


FIGURE 14.6
Important coupler
types.



Wavelength-selective couplers distribute signals according to their wavelength.

Wavelength-selective couplers distribute signals according to their wavelengths. Their main uses are to route WDM signals to their proper destinations and to separate wavelengths transmitted for different purposes through the same fiber, such as separating the light pumping an optical amplifier from the amplified signal. Wavelength-selective couplers are supposed to block other wavelengths from reaching the wrong destination. Chapter 15 covers them in detail.

Bulk and Micro Optics

In the world of fiber optics, *bulk optics* are conventional lenses, mirrors, and diffraction gratings, the sort of things you can hold in your hands. Bulk optics do not have to be large; they may be made quite small to match the dimensions of optical fibers and light sources. Such *micro optics* may be tiny, but they are still based on the same optical principles as larger bulk optics, so I will cover them together.

Beamsplitter transmits half the input light, reflects the other half.




FIGURE 14.7

Bulk optical coupler: A beamsplitter divides a signal in half.

Bulk optics were the basis of many early types of couplers, and they still work. Figure 14.7 shows a simple example, the use of a device called a *beamsplitter* to split one input signal into two outputs. Like a one-way mirror, the beamsplitter transmits some light that hits it and reflects the rest. Collect the light from the two outputs in fibers, and you have a T coupler.

Bulk optical couplers often include lenses that expand, collimate, or focus light. The simple coupler of Figure 14.7 generally works better if a lens expands the light emerging from a fiber and focuses it onto a large area of the beamsplitter. This function is *collimation*, and such optics are called *collimators*. Then additional lenses focus the output beams into output fibers. Standard lenses with curved surfaces may be used; generally they are tiny, to match the sizes of fibers.

Alternatively, gradient-index (GRIN) lenses (or rods) may be used. These are rods or fibers in which the refractive index of the glass changes either with distance along the rod or with distance from the axis. The refractive-index gradient makes GRIN lenses focus light in a way functionally equivalent to ordinary lenses, but GRIN lenses are smaller and easier to adapt to fiber systems.

Another application of bulk optics is the use of a diffraction grating to separate wavelengths. A diffraction grating is an array of closely spaced parallel grooves, which act together to scatter light at an angle that depends on its wavelength, generating a rainbow of colors.

Fused-Fiber Couplers

Normally you can't transfer light between fibers just by touching them together. The light-guiding cores are covered by claddings that keep light from leaking out. If you want to couple light between fibers, you have to transfer it between the cores. That means you have to remove the claddings from one side of each fiber so the cores can touch. That is the basis of fused-fiber couplers, made by melting together fibers, usually with claddings removed partly or totally from one side, as shown in Figure 14.8. Often the fibers are twisted together to improve light transfer. Fused-fiber couplers sometimes are called *biconic* couplers, which should not be confused with biconic connectors, an early type in very limited use today. They are the most common technology used to make couplers.

Although Figure 14.8 shows the cores merged, they don't have to merge completely in the middle zone as long as they come close. A phenomenon called *evanescent-wave coupling* can transfer light energy through a thin cladding or material with lower refractive index than the light-guiding zone. As learned in Chapter 4, a small amount of the light energy guided


Micro optics are tiny versions of conventional lenses and other optical components, shrunk in size to work with fibers.

GRIN lenses are rods or fibers with refractive index graded so they act like ordinary lenses.

Fused-fiber couplers are the most widely used type.

FIGURE 14.8

A 2×2 fused-fiber coupler.



in the core of an optical fiber actually penetrates a short distance into the cladding. This light is called the *evanescent wave*, and you can see its effect in single-mode fibers, where it makes the mode-field diameter larger than the core of a step-index single-mode fiber. Evanescent-wave coupling is important in both fused-fiber and waveguide couplers.

Fusing two fibers produces a 2×2 coupler with two inputs and two outputs. In practice, these are often turned into 1×2 couplers by cutting one fiber end inside the case. This design is functionally directional, although it is bidirectional in the sense that light can go through it in either direction. If light enters the fiber end at upper left in Figure 14.8, the only way light can reach the fiber end at lower left is by reflection or scattering. Directivity is measured by comparing the input power, P_1 , to the power reflected back through the other fiber end on the input side, P_4 :

$$\text{Directivity (dB)} = -10 \log \left(\frac{P_4}{P_1} \right)$$

For a typical fused-fiber coupler, the directivity is 40 to 45 dB.

The details of fused-fiber coupler operation depend on whether the fibers are multimode or single-mode. In multimode couplers, the higher-order modes leak into the cladding and into the core of the other fiber; the degree of coupling depends on the length of the coupling zone, and does *not* depend on wavelength. In single-mode fibers, light transfers between the two cores in a resonant interaction that varies with length. If all the light enters in one fiber, it gradually transfers completely to the other, then transfers back as it travels farther, shifting back and forth cyclically. The distance over which the cycling takes place depends on the coupler design and the wavelength, as Chapter 15 will describe in more detail.

The fused-fiber coupler design can be extended to multiple fibers using the same basic principles. The important change is adding more fibers, so signals from all input fibers mix in the coupling zone and emerge out of all the output fibers. This approach can be used to make star couplers with many distinct inputs and outputs. Such multifiber fused couplers are bidirectional.

Planar Waveguide Couplers

As you learned earlier, optical fibers are not the only type of optical waveguides. Like fibers, planar waveguides confine light in a region of high refractive index surrounded by material with a lower refractive index. The planar waveguide may be a thin strip embedded in the surface of a flat substrate, as you saw in Figure 6.12, or it could be a strip deposited on the top of a flat substrate. Air and the substrate combine to serve the function of the cladding in a fiber.

Waveguide patterns are written using the same techniques that write integrated electronic circuits onto semiconductor wafers. They can branch or merge, making them the equivalent of fused-fiber couplers. A simple example is a Y-shaped structure that divides one input waveguide to form two outputs, as shown in Figure 14.9. (An actual split would be much smaller than shown.) If the outputs split at equal angles, as shown, the light divides equally between them. This approach can be extended to more outputs by adding Y couplers to divide the output signals. Alternatively, the input waveguide could be divided to give more outputs, but the power would not be evenly divided.

Two closely spaced waveguides on the same substrate can also transfer light by evanescent-wave coupling through a thin layer of lower-index material, as in fused-fiber

Simple waveguide couplers are branched planar waveguides.





FIGURE 14.9

Planar waveguide splits in two, so light divides equally between arms of the Y coupler.

FIGURE 14.10

Light transfer between two evanescently coupled waveguides.



couplers. This type of waveguide device is called an *evanescent-wave coupler* and shown in Figure 14.10.

An evanescent-wave coupler gradually transfers light between the two waveguides, along the region where the waveguides are closely spaced. In Figure 14.10, light enters in the top waveguide, and gradually transfers to the lower waveguides. This continues until all light shifts from the upper to the lower waveguide at a point called the *transfer length*, which depends on the optical characteristics of the waveguide. At this point, all the light is in the lower waveguide, and the process reverses, with the light starting to shift back from the lower waveguide to the upper one. Thus the distribution of light energy between the two waveguides oscillates back and forth between them with distance, as shown in the lower part of Figure 14.10. This oscillation stops at the end of the coupling region, determining the final distribution of light. The same process occurs in single-mode fused-fiber couplers.

Designers select lengths and optical properties of the two parallel waveguides to give the desired distribution of light (e.g., 50/50 or 75/25). In practice some light is lost within the waveguide and in transferring between the two waveguides.

Surface waveguides can be fabricated in complex patterns on a variety of materials. When they are made on the same substrate with many other devices, they are often called *integrated optics*, but then they usually contain active devices such as lasers, switches, or modulators. Chapter 16 covers such devices.

Active Couplers

Devices called *active couplers* also look to the user as if they split signals from fiber-optic transmission lines, but if you look closely, they work quite differently. An active coupler is essentially a dedicated repeater in which the signal from a receiver drives two (or more)

Evanescent-wave couplers depend on light leakage between two closely spaced waveguides.

