher speed, as is done with telephone signals. Signals traveling through the system may come from different standard layers, and the carrier may repackage the signals for transmission on a different layer.

Transmission speeds for regional and metro networks cover a wide range. Small towns have a few T1 connections, larger towns have T3 lines, cities have OC-3 or OC-12 links, and so on up the data-rate hierarchy. The highest rates in regional and metro networks are usually 2.5-Gbit/s OC-48 lines, although higher rates are used in special cases. Generally, most of the transmission lines are at lower speeds, with the high-speed lines concentrated in the most densely populated areas. Metro networks tend to have higher-speed inputs, with switching offices providing input at OC-3 rates or higher.

Regional links carry data at speeds to 2.5 Gbit/s

Because metro and regional networks carry signals shorter distances than long-distance systems, designers face different trade-offs in selecting transmission speeds and protocols. Expensive transmitters and receivers can be justified in long-haul systems because they can send signals a long distance. However, it's harder to justify purchasing expensive terminal equipment—like time-division multiplexers and demultiplexers—to span the shorter distances in metro networks. That is, for short networks adding extra transmission lines to carry low-speed signals may cost less than installing the electronics needed to time-division multiplex the input to higher speed for transmission over one fiber.

Network Access

There are two distinct ways of accessing a metro or regional network, through a hub and through an add-drop multiplexer. We have glossed over these differences so far, but they impact system function.

A *hub* is a point where all (or most) of the signals in a system are switched and organized. Hubs include local switching offices and correspond to the terminal points on submarine cables or long-distance systems. In the metro network of Figure 24.5, there are two hubs: the network operation center at the top and the urban telephone switching office at the lower right. Other metro networks may include several switching offices.

Hubs or add-drops provide network access.

Add-drops are points where some signals are picked up and others dropped at an attached node. An add-drop multiplexer along the network diverts only part of the signals to the node. In Figure 24.5, the only signal picked up from the television station is one video signal directed to the cable-television company, which will distribute the signal to its subscribers. The figure also shows a two-way link between City Hall and an Internet Service Provider, which operates the city's Web site and provides officials with Internet access.

The metro network in Figure 24.5 is a ring that contains many add-drops and a few hubs. This means it functions largely as a network, interconnecting many nodes and directing signals between pairs of nodes. In a metro network, the signals often are transmitted separately, at different wavelengths or on different fibers. They may pass through add-drops, but not through the nodes attached to the network. In this they differ from

Metro networks typically are rings with many add-drops.

FIGURE 24.5
Hubs and add/drop nodes in a metro network.