An active coupler is a repeater with two or more outputs.

transmitters, which can generate optical and/or electronic output signals. This means that active couplers are not passive devices. However, they do function as couplers, so they are mentioned here.

Active couplers are mostly used in local-area networks. For example, a fiber that runs to a network node may drive a receiver that generates two electronic outputs. One goes to an optical transmitter, which generates an optical signal to send through the next length of fiber in the network. The other is transmitted in electronic form to the terminal attached to that network node. Other configurations also are possible.

Attenuators

As you learned in Chapter 11, too much light can overload a receiver. Attenuators reduce light intensity, by transmitting only a fraction of the input light. They are needed when a transmitter could deliver too much light, such as when it is too close to the receiver.

An *attenuator* is a type of optical filter, which should affect light of all wavelengths transmitted by the system equally. Attenuators are like sunglasses, which protect your eyes from being dazzled by bright lights. Fiber-optic attenuators generally absorb the extra light energy, which is too little to heat the attenuator noticeably. They should not reflect the unwanted light, because it could return through the input fiber to cause noise in a laser transmitter.

Most attenuators have fixed values that are specified in decibels. For example, a 5-dB attenuator should reduce intensity of the output by 5 dB. Attenuators designed for general optics use may have attenuation specified as the percent of light transmitted (T) or as *optical density*. Optical density is defined as:

$$\text{Optical Density} = \log_{10}\left(\frac{1}{T}\right)$$

This should look familiar, because it's close to the formula for attenuation in decibels, without the factor of 10. You can think of optical density as 0.1 times attenuation in dB, so a filter with optical density of 2 has a 20-dB loss.

Variable attenuators also are available, but they usually are used in precision measurement instruments.

If you're familiar with electronics, it may be tempting to think of an attenuator as an optical counterpart of a resistor. This is not a good general analogy. An attenuator does limit the flow of light like a resistor limits current flow—but resistors also serve other circuit functions, such as providing voltage drops, and controlling circuit loads. The only job of attenuators in a fiber-optic system is to get rid of excess light.

It's important to distinguish between attenuators and other types of optical filters. Attenuators should have the same effect on all wavelengths used in the fiber system. That is, if the attenuator reduces intensity at one wavelength by 3 dB, it should do the same at all wavelengths. Other types of filters typically do not affect all wavelengths in the same way. For example, a filter might transmit light in the 1530 to 1565 erbium-amplifier band, but have 50 dB attenuation in the 980-nm pump band. In fiber-optic systems, the term *filter* is used for filters in which light transmission varies significantly with wavelength; they are used in wavelength-division multiplexing, and covered in Chapter 15.

Attenuators reduce light intensity uniformly across the spectrum.

Optical Isolators

Optical isolators transmit light only in one direction.

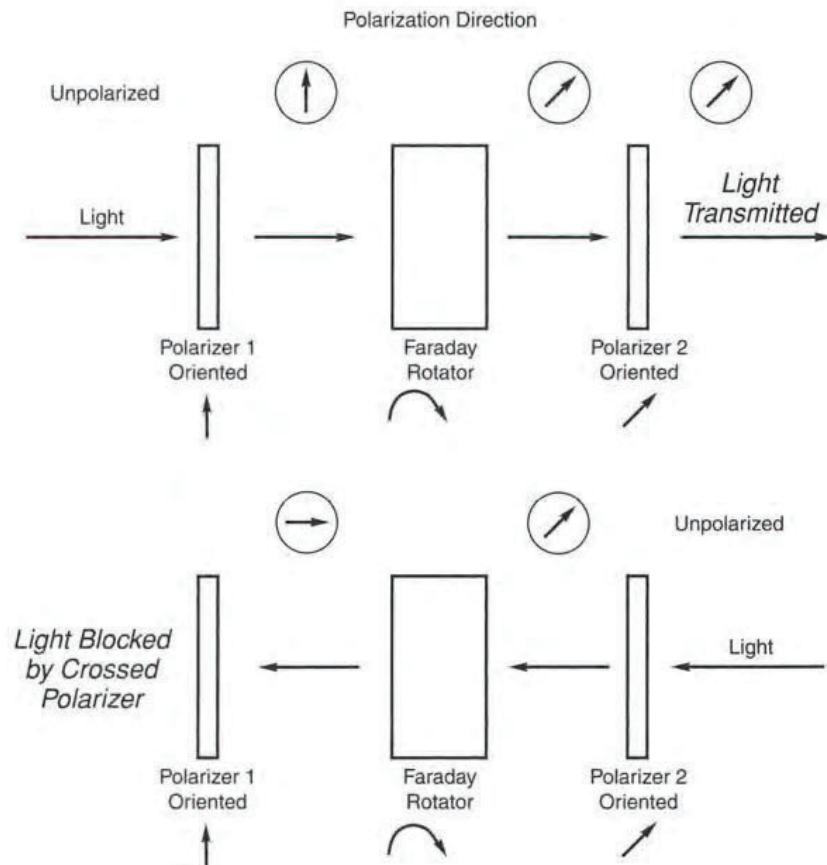
Optical isolators are devices that transmit light only in one direction. They play an important role in fiber-optic systems by stopping back-reflection and scattered light from reaching sensitive components, particularly lasers. You can think of them as optical one-way streets with their own traffic cops or as the optical equivalent of an electronic rectifier (which conducts current only in one direction).

The operation of optical isolators usually depends on materials called *Faraday rotators*, which rotate the plane of polarization of light. The rotation is always by the same angle when seen from the viewpoint of the light source. But for light transmitted from sources on opposite sides of the Faraday rotator, the angles are in different directions. With a bit of smart design, this feature can separate light going in different directions, so light going in the desired direction gets through, but light going the wrong way is stopped.

Figure 14.11 shows a simple example. First consider light going from left to right, the desired direction. The input light is unpolarized, but it passes through a linear polarizer, which transmits only light polarized vertically. Then the Faraday rotator twists the plane of polarization 45° to the right. The light then encounters a second polarizer, which is oriented so it transmits only light with its plane of polarization oriented 45° to the right of

FIGURE 14.11

An optical isolator transmits light in only one direction.



vertical. That's all of the light going from left to right, so the signal goes through unimpeded except for a 3-dB loss because the input polarizer blocked half of the unpolarized input signal.

Now consider light going in the opposite direction, from right to left. The polarizer at right transmits only light polarized at 45° to the vertical, and the Faraday rotator turns the plane of polarization another 45° to the right. That makes the plane of polarization horizontal, so the light is blocked by the vertical polarizer at the left. A little stray light does leak through, but light headed in the wrong direction can be attenuated by 40 dB or more, protecting lasers from stray light that could induce noise.

One drawback of this simple design is that it's polarization sensitive. The input polarizing filter blocks half the input light that is not vertically polarized, causing 3-dB loss. More refined polarization-insensitive designs instead separate the input signal into two beams: one vertically polarized and the other horizontally polarized. One approach uses transparent crystals in which light travels at different speeds depending on its polarization. Prisms of such strongly birefringent materials separate vertically and horizontally polarized light so they follow different paths; these prisms are sometimes called *beam displacers*. They can be combined with focusing elements and Faraday rotators so that light traveling in one direction is focused from the input fiber into the output fiber, while light traveling in the opposite direction is defocused to prevent it from going into the input fiber. Although this design is a bit more complicated, it avoids 3 dB of loss.

● Polarization-insensitive optical isolators do not have 3-dB internal loss.

Optical Circulators

The optical circulator is a cousin of the optical isolator in both its function and design. Its function is to serve as a one-way street for light passing through a series of ports, so light that enters in port 1 must go to port 2, and any light entering at port 2 goes to port 3, and so on. Like the optical isolator, it uses polarization to do its job.

One way to make an optical circulator is with a pair of optical isolators. One can be inserted between port 1 and port 2, blocking light going backwards from port 2. A second can be inserted between ports 2 and 3, blocking light trying to go back from port 3 to port 2. However, these designs lose the blocked light.

Figure 14.12 shows a more elegant and efficient optical circulator, which is assembled from three types of components. Faraday rotators and birefringent beam displacers also are used in the optical isolators described above. Recall that the displacers separate light of different polarization, while the Faraday rotators always rotate the polarization by 45° from the viewpoint of a photon passing through them. If light goes back and forth through the same Faraday rotator, its polarization changes a total of 90° .

Optical circulators also include devices called *waveplates*, which also rotate the polarization by 45° , but in a different way. If light passes through a waveplate one way, it's rotated 45° to the right; if it goes through the other way, the light rotates 45° to the left. That means the net change on a round trip is 0° , so light that makes a round trip through a waveplate emerges with the same polarization it started with.

Now go back to Figure 14.12 and trace the paths of the two polarizations from port 1 to port 2. The vertically polarized input is deflected up, then rotated $+45^\circ$ by the waveplate

● An optical circulator sends light in one direction through a series of ports.

Faraday rotators shift polarization $+45^\circ$ when light goes in either direction.

Waveplates rotate polarization -45° when light goes one way, $+45^\circ$ when light goes the other way.




FIGURE 14.12

Optical circulator directs light from part 1-2, 2-3, and 3-4, without allowing it to go backwards.

and another $+45^\circ$ by the Faraday rotator, a total of 90° , making it horizontally polarized so it passes straight through the second beam displacer. Then it is rotated -45° by the waveplate and $+45^\circ$ by the Faraday rotator, a net change of zero, so it remains horizontally polarized through the third displacer and out port 2. The horizontally polarized input, in contrast, goes straight through the first displacer, and is rotated -45° by the waveplate and $+45^\circ$ by the Faraday rotator, a net of zero. It then goes straight through the second displacer. At the second rotator, it is rotated $+45^\circ$ by the waveplate and $+45^\circ$ by the Faraday rotator, a total of 90° , which makes it vertically polarized so the displacer bends it upwards—and aims it out port 2, where it is supposed to go.

The tricky part is following the path from port 2 to port 3 (right to left in Figure 14.12). The beam displacer splits the two polarizations so they pass through the second rotator. This time the top (horizontally polarized) beam is rotated $+45^\circ$ by the waveplate and $+45^\circ$ by the Faraday rotator, changing its polarization to vertical. The middle beam displacer bends the beam upward, and it emerges from the upper side of the middle beam displacer on its way to port 3. In this case, it is rotated -45° by the waveplate and $+45^\circ$ by the Faraday rotator, so it remains vertically polarized, and is deflected downward to port 3. The bottom (vertically polarized) light from port 2 is deflected downward, where the waveplate rotates it -45° and the Faraday rotator rotates it $+45^\circ$, leaving it vertically polarized as it enters the middle beam displacer. It's bent upward, and arrives at the lower position on its way to port 3. Here the polarizer rotates it $+45^\circ$ and the Faraday rotator rotates it $+45^\circ$, changing it to horizontally polarized light that goes straight through the beam displacer to port 3.

Each level of the optical circulator is identical, so the steps can be repeated as long as you want. The crucial tricks are separating the polarizations, routing them through different components, and taking advantage of the different ways Faraday rotators and waveplates rotate polarization.

What Have You Learned?

1. Couplers connect three or more fibers or ports. Dividing an optical signal among two or more ports reduces its strength because it divides the photons in the signal.
2. Several different types of couplers are used; their design depends on the application.
3. Direction is important in couplers. Most couplers are directional in the sense they transmit signals from one or more inputs to one or more outputs, with little light going from one input to another. Most designs also are bidirectional, in the sense that the input and output ports could be reversed to change light coupling.
4. T or Y couplers, or taps, are three-port devices. Tree, or 1-to- n , couplers divide one input among n output ports. Star couplers have multiple inputs and outputs. Outputs are usually distinct from inputs.
5. Wavelength sensitivity is important in couplers. It is desirable for wavelength-division multiplexing, but not for most other applications.

6. Many couplers are made from bulk or micro optics, such as beamsplitters.
7. GRIN lenses are rods or fibers with refractive index graded so they refract light like ordinary lenses.
8. Fused-fiber couplers transfer light between the cores of two fibers melted together. Single-mode fused-fiber couplers work differently than multimode versions. Multifiber fused-fiber couplers are possible.
9. There are two types of planar waveguide couplers. Some simply divide light between two waveguides branching in a Y from a single-input guide. Others rely on evanescent-wave coupling to transfer light between two parallel waveguides. Evanescent-wave couplers are sensitive to wavelength.
10. Active couplers are repeaters with two or more outputs.
11. Attenuators block light uniformly across a range of wavelengths to reduce signal strength at the receiver.
12. Optical isolators transmit light in only one direction. They rely on polarizing optics and Faraday rotators.
13. Optical circulators route light through a series of ports, feeding output from one port to the next, and taking input from the second port and routing it to a third. They rely on birefringent crystals, Faraday rotators, and other polarization rotators.

What's Next?

Chapter 15 covers wavelength-selective optics used for wavelength-division multiplexing.

Further Reading

Morris Hoover, "New coupler applications in today's telephony networks," *Lightwave*, Vol. 17, March 2000.

Luc B. Jeunhomme, *Single-Mode Fiber Optics: Principles and Applications* (Marcel Dekker, 1990). See Chapter 6, "Passive Components."

Questions to Think About

1. Suppose your input signal is -10 dBm and your receivers require a signal of at least 0 dB μ . You want to distribute signals to as many terminals as possible. If there is 3 dB of fiber loss between you and each receiver, how many terminals can you deliver signals to? How much coupler loss does this correspond to on each channel? Assume you can buy a star coupler with as many ports as you want, which has no excess loss.

2. Suppose that all star couplers available for the system described in Question 1 have excess loss of 3 dB. How many terminals can you reach with these couplers, and what is the total loss per channel?
3. An alternative design is to cascade a series of 3-dB T couplers. The first splits the signal in half, then each output has its own 3-dB coupler, dividing that output in half, yielding one-quarter of the original output. Adding more layers further divides the signal. Suppose you can get as many 3-dB couplers as you want and each one has no excess loss. How many terminals can you divide signals among in Question 1?
4. A local-area network includes 90/10 couplers, which split 10% of the input signal and deliver it to a local terminal. Suppose you have 10 of them in series and the input power is -10 dBm. What is the power delivered by the last coupler out each of its ports?
5. If your receiver requires $1 \mu\text{W}$ of power, how many more 90/10 couplers could you have in series before the 10% side does not deliver enough power for reliable operation? Assume the same -10 dBm input as in Question 4.
6. An optical amplifier delivers 0.5 mW/channel on each of 32 channels. You want to monitor its performance by diverting a small portion of its output to an optical performance monitor that requires $1 \text{ dB}\mu$ input per optical channel. What fraction of the output power do you need to divert to the performance monitor?
7. Neglecting excess internal losses, what is the difference in attenuation between the following two optical circulators? The first uses the simple optical isolators of Figure 14.11—one oriented from port 1 to 2 and the second from port 2 to 3, with a 50/50 T coupler splitting input signals from port 2 between two routes. The other is the more complex optical circulator of Figure 14.12.