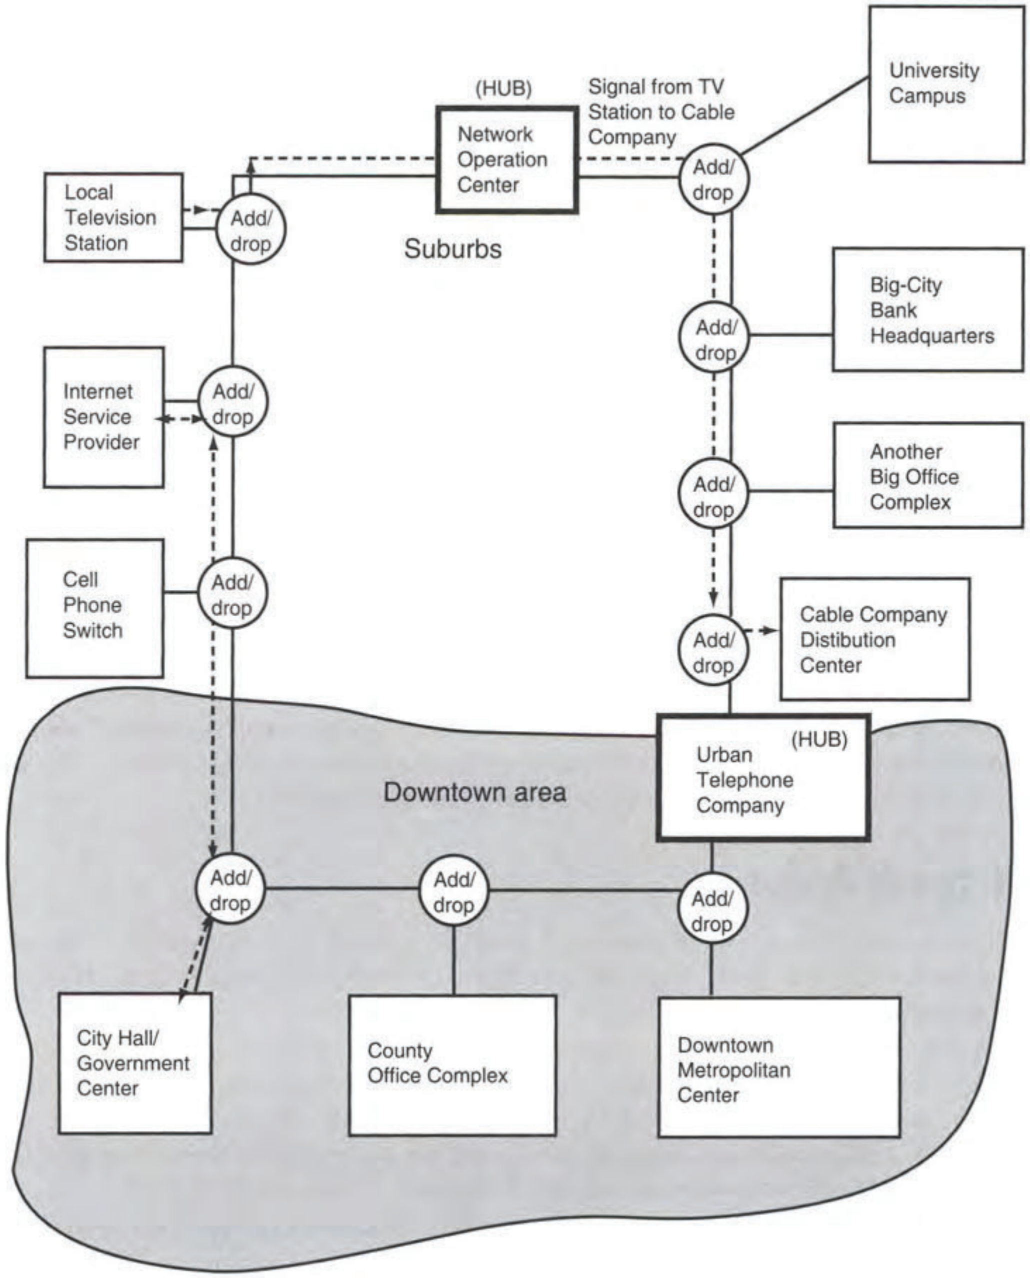

other networks, which route a combined signal through all nodes, but each node only receives the signals directed to it. This subtle distinction can be important to some users, such as banks, that are particularly concerned with security.

Other metro and regional networks may contain mostly hubs. In the rural regional network of Figure 24.1, you could consider all the towns to be hubs, some smaller than the others. You can also think of a regional telephone network as an array of point-to-point links, with each terminating at a telephone switch in a different location. Telephone companies often treat their regional or metro networks as a grid of point-to-point links, so they upgrade parts at different times.

Capacity and WDM in Metro Networks

Metro and regional links require lower transmission capacity than long-haul links, a fact that has important consequences. Network operators have installed metro cables with fiber counts up to 400 pairs in highly developed areas, but the cable installations generally were much less costly, and the resulting capacity was not as far beyond real requirements as for long-distance systems.

The modest capacity requirements mean that a single channel at 1310 nm can meet the needs of most links in metro and regional networks. This also reduces terminal equipment cost by avoiding expensive WDM optics and costly transmitters in the 1550-nm region.

Most metro links carry one channel at 1310 nm.

Wavelength-division multiplexing may make sense for metro and regional links carrying heavy traffic if no additional fibers are available on the route, the time needed for new construction is excessive, or the cost is higher than installing WDM equipment. Coarse-WDM is preferred over dense-WDM if the transmission distance is short enough to avoid using optical amplifiers and large numbers of channels are not needed. Widely spaced CWDM channels cost less to install than narrow-line DWDM channels, and cover a broader range of the spectrum. Low-water fibers allow use of more wavelengths than older standard single-mode fibers.