Chapter Quiz

1. You have a coupler that divides an input signal equally among 16 outputs. It has no excess loss. If the input signal is -10 dBm, what is the output at any one port?
 - a. -12 dBm
 - b. -20 dBm
 - c. -22 dBm
 - d. -26 dBm
 - e. -30 dBm
2. A 1×20 coupler has output signals of -30 dBm at every port if the input signal is -10 dBm. What is its excess loss?
 - a. 0 dB
 - b. 1 dB

- c. 2 dB
 - d. 4.2 dB
 - e. 7 dB
- 3.** A coupler splits an input signal between two ports with a 90/10 ratio. If the input signal is -20 dBm and the coupler has no excess loss, what is the output at the port receiving the smaller signal?
- a. -21 dBm
 - b. -29 dBm
 - c. -30 dBm
 - d. -31 dBm
 - e. -110 dBm
- 4.** What type of coupler could distribute identical signals to 20 different terminals?
- a. T coupler
 - b. tree coupler
 - c. star coupler
 - d. $M \times N$ coupler
 - e. wavelength-selective coupler
- 5.** What type of coupler divides one input signal between two output channels?
- a. T coupler
 - b. tree coupler
 - c. star coupler
 - d. $M \times N$ coupler
 - e. wavelength-selective coupler
- 6.** You find a coupler with four ports and no label on it. You measure attenuation from port 1 to the other three ports. The values are: -40 dB to port 2, 3 dB to port 3, 3 dB to port 4. What type of coupler do you have?
- a. star coupler with three unequal outputs
 - b. tree coupler with three unequal outputs
 - c. a directional 2-by-2 coupler with two inputs and two outputs
 - d. a nondirectional T coupler
 - e. a broken coupler
- 7.** A Y coupler that equally divides light between two outputs has a 3-dB loss on each channel. What is the right explanation?
- a. The 3-dB figure is excess loss.
 - b. Half the photons that enter the coupler go out each output, corresponding to a 3-dB loss on each channel.
 - c. The coupler polarizes the light going out each port, causing 3-dB loss.

- d. Every coupler has at least 3-dB loss no matter how it divides the input light.
 - e. The optics are dirty, causing loss of half the input light.
- 8.** Evanescent waves cause light energy to transfer between channels in what type of coupler?
- a. planar-waveguide coupler
 - b. single-mode fused-fiber coupler
 - c. bulk optical coupler
 - d. multimode fused-fiber coupler
 - e. a and b
 - f. b and d
- 9.** An attenuator
- a. is a filter that blocks one wavelength but transmits others.
 - b. polarizes input light, causing loss of the other polarization.
 - c. reduces light intensity evenly across a range of wavelengths.
 - d. selectively blocks photons produced by spontaneous emission.
- 10.** How many polarizers does a simple polarization-sensitive optical isolator use to block transmission of light in the wrong direction?
- a. none
 - b. 1
 - c. 2
 - d. 3
 - e. 4

Wavelength-Division Multiplexing Optics

About This Chapter

Wavelength-division multiplexing (WDM) multiplies transmission capacity by allowing a single optical fiber to carry separate signals at multiple wavelengths, but that benefit comes at a cost in complexity. Additional optical components are needed to combine and separate optical channels at closely spaced wavelengths, and to process the signals as they are transmitted. Specific requirements vary with system design, and several technologies are in use.

This chapter first outlines the basic requirements, then explains how these requirements relate to the operation of WDM systems. Then it covers specific technologies, their operation, and their capabilities. The chapter introduces the important concepts of wavelength-selective optical filtering, and describes the types of devices that can perform it. You will recognize a little overlap with the couplers and attenuators described in Chapter 14, which can be adapted for wavelength selectivity. The current chapter concentrates on passive technologies like those in Chapter 14, but does mention some active technologies that serve similar purposes in “dynamic” versions of passive components. Chapter 16 will cover switching, modulation, and other “active” technologies.

WDM Requirements

At first glance, wavelength-division multiplexing looks simple; it's easy to shine light of many different colors into an optical fiber. The hard part seems to be the *demultiplexing*, when you have to separate the output signals by their wavelength. Unfortunately, reality is not quite that simple. Demultiplexing is a more difficult task, but care does

Demultiplexing is harder than multiplexing.

need to be taken in multiplexing, to control background noise that might otherwise be introduced into the fiber. Other types of processing also require some type of wavelength selection.

There are a number of basic requirements for WDM, which often differ in degree:

- *Wavelength multiplexing* or combination, which must transmit the desired signal while blocking noise at other wavelengths, which might interfere with other optical channels.
- *Wavelength demultiplexing*, to isolate individual optical channels for detection or further processing.
- *Add-drop multiplexing*, often called “optical add-drop multiplexing,” which separates and combines only a few wavelengths from many wavelengths transmitted by a system.
- *Wavelength separation*, typically to isolate a signal wavelength from a pump wavelength, as in an optical amplifier.
- *Wavelength-selective processing*, such as attenuating signals at some wavelengths to balance the strength of all optical channels in a system. This may be static or dynamic.
- *Wavelength conversion*, to change a signal from one wavelength to another. (See Chapter 12.)
- *Wavelength switching*, to redirect signals at one or more wavelengths while transmitting others unchanged. (See Chapter 16.)

The specific requirements depend on system properties such as the number of channels, the spacing between them, the distance spanned, the amplification required, and the system configuration. Requirements differ greatly between a system that carries eight optical channels between two office buildings 10 kilometers apart, and a network that carries 32 optical channels several hundred kilometers, dropping and adding signals at cities along the way. A few of these requirements are not exclusive to WDM systems, such as separating the pump light from the signal in an erbium fiber amplifier, which could be amplifying only a single wavelength.

During the telecommunications bubble, there was much talk of dynamically reconfigured optical networks carrying large numbers of optical channels, which they would routinely switch and convert to different wavelengths. Most of these ideas were never implemented outside the laboratory, and little demand for them seems likely in the near future. We will talk about some of these ideas that may find practical applications, but not in much detail.

● Many advanced concepts were never implemented outside the lab.

WDM Systems

The basic elements of a WDM system are shown in Figure 15.1. Optical signals from eight transmitters at separate wavelengths are combined in a *multiplexer* at the left. The signals travel through the same fiber from the multiplexer to an *add-drop multiplexer* in the middle, which extracts one wavelength, λ_4 , from the main signal and directs it to a local receiver. It also picks up another signal from the local transmitter at the same wavelength, called λ_{4^*} , to show it is a different signal.

● Multiplexers combine optical channels; demultiplexers separate them.

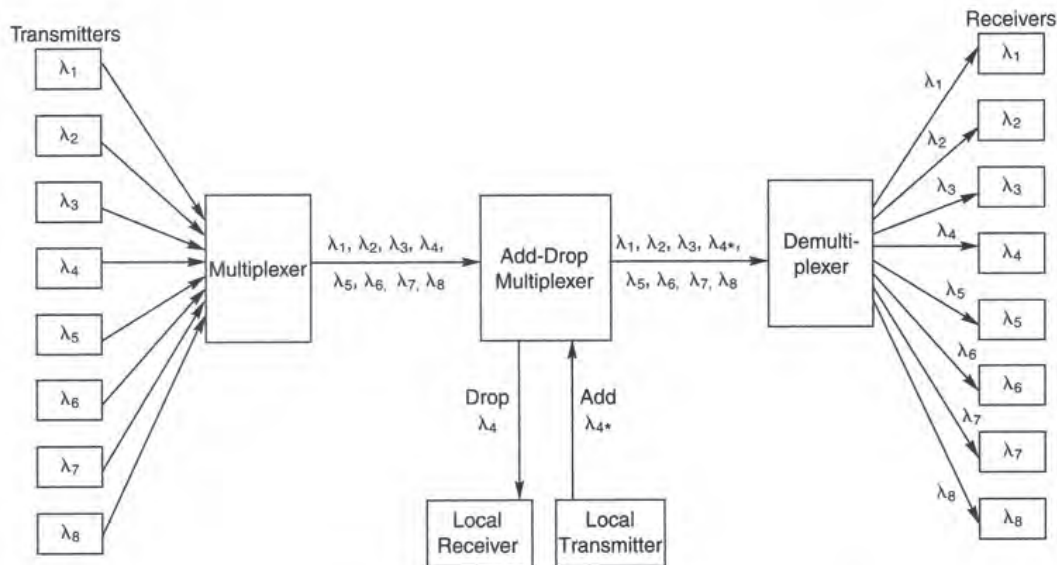


FIGURE 15.1
A WDM system.

At the right side, the eight signals are split in a *demultiplexer* and routed to separate receivers, one for each wavelength. As you learned in Chapter 11, receivers are color-blind in the sense that they respond in the same way to all wavelengths they can detect. This means that the signals must be completely separated, because any stray light at a different wavelength will show up as noise in the receiver output. If some light from the signal at λ_5 reached Receiver 6, the receiver would think it belonged in channel 6, and it would interfere with the actual λ_6 signal.

You can think of the multiplexer and demultiplexer as mirror images, but they are not identical. The multiplexer takes separate wavelengths and combines them, and the demultiplexer takes combined wavelengths and separates them. Key operating considerations differ between the two. Multiplexers should have low insertion loss and avoid scattering light back to any of the transmitters. Demultiplexers must reliably separate the optical channels, with low leakage of light from one optical channel into an adjacent channel. In practice, the two devices can be used as mirror images of each other, although sometimes the multiplexers may have wider channel spacing than the demultiplexers to reduce insertion loss.

The add-drop multiplexer serves a different function, picking out one or more wavelengths from a combined signal so they can be “dropped” at a location partway along a system. It also can “add” signals transmitted from the midpoint station onto empty channels. In Figure 15.1, it is adding a signal λ_{4^*} to replace the signal that has been dropped at the λ_4 wavelength.

The way these system elements operate depends on various factors including the number and density of optical channels, and the topology of the system—that is, how it collects and distributes signals.

Optical Channel Density

WDM optics provide a number of uniformly spaced slots for optical channels, although the system operator may not populate all those slots with transmitters and receivers. The

Demultiplexers must separate optical channels completely, with low crosstalk.

spacing of these slots determines the potential optical channel density. The maximum number of channels that can fit into a given spectral width is the bandwidth divided by the channel spacing:

$$\text{Channel Capacity} = \frac{\text{total bandwidth}}{\text{channel spacing}}$$

System designers can pick any channel spacing they want, but in practice most WDM systems fall into two specific categories based on industry standards:

● DWDM channels are spaced 200, 100, or 50 GHz apart.

● CWDM channels are spaced 20 nm apart.

Dense Wavelength-Division Multiplexing or *DWDM* is based on channel spacing of 200 GHz or less. Normal center-frequency spacings are 200, 100, or 50 GHz, based on a standard grid developed by the International Telecommunications Union. A 100-GHz spacing corresponds to 0.8 nm in the erbium-amplifier band. Most DWDM systems operate in the erbium-amplifier band around 1550 nm, which is used for long-haul, high-capacity transmission.


Coarse Wavelength-Division Multiplexing or *CWDM* is based on channel spacing of 20 nm across a range from 1270 to 1610 nm, set by the ITU under standard G.694.2. The CWDM standard creates 18 channels across the low-loss window of low-water fibers. CWDM is intended for use in metro networks where transmission distances are at most tens of kilometers and amplifiers are not required.

Note that CWDM spacing is in wavelength units while DWDM spacing is in frequency units. This means that DWDM channel spacing is uniform in frequency across its operating range, but not uniform in wavelength. Likewise, CWDM channel spacing is uniform in wavelength but not in frequency units.

The difference in fiber capacity of these two spacings is dramatic, as shown in Figure 15.2. Forty 100-GHz DWDM channels fit into the 1530–1565 nm C-band of erbium amplifiers.

FIGURE 15.2

Channel spacing in DWDM and CWDM systems.



Two CWDM channels don't quite fit into the same space; channels are centered at 1530, 1550, and 1570 nm.

Wide channel spacing reduces the costs of multiplexing and demultiplexing optics; it also allows the use of inexpensive uncooled directly-modulated laser transmitters. Narrow channel spacing raises costs, but allows a single fiber to carry more bandwidth, which in practice is more important over long distances or where the number of available fibers is limited.

Channel Separation Requirements

Demultiplexing optics should separate optical channels cleanly. Ideally the optics should transmit all light at the center wavelength of the optical channel, but block adjacent channels completely. In practice, DWDM optics have some attenuation at the center of the channel, typically 3 to 5 dB, and adjacent channels are attenuated by 20 to 40 dB, as shown in Figure 15.3. Normally the channels are equally spaced, so the points where the curves intersect match the channel widths or, equivalently, the channel separation. Actual transmission curves depend on the technology used for the WDM optics, and in practice spread out more at the bottom than the simplified curves of Figure 15.3.

The overlap of the transmission curves shows where crosstalk occurs. The crosstalk can be reduced by using optics with spectral width narrower than the optical channel, which reduces the amount of signal transmitted but reduces background noise. The amount of

DWDM optics don't isolate channels completely.




FIGURE 15.3

Channel isolation in WDM system.

crosstalk also depends on the range of wavelengths emitted by the transmitter on each channel, which is not shown.

In practice, DWDM systems generally have some empty channels. Some are initially left open to allow later capacity upgrades. Others are in gaps intentionally built into the system to avoid possible crosstalk, such as the gap normally left between C- and L-band erbium amplifiers. Standard C-band amplifiers stop at 1565 nm, and L-band amplifiers don't start until 1570 nm, leaving a 5-nm gap. Some DWDM systems leave similar gaps between blocks of wavelengths, perhaps dropping one or two 100-GHz slots between blocks of eight optical channels.

The wider spacing of CWDM channels eases fabrication of the optics, so crosstalk and channel separation are not as serious concerns as they are for DWDM optics.

Add-Drop Multiplexers


● An add-drop multiplexer diverts one or more optical channels, and may add new signals in their place.

A full demultiplexer separates all the optical channels in a fiber, but in many cases you may want to separate only one or two channels from a larger number of channels. This is the job of an *optical add-drop multiplexer* (OADM), like the one shown in the middle of Figure 15.1.

To pick off one optical channel, you pass the light through an optical device that treats the selected channel differently from other channels, such as a wavelength-selective filter. If it reflects the selected channel, it should transmit all other wavelengths; likewise, if it transmits the selected channel, it should reflect all other wavelengths in the system. Figure 15.4 shows how an add-drop multiplexer reflects a selected channel while transmitting other channels, so it can drop one optical channel at a local node while transmitting all other channels. Other optics may be used to add a signal at the same wavelength to replace the dropped channel.

FIGURE 15.4

Add-drop multiplexer reflects one channel and transmits others.



So far we have assumed that an add-drop multiplexer is a fixed device that always adds and drops the *same* wavelengths at the same point. It is possible to build *reconfigurable add-drop multiplexers (ROADMs)*, using tunable or active optics. However, they aren't common, and we won't stop to explore their design.

Wavelength Routing

Wavelength routing is another type of optical demultiplexing that works in a different way than the optical demultiplexers we've described so far. Instead of separating optical channels and sending each one to a separate receiver, as shown in Figure 15.1, wavelength routing separates the optical channels and routes them to different destinations. You can visualize this as sending signals at each wavelength to a different town. (As you will learn later, this is not the same as the routing used in Internet transmission.)

● A *wavelength router* directs signals according to their wavelength.

Channel Equalization

Not all wavelength-selective optics in WDM systems are used to separate optical channels. Optical amplifiers do not have uniform gain across all wavelengths in their operating bands, so they amplify some channels more than others. This is a minor problem if the system has only one amplifier, but it can build up to a serious problem if the system has a long chain of amplifiers.

Channel equalization optics balance this gain differential by attenuating the wavelengths that are amplified the most, so overall gain is uniform across the system's operating range. Alternatively, they may include supplementary optical amplifiers that amplify the weaker wavelengths.

Advanced Optical Networking

Plans to develop advanced "all-optical networks" came to little when the telecommunications bubble collapsed, but some progress was made on the technology. It's worth noting a few WDM-related ideas that eventually may be implemented in future networks.