Metro systems use CWDM rather than DWDM.

The less-stringent requirements on metro and regional transmission have also reduced the layers of interface electronics. The rationale is that the cost of the electronics needed to convert signals to more efficient formats may be greater than the gain obtained from the higher efficiency. This has led to the development of *protocol agnostic* systems in which carriers assign customers their own optical channel or fiber to use as they see fit. Instead of trying to pack as many channels as possible onto each fiber, the operator reduces the number of layers to cut costs. Figure 24.6 shows an example, where four customers pick different formats: a 2.5-Gbit/s OC-48 line, digital video, Gigabit Ethernet, and Fibre Channel. Using this approach, a video studio can transmit its signal to a cable operator in video format without having to convert the signal.

The shorter links in metro and regional networks also reduce limitations imposed by dispersion and nonlinear effects on system operation. Recall that dispersion is a cumulative spreading of signal pulses that limits maximum transmission speed, so the shorter the distance spanned, the higher the maximum speed. Although very few metro or regional links require 40-Gbit/s transmission, this speed is easier to achieve in metro networks than in long-haul systems.

Shorter spans suffer less dispersion and can transmit higher speeds.

Nonlinear effects also accumulate with distance, so their impact decreases in metro and regional networks. Reducing the cumulative impact of nonlinearities in turn allows transmission of higher optical powers, providing more power per channel in WDM systems. Thus potentially more channels could be transmitted, although at this point there is little interest in increasing the number of optical channels in metro networks.

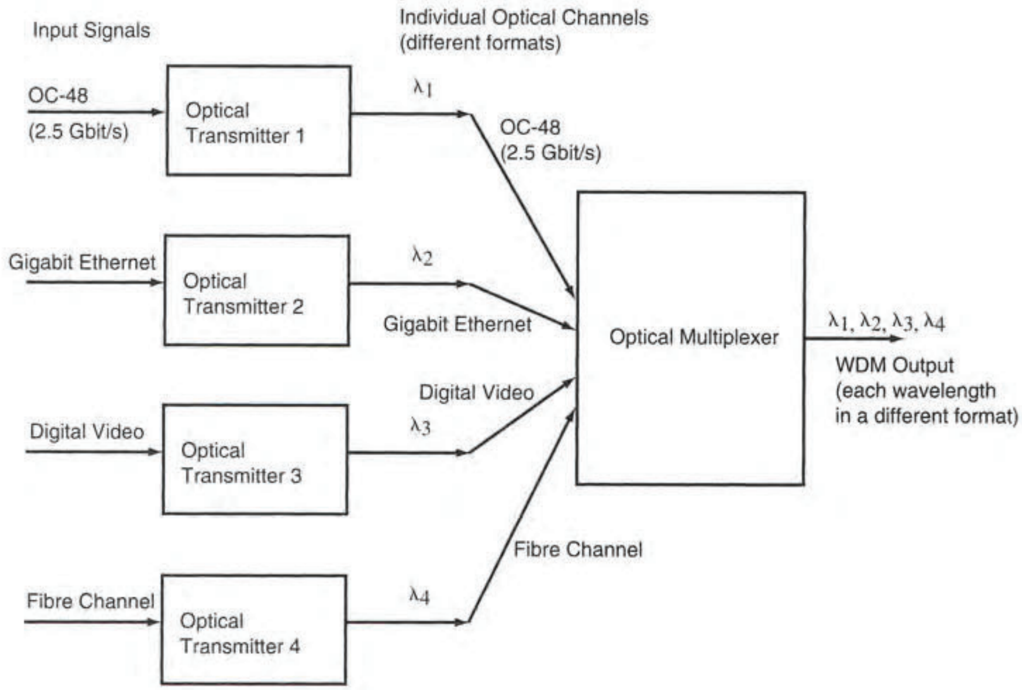

FIGURE 24.6 *Metro WDM system transmits optical channels in different formats.*

What Have You Learned?