- *Optical wavelength conversion* that uses an optical device to shift a signal from one wavelength to another for transmission through other parts of the network.
- *Wavelength tuning* during network operation, so signals could be shifted to other wavelengths remotely by adjusting the transmitter.
- *Dynamic equalization* of optical channel strength across the operating range when the network configuration changes. Adding new optical channels changes the response of optical amplifiers, requiring technicians to adjust them. Dynamic equalizers would measure the change and adjust equalizing filters to compensate.
- *Reconfigurable add-drop multiplexers*, mentioned briefly above, which could automatically adjust themselves, or could be controlled remotely, to adapt to changing transmission needs, such as dropping signals at a new location.
- *Dynamic gain adjustment*, which would automatically adjust gain of optical amplifiers when new optical channels were added to a system.

Optical Filters and WDM

The optical devices most often used to selectively transmit certain wavelengths are called *filters*. The term covers a broad range of devices, including the attenuators described in Chapter 14, and you should understand what they are and how they work. Filters play important roles in WDM systems, although other technologies also may be used.

Sunglasses are a familiar type of optical filter, and like the filters used for WDM, sunglasses come in many varieties. Ordinary gray-green sunglasses are simple attenuators that block a uniform fraction of the light across the spectrum, and don't obviously change the colors of the world. Polarizing sunglasses transmit light of only one polarization, blocking the other polarization. The world doesn't look obviously different through polarizing sunglasses unless you look at certain parts of the sky or surfaces that look unusually bright or dark. Colored sunglasses and some photographic filters make the world look colored because they block other shades. Thus blue or red sunglasses make other objects seem to be those colors.

In the world of optics, "filter" often is a broad term applied to components that filter out part of the incident light and transmit the rest. Many types, such as photographic filters and most sunglasses, absorb the light they don't transmit. The only places such absorbing filters are used in fiber-optic systems are where it's important to absorb undesired light, such as in attenuators and optical isolators. In WDM systems, the wavelengths that are not transmitted through the filter normally are reflected so they can go elsewhere in the system. Such filters are like mirror shades and one-way mirrors, which reflect most incident light, but transmit enough for you to see through them (if you're looking into a brighter area).

The term "filter" is used a little differently in WDM. Typically it means one specific type of filter, the *interference filter*, which I describe below. Other types of WDM optics such as fiber Bragg gratings *act* like filters, in the sense that they block some light, but they are not considered quite the same. I'll use interference filters as a way to explain the operation of WDM optics, then describe other types of wavelength-selective optics.

Interference Filters

Interference filters are made by depositing a series of thin layers of two materials with different refractive index on a flat piece of glass. Alternating layers are deposited of each material. Typically the materials are insulators or dielectrics, which do not conduct electricity, so these filters are sometimes called *dielectric filters*.


The difference in refractive index between the two layers causes reflection at each surface. (The basic phenomenon is the same as the Fresnel reflection that occurs in connectors with an air gap between the fibers.) The more identical pairs of alternating layers, the more the reflection builds up—at most wavelengths.

Light is transmitted only at certain wavelengths, which are selected by the optical characteristics of the layers. As in a resonant laser cavity, the light wave has to make a round trip between the layers in an integral number of wavelengths. Light waves at these wavelengths are in phase with each other, so they add constructively in the transmitted beam. Transmitted wavelengths λ are given by the formula

$$N\lambda = 2nD \cos \theta$$

WDM filters transmit selected wavelengths and reflect others.

The layers in an interference filter selectively transmit a narrow range of wavelengths, and reflect other light.

**FIGURE 15.5**

Wavelength selection in an interference filter.

where N an integer, n is the refractive index of the layer, D is the layer thickness, and θ is the angle the incident light makes to the normal. Figure 15.5 shows how an interference filter transmits light with one wavelength per pair of layers (or one-half wave per layer), and reflects other wavelengths. Note that the wavelength transmitted depends on layer thickness, refractive index, and angle of incidence on the filter.

From an optical standpoint, the transmitted wavelengths are in phase and interfering constructively, so the waves add in intensity. Waves at other wavelengths are out of phase, so they interfere destructively, canceling their amplitude in the transmitted beam. Instead they are reflected.

From the user's standpoint, these effects are analogous to the wavelength selection effects of fiber gratings covered in Chapter 7. However, there is an important difference. Fiber gratings selectively *reflect* a narrow range of wavelengths, while interference filters selectively *transmit* a narrow range of wavelengths. This becomes important in designing demultiplexers, as you will learn later in this chapter.

The precise selection of wavelengths transmitted and reflected, and the shape of the reflection and transmission curves, depend on details of filter design including thicknesses and compositions of the layers, and the numbers of layers in the "stack." Normally the more layers, the finer the resolution, and the narrower the range of wavelengths selected.

The design of interference filters is a well-developed and highly specialized art, and it can produce carefully controlled results. With the right choice of material compositions and layer thicknesses, engineers can coat thin glass plates with interference filters that strongly reflect one wavelength while transmitting almost all the light at nearby wavelengths, so they are widely used in optical multiplexers and demultiplexers.

Although interference filters are considered “bulk optics” because they are discrete components, the filters used in WDM optics are quite small, typically a few millimeters across.

Line, Band, and Cutoff Filters

Interference filters can be made with various transmission characteristics by adjusting their composition, the thickness of layers, and the number of layers. Three types of filters important in WDM optics are the *line filter*, the *band filter*, and the *cutoff filter*, Figure 15.6 shows their transmission characteristics.


Line and band filters either reflect or transmit light in a selected range of wavelengths. If the range of wavelengths is narrow, they are called *line filters*; an example would be a filter to pick out one 100-GHz optical channel. Filters that select a broader range of wavelengths are called *band filters*; an example would be a filter that selects a 10-nm chunk of the erbium-fiber amplifier band. Line-rejection or band-rejection filters reflect the selected band while transmitting nearby wavelengths; line-pass or band-pass filters transmit the selected wavelengths while reflecting adjacent wavelengths.

Note that you can arrange filters in various ways so one type can serve different functions. For example, a line-pass filter normally transmits one wavelength while reflecting other light. However, you can use a line-pass filter as a line-rejection filter simply by collecting the *reflected* light rather than the transmitted light.

Line and band filters select a range of wavelengths.

FIGURE 15.6

Transmission of line, band, and cutoff filters.



Cutoff filters are designed to make a sharp transition between transmitting and reflecting at a certain wavelength. For example, a filter designed to separate optical channels directed to C- and L-band erbium-fiber amplifiers would have a cutoff wavelength at 1567 nm, with shorter wavelengths reflected to the C-band amplifier and longer wavelengths transmitted to the L-band amplifier. The cutoff filter in Figure 15.6 transmits short wavelengths but blocks long ones.

Filter transitions don't have to be sharp; in fact, it's easier to make filters with more gradual transitions because they don't require as many layers. Such gradual cutoff filters can be used to divide widely separated wavelengths, such as the pump band of an erbium optical amplifier and the wavelengths it amplifies.

These filters are made for use at specific wavelengths. Outside that range, their transmission may vary. Thus you can't be sure that a cutoff filter that reflects wavelengths shorter than 1567 nm will also reflect light at 1300 nm unless you have checked its properties at the shorter wavelength.

Equalizing Filters

Earlier we mentioned the need to compensate for the uneven gain of optical amplifiers. This normally is done by equalizing filters, which attenuate the wavelengths that are amplified most strongly. As shown in Figure 15.7, attenuation by the filter offsets the extra gain of the amplifier. The higher the gain, the higher the compensating attenuation. For example, if amplifier gain is 2 dB higher at 1535 nm than at 1550 nm, the filter should transmit 2 dB more light at 1550 nm. The filter may be placed before or after the amplifier, but the idea is the same. After the light passes through both filter and amplifier, the output power should be uniform across the gain band, as shown in Figure 15.7.

Fixed and Tunable Filters

The standard interference filters described above always transmit light in the same way as long as light strikes them at the same angle. Such fixed optical filters are fine for many applications, but tunable filters also are attractive for use in instruments or systems that require adjustment. A few different approaches are possible.

One simple approach is to tilt an interference filter, because the wavelength it selects depends on the angle at which light strikes it. Another is to use a prism or diffraction grating to spread out a spectrum and pick a narrow range of wavelengths from that spectrum. However, neither approach has proved able to meet the stringent requirements of dense wavelength-division multiplexing.