1. Metro and regional networks are distribution systems, not information pipelines. They serve similar purposes, but differ somewhat in detail.
2. Regional phone networks connect local phone switches with each other and long-distance networks. They operate on different levels, ranging from low-capacity connections to small towns to high-capacity connections between suburbs and big cities.
3. A metro network interconnects many points in a metropolitan area. These include telephone switching centers and other facilities that require high-speed connections, such as large businesses, cable television companies, universities, and government offices.
4. Regional networks link switching offices; they tend to concentrate traffic to the largest switches. They transmit signals at a wide range of data rates.
5. Originally built for fixed telephone service, regional and metro networks now carry a variety of other signals and must connect with other equipment, including pagers, cell phones, Internet Service Providers, and competitive phone companies.
6. Regional phone systems are part of the public switched telephone network, which collects, packages, switches, and distributes telephone signals. They carry signals at a hierarchy of data rates.

7. Trunk lines run between switching offices. Switching offices are on the edge of the phone network.
8. Metro telecommunications networks are enhanced regional networks serving metropolitan areas. Like regional networks they have extensive connectivity. Wholesale metro networks sell service to other companies rather than provide public telephone service, and generally have fewer connections.
9. A metro core ring links metro access loops serving a metropolitan area.
10. Metropolitan-area networks (MANs) and wide-area networks (WANs) connect local-area networks (LANs) operated by an organization.
11. Rural regional networks span more territory than metropolitan networks but carry less traffic.
12. Regional and metro networks generally carry signals at up to 2.5 Gbit/s, although higher speeds are technically possible. Traffic demand is lower than in long-haul systems.
13. Network access is provided through hubs or add-drop multiplexers. Many metro networks are rings with many add-drops and few hubs. Regional networks typically have many hubs.
14. Most metro and regional links transmit one channel at 1310 nm and are not long enough to require repeaters. Where WDM is needed, they use coarse-WDM rather than more expensive dense-WDM.
15. Metro networks generally require less expensive terminal equipment than do long-haul networks.
16. Transmission distances in metro and regional networks typically are a few kilometers to about 200 km; most links are shorter than 100 km. Amplifiers are rare, but may be needed in sparsely populated rural areas or if coupling losses are high.
17. Dispersion and nonlinear effects are low because metro and regional systems are short.

What's Next?

In Chapter 25 you will learn about local telephone networks. Chapters 26 and 27 cover video and data networks.

Further Reading

John C. Ballamy, *Digital Telephony,* 3rd ed. (Wiley InterScience, 2000)

Yi Chen et al., "Metro Optical Networking," *Bell Labs Technical Journal* (January–March 1999) pp. 163–186.

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

Questions to Think About

1. Telephone calls that go outside of your community pass between your local-telephone switching office and other switching offices on the regional telephone network. What makes some of these long-distance calls?
2. A fire destroys the local switching office in one of the "larger towns" of Figure 24.1. How can the smaller towns connected to it still receive phone service?
3. Your company needs more transmission capacity between two offices 20 km apart. You can lease wavelengths at $20,000 per wavelength per year on a metro network with a five-year lease. You need to transmit three channels, two containing Gigabit Ethernet, and one containing TDM phone signals at 622 Mbit/s. They could fit on a single 2.5-Gbit/s line, but you would need to buy a time-division multiplexer and demultiplexer for the job. What's the most you could pay if you have to amortize costs over 5 years?
4. You're dealing with the same metro network as in Question 3. Their base rate is $1000 per kilometer per wavelength per year. What's the most you would pay for a time-division multiplexer and demultiplexer if you were sending signals only 5 km?
5. A metro network uses low-water single-mode fiber with zero dispersion at 1310 nm, and dispersion of 17 ps/nm-km at 1550 nm. If it transmits signals 100 km, what is the maximum data rate it can handle in the 1550-nm window with a transmitter having linewidth of 0.1 nm? Assume polarization-mode dispersion is insignificant, and that the maximum data rate is

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

6. You're using nonzero dispersion-shifted fiber in a metro network. If the chromatic dispersion is 1 ps/nm-km and the source bandwidth is 0.1 nm, how far can you transmit signals at 40 Gbit/s? Assume you can neglect polarization-mode dispersion.

Chapter Quiz

1. The best analogy for a regional telephone network is
   a. a pipeline pumping information from coast to coast with no detours.
   b. a superhighway crossing a large unpopulated area with few off-ramps.
   c. surface main streets connecting rural towns.
   d. residential streets in suburbia.
   e. a radio station broadcasting signals in all directions.
2. Which of the following is *not* connected directly to regional telephone networks?
   a. home and office telephones
   b. telephone switching offices
   c. cellular telephone networks

d. long-distance phone lines

e. Internet Service Providers

3. A regional network is a(n)

a. high-speed subway system in Washington, DC.

b. mesh connecting many telephone switching offices.

c. mesh connecting major cities across the country.

d. loop interconnecting many points in an urban area.

e. obsolete concept.

4. The hubs in a regional telecommunications network are

a. connections to international phone lines.

b. points from which cable television signals are distributed.

c. switching offices in towns and cities served by the network.

d. customers connected by add-drop multiplexers.

e. individual phone lines.

5. A regional telecommunications network serves

a. small towns in rural areas.

b. small cities and county seats in rural areas.

c. suburban communities.

d. large cities.

e. all of the above

6. Customers of wholesale metro networks may include

a. cable television operators.

b. regional telephone companies.

c. large businesses.

d. state and local governments.

e. Internet Service Providers.

f. all the above, if they have enough money.

7. A typical distance spanned by a metro network is

a. under 5 km.

b. a few to 200 km.

c. 100 to 200 km.

d. at least 200 km.

e. 200 to 1000 km.

8. Which of the following is true for the structure of signals on a regional network?

a. All signals must be at the same speed.

b. Signals can be transmitted at a hierarchy of speeds from T1 up.

c. Signals cannot be transmitted in different formats on the same fiber.

d. Signals always are transmitted at 2.5 Gbit/s.

e. All signals are time-division multiplexed together for transmission in a single data stream.

9. A metro network operates over a cable containing 864 individual fibers. If the system uses DWDM, each fiber can transmit 100 optical channels at 2.5 Gbit/s. What is the maximum capacity of the entire cable for two-way DWDM transmission?

a. 250 Gbit/s

b. 2.16 Tbit/s

c. 86.4 Tbit/s

d. 108 Tbit/s

e. 216 Tbit/s

10. The metro network described in Question 9 provides broadband service to a community of one million households. If the capacity is divided equally, what capacity can the cable offer to each home?

a. 250 kbit/s

b. 2 Mbit/s

c. 86.4 Mbit/s

d. 108 Mbit/s

e. 216 Mbit/s

11. What type of transmitter is most commonly used in metro and regional networks?

a. 850-nm VCSEL

b. 1310-nm single-channel laser

c. 1550 nm single-channel laser

d. coarse-WDM

e. dense-WDM

12. Which factor is least important in selecting coarse-WDM for use in a metro network?

a. Terminal equipment costs less than that for DWDM.

b. Shorter metro links do not require amplifiers.

c. Metro networks do not require many channels per fiber.

d. Shorter metro links have low dispersion.

e. All are equally important.

CHAPTER

25

Local Telephone or "Access" Networks

About This Chapter

The most visible part of the telephone system is the local network, connecting individual phones to the local switching office. Technological and industrial changes are reshaping the local phone network, as well as the rest of the telephone system. This chapter introduces you to the local phone system, sometimes called the access network, and shows how it serves home, business, and other users. It will explain the changes gradually spreading through the system, and the increasing role of fiber optics in delivering services to business and home customers.

Structure of the Local Phone Network

The local telephone network distributes signals to and from individual users. Traditionally, it was called the *subscriber loop*, industry jargon for the wires that form a circuit or loop from the local switching office to your home phone. At the switching office, local lines connected to a switch that could route signals to other local lines, the regional phone network, or the long-distance network.

That picture has become far more complicated as the network has grown and added new capabilities. The modern version is sometimes called the *access network* because it gives subscribers access to the global telecommunications network, and no longer carries only standard wire-line telephone calls. Facsimile and data traffic also goes through standard telephone lines. Cellular or mobile phone networks also carry voice traffic, although the signals don't go through wires. Voice also may be digitized and converted into Internet format at the phone itself, a new type of telephone traffic. The local telephone network also may carry additional digital signals from computer networks and broadband terminals.