A more complex approach is to move an interference filter that is made so its optical characteristics vary along its length. This approach can meet high-resolution requirements, but requires special filters and mechanical movement of the filter.

A more common approach is the Fabry-Perot interferometer, which is essentially an optical cavity similar to those used as laser resonators, but without the laser medium inside. It consists of two partially transparent mirrors aligned parallel to each other, so light bounces back and forth between them. Normally air fills the space between them. Light bounces back and forth many times once it enters the cavity, so interference effects select wavelengths that


Cutoff filters make a sharp transition between transmitting and reflecting at a certain wavelength.

Equalizing filters compensate for uneven gain of optical amplifiers.

Some filters can be tuned to select different wavelengths.

A Fabry-Perot interferometer can be a tunable filter.

FIGURE 15.7
*Equalizing filter
 balances uneven
 gain in a fiber
 amplifier.*



resonate in the cavity. That is, an integral number (N) of wavelengths λ equals a round-trip distance in the cavity ($2L$):

$$2L = \frac{N\lambda}{n}$$

where n is the refractive index of the material between the mirrors, a quantity needed to account for the difference between the wavelength in empty space and the material.

The cavity transmits light at wavelengths that match this resonant condition, like an interference filter. In fact, the Fabry-Perot interferometer is just a simple version of an interference filter, with a single cavity instead of a stack of layers. Normally the Fabry-Perot cavity is short, so the spacing between wavelengths is large. Adjusting the cavity length changes the wavelength selected, tuning the filter. You adjust the length either by moving the mirrors or by tilting them so light follows a longer path between the mirrors.

Another common approach is the acousto-optic filter, where acoustic waves travel through a transparent material such as glass. The atomic vibrations produced by the acoustic waves create regions of higher and lower density within the glass. The denser regions have higher refractive index, creating a multilayer structure in the glass. As in an

Acoustic waves can create density waves in glass; changing the sound frequency adjusts the wavelength selected.

interference filter or fiber grating, these regular high-index zones selectively scatter light of certain wavelengths selected by the spacing between them. Tuning the acoustic frequency changes the grating spacing, and hence the selected wavelength—making a tunable filter.

The big advantage of tunable filters is obviously their tunability. Their disadvantages are much greater cost and complexity than fixed wavelength filters.

WDM Technologies

Interference filters are one of several technologies available for WDM optics. Each of these technologies has distinct characteristics that fit into certain application requirements. Some are easy to implement for demultiplexing a few optical channels, while others easily handle large numbers of channels. Some work better for narrow optical channel separation than others. All achieve the basic goal of optical multiplexing and demultiplexing for some applications. Let's look at these diverse technologies one by one.

Several technologies are available for WDM optics.

Interference Filters for WDM

Using interference filters for WDM requires taking light out of the fiber and passing it through a set of filters that sorts the light out by wavelength. Typically a lens collimates or focuses the light emerging from the input fiber, which then passes through one or more filters. When the demultiplexing is finished, separate lenses collect the separated optical channels and focus them into individual output fibers.

A narrow-line interference filter typically transmits a single optical channel while reflecting other wavelengths. Several interference filters can be cascaded to pick off a series of six wavelengths, as shown in Figure 15.8. The first filter transmits channel λ_1 while

Cascaded interference filters can pick off one wavelength at a time for demultiplexing.





FIGURE 15.8

Interference filter WDM picks off one wavelength at a time.

FIGURE 15.9

Demultiplexing 40 channels by separating blocks of channels.



reflecting all other channels. The remaining channels hit the second filter, which transmits channel λ_2 while reflecting the four remaining channels. In this arrangement you need $n - 1$ filters to isolate n optical channels.

The concept is simple and straightforward, but interference filters are not perfect. Although they reflect *virtually* all the incident light at other wavelengths, some is lost, and these losses add up after a series of reflections. Picking off one wavelength at a time works fine for 8 channels, but the losses could grow excessive if you have 16 or 32 channels.

To prevent such losses, optical signals can be divided into groups of channels, which are then split up individually. Figure 15.9 shows such a system built from high- and low-pass filters plus 8-channel demultiplexers that pick off one channel at a time, as in Figure 15.8. In this scheme, incoming light first hits a high-pass filter, which reflects all light with wavelength less than λ_{17} . The shorter wavelengths are diverted to a low-pass filter, which reflects light with wavelengths longer than λ_8 . Each of those sets of 8 channels is directed to an 8-channel demultiplexer. Wavelengths from λ_{17} to λ_{40} are routed to another low-pass filter, which reflects all light with wavelengths greater than λ_{24} . Channels λ_{17} to λ_{24} then go to an 8-channel demultiplexer, while the longer wavelengths are sent to another long-pass filter, which splits them into 8-channel groups for demultiplexing.

This approach does not reduce the *total* number of filters needed, but it does reduce the number of filters any optical channel is going to encounter before reaching a receiver. The upper limit for the configuration shown in Figure 15.9 is 10.

The arrangement shown in Figure 15.9 has another important advantage. It's modular, so you don't need to start with all 40 channels. You could start with channels λ_{17} to λ_{24} , then add the high- and low-pass filters to split off other wavelengths as you needed to add the extra channels. This is the way telephone companies like to work, adding capacity only as they need it. They save money in the short term, and retain the option for expanding capacity in the long term.

Filters can divide optical channels into groups, then separate the groups into individual channels.

Interference filters are widely used for WDM, and it's worth reviewing their advantages. First, the underlying technology is well developed. Interference filters have been around for many years, although the extremely narrow-line filters used in DWDM systems were only developed very recently. Filters can be made very small—a few millimeters across—a good match for fiber-optic systems. They have good performance and can be assembled in modular units, so users can upgrade their systems several channels at a time, instead of jumping from 1 to 40 channels. On the down side, you need roughly as many filters as you have optical channels—adding to costs, complexity, and optical losses.

Remember also that interference filters do not always have to separate every single wavelength out of an optical signal. A single optical filter could transmit a single optical channel in an add-drop multiplexer, with the remaining channels reflected and collected for transmission through the rest of the system.

Fiber Bragg Gratings and Optical Circulators

Fiber Bragg gratings can be grouped together in ways similar to interference filters, but they have some significant functional differences. Generally they reflect a single selected wavelength and transmit the rest, as shown earlier in Figure 15.10. Recall that interference filters instead generally *transmit* the selected wavelength.

This ability of fiber gratings to reflect a selected wavelength can be used to build an optical demultiplexer that resembles one based on interference filters. The important differences come from the facts that a fiber grating is a fiber, and that it reflects rather than transmits the wavelength it selected.

We'll start with the example shown in Figure 15.10, where the input signal includes eight wavelengths, 1546, 1548, 1550, 1552, 1554, 1556, 1558, and 1560 nm. The input signal enters the fiber grating through an optical circulator, described in Chapter 14, which acts like an optical traffic circle, transferring the signal from the input port to the grating. The grating transmits all the wavelengths except 1552 nm, which it reflects back to the optical circulator. The circulator is a directional device, so it transmits the light reflected from the fiber grating onto the drop port at bottom, where the light exits into another fiber that carries the signal to the 1552-nm receiver.

In a demultiplexer, the remaining seven input signals pass through the grating and become the input signal to a second fiber grating, which selectively reflects one of the remaining wavelengths. That wavelength is dropped to the appropriate receiver, and the remaining six channels continue, with one wavelength being picked off at a time. Like an




FIGURE 15.10


Fiber Bragg grating filter reflects the one wavelength it selects.

Filter WDMs can be upgraded modularly, but require about as many filters as the system has channels.

Fiber gratings reflect the selected wavelength and transmit other wavelengths.

FIGURE 15.11

A single fiber grating with optical circulators at both ends serves as an add-drop multiplexer.



Fiber gratings are natural choices for add-drop multiplexers.

Fiber gratings can serve as optical filters as well as WDM filters.

interference-filter demultiplexer, the fiber Bragg grating demultiplexer requires seven stages to separate eight channels.

Fiber gratings are a natural choice for use in add-drop multiplexing because the grating reflects the same wavelength when it enters from either end, so it can perform both add and drop functions, as shown in Figure 15.11. In this case, the fiber is carrying 1550, 1552, 1554, and 1556 nm. The 1556-nm signal is carrying information from town A to town B, while the others are carrying data from town A to town C. The system also is adding another 1556-nm signal from town B to town C. The grating first reflects the 1556-nm signal so it can be dropped at town B, transmitting 1550, 1552, and 1554 nm so they go to town C. The input signal at 1556 nm added at town B is coupled to the other end of the fiber grating through a second optical circulator. The grating reflects this signal back in the opposite direction, adding it to the original 1550, 1552, and 1554 nm signals on their way to town C through the circulator.

Fiber gratings have good wavelength selectivity, and can be used to demultiplex signals with channel spacing as fine as 25 GHz. However, a drawback for demultiplexing is that separation of the reflected wavelength requires optical circulators, which are complex and expensive. That leads to fiber gratings being picked for the narrow-band applications, where their selectivity is particularly important, while interference filters are used for applications, where narrow-band selectivity is not critical.

The fiber-grating principle can be used to make other types of optical filters as well as WDM filters. Fiber gratings can be made with multiple grating sections along their length, each reflecting a different wavelength. As you learned in Chapter 7, one application is an optical delay line that compensates for chromatic dispersion in transmission fibers. Another is in an add-drop multiplexer that picks off two or more channels. Other designs allow fiber gratings to have complex attenuation properties spanning a range of wavelengths so they can serve as gain-equalization filters.

The fiber geometry of the fiber grating is at best a mixed blessing. As a fiber, it's easy to connect to transmission fibers. However, as a fiber it also reflects light back in the direction that the input signal arrived, so demultiplexers require an optical circulator to extract the selected wavelength. That's a problem because optical circulators are complex and expensive devices that can add considerably to system costs.

Fused-Fiber Couplers

The fused-fiber couplers described in Chapter 14 are inherently sensitive to wavelength. As in waveguide couplers, the amount of light transferred between the fused fibers depends on the length of the coupling region, as measured in wavelengths. Over some characteristic distance, the light is transferred completely from one output to the other. This distance is longer when measured in shorter wavelengths, because more of them fit into the same distance, opening a way to separate wavelengths.

The process works best for two wavelengths that are not closely spaced, so fused-fiber couplers are not used for separating optical channels in dense-WDM systems. However, it works fine for widely spaced wavelengths, like the pump and signal wavelengths in erbium-doped fiber amplifiers, as shown in Figure 15.12. Light initially enters the top of the two

Fused-fiber couplers can separate wavelengths by directing them out different ports.




FIGURE 15.12

Fused-fiber coupler splits two wavelengths.

fused fibers. Gradually, the light shifts to the bottom fiber. If the fused region is long enough, all the light transfers into the lower fiber, and the process starts over again, this time shifting from the bottom to the top. The degree of shifting depends on how many wavelengths the light has travelled, so shorter wavelengths shift back and forth first, with longer wavelengths following. In Figure 15.12, the 980-nm light shifts from the top to the bottom fiber and back to the top at the end of the fused region, while the 1550-nm light has shifted only from the top to the bottom fiber. This process completely separates the two wavelengths.


Mach-Zehnder Interleavers

Mach-Zehnder interferometers interleave wavelengths, separating odd and even optical channels.

Fused-fiber couplers are essential components in another type of device for dense wavelength-division multiplexing called a *fused-fiber Mach-Zehnder interferometer*, or more simply an *interleaver*. Unlike other multiplexers and demultiplexers, these devices split groups of evenly spaced optical channels into sets of odd and even channels, by using the interference of light in a fiber structure. To see how they work, we'll start with a look at the concept of a Mach-Zehnder interferometer, named after the physicists who invented it.

Interferometers pass light waves along two different paths, causing interference between the waves. As you learned earlier, coherent light waves can add or subtract their amplitude, producing constructive or destructive interference. In a Mach-Zehnder interferometer, a device called a *beamsplitter* splits an input beam into two parts, which pass along different routes, then are combined in a second beamsplitter. Figure 15.13 shows how this can be done with fused-fiber couplers serving as the beamsplitters. Input enters through one fused-fiber coupler, where it is divided between two fibers that form the arms. Light in the two arms recombines in a second fused-fiber coupler.

FIGURE 15.13
Fused-fiber coupler Mach-Zehnder interferometer interleaves wavelengths.



The relative phase of the light emerging from the interferometer arms determines its distribution between the two outputs of the output coupler. This phase depends on wavelength as well as the length of the arms. As the wavelength changes, the distribution of light between the two arms changes. In the case shown in Figure 15.13, at λ_1 all the light emerges from the top arm, but at λ_2 all the light emerges from the bottom arm. Every time the wavelength changes by that increment, the output light shifts between arms of the second fused-fiber coupler.

The increment in wavelength that shifts the output depends on the difference in effective length ΔL between the two interferometer arms. If we assume the two arms have the same refractive index n , this equals

$$\Delta L = \left(2n \left[\frac{1}{\lambda_1} - \frac{1}{\lambda_2} \right] \right)^{-1} = \frac{c}{2n\Delta\nu}$$

(The two arms could have different refractive indexes, which would complicate things.) Note that it's simpler to express the difference in terms of the change in *frequency*, $\Delta\nu$, rather than a change in wavelength. In fact, the increment in spacing of optical channels is uniform only if the spacing is measured in terms of a change in frequency $\Delta\nu$ rather than a change in wavelength. Thus the signals are at frequencies ν_1 , $\nu_1 + \Delta\nu$, $\nu_1 + 2\Delta\nu$, $\nu_1 + 3\Delta\nu$, and so on. The difference between taking increments in terms of wavelength and in terms of frequency is small, but can be significant if you're designing systems. I label optical channels by their wavelength for convenience, but remember that actual DWDM channel spacing is uniform in frequency, *not* wavelength.

Interleavers essentially split odd and even optical channels. Thus in Figure 15.13, signals at λ_1 , λ_3 , λ_5 , and λ_7 emerge entirely from the top output, while signals at λ_2 , λ_4 , λ_6 , and λ_8 emerge entirely from the bottom output. This makes a fused-fiber Mach-Zehnder interferometer an effective wavelength interleaver that can demultiplex a set of uniformly spaced optical channels. The interleaver also can work backwards, shuffling odd and even optical channels together.

A single interleaver does not completely demultiplex the signals. In the example of Figure 15.13, you still have λ_1 , λ_3 , λ_5 , and λ_7 in the top output, and λ_2 , λ_4 , λ_6 , and λ_8 in the bottom output.

Those channels also must be separated, but you can do that by repeating the process, as shown in Figure 15.14. In this case, the first interleaver has the finest resolution—100 GHz (about 0.8 nm) in the example. The next interleaver needs to split the remaining channels, so it splits channels twice as far apart—200 GHz. The final interleaving demultiplexer must split the two channels left in each output fiber, so it needs 400-GHz spacing (about 3.2 nm).


Note that Mach-Zehnder interferometers can be built with planar-waveguide technology as well as with fused-fiber couplers. So far, the main application of planar waveguides has been in the multi-arm arrayed waveguide devices described below, but two-arm Mach-Zehnder interleavers also can be made using planar waveguides.

As you can see, interleaving is a different process than picking one signal off at a time with an interference filter or fiber Bragg grating. Demultiplexing eight channels requires the same number of components—seven—whether you use interference filters or interleavers. However, with interleavers all signals pass through three components; with the

Incrementing frequency of an optical channel by $\Delta\nu$ shifts output between arms of the interleaver.

Interleavers separate or combine signals in several stages.

FIGURE 15.14
Interleavers demultiplex optical channels in stages.



interference