Some of the new signals are additions to the public switched telephone network and are routed through standard switching offices or their equivalents. Fax traffic goes through standard telephone lines; cell phones are directed through cellular switches. Much data traffic is routed through cables originally laid for telephone traffic because the lines are convenient. In this modern topology, the access network runs between individual subscribers and the *network edge*, the point where signals enter the regional or global telecommunications network.

The major difference between the traditional subscriber loop and the new access network is that the access network includes services and options that have been added to the traditional local phone network. Many details of access networks and traditional subscriber loops differ. However, the access network has its roots in the traditional telephone network, and the two retain a common functional structure. To understand that structure, let's look at it step by step, moving toward the individual subscriber.

The subscriber loop distributes telephone signals to individual users.

(Over the past two decades, the local telephone network has converged in many ways with the cable-television networks described in Chapter 27 and the data networks described in Chapter 26. You'll learn about this convergence later.)

The Network Edge

The "edge" of the network is the point where signals enter the regional or global telecommunications network. As Figure 25.1 shows, the edge includes traditional local telephone

FIGURE 25.1
The network edge.

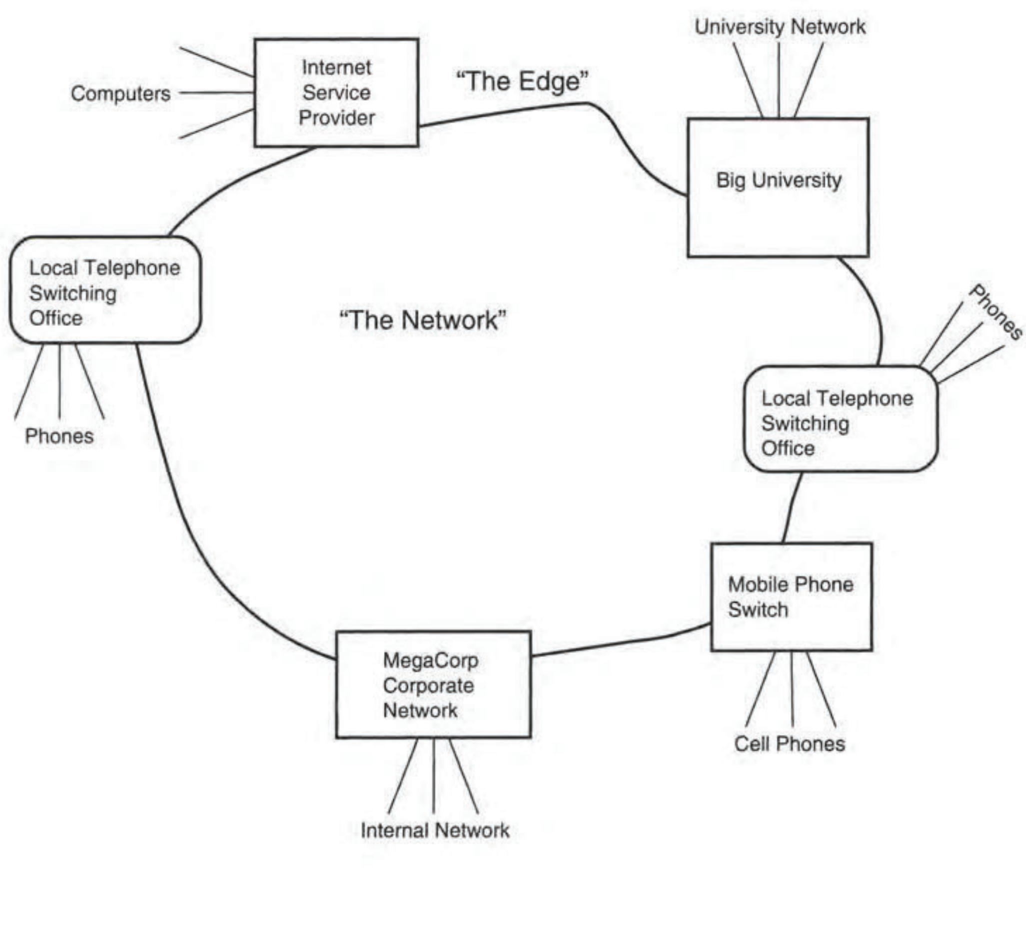

switching offices as well as mobile telephone switches, Internet Service Providers, and large organizations that make heavy use of telecommunication services. The figure greatly oversimplifies, by hiding the core of the network and covering the edge in only one region. The point is that equipment on the edge serves as a link between individual users and the telecommunications network as a whole, whether the user has a new mobile phone, an old wire-line phone, or some other connection.

The network edge is the interface where signals are directed into the global network.

The edge is an interface from which signals are directed. Switches in a central office direct individual telephone calls elsewhere in the telecommunications network. Switches operated by mobile phone carriers and competitive phone companies serve the same function. Likewise, an Internet Service Provider directs data packets, typically to routers somewhere inside the network. Corporate and university networks have come to serve similar functions internally rather than directing signals through a telephone switching office or Internet Service Provider.

The Telephone Switching Office

A *switching office* or *central office* is the facility where a wire-line telephone company makes connections to and from individual telephone customers. It is a central point in the community from which telephone services radiate outward over cables to individual users. Figure 25.2 shows the concept, which dates back to the early days of the telephone industry. The dark lines show where fiber has replaced copper.

A switching office makes connections to individual customers.

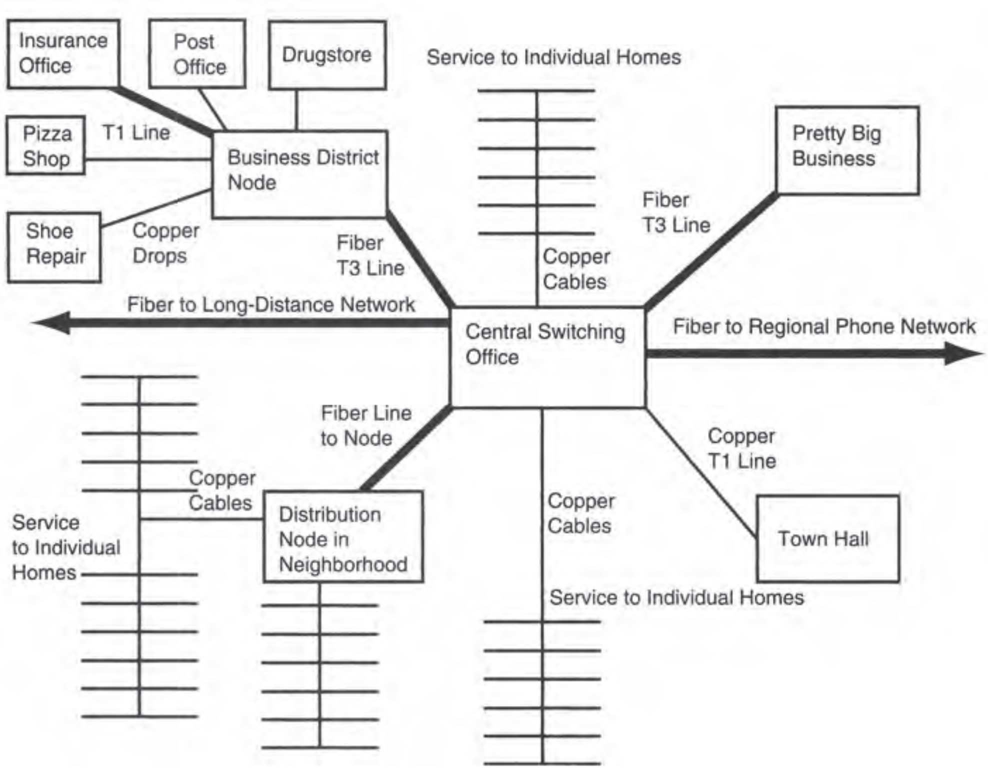

FIGURE 25.2
A telephone switching office has a central role in distributing signals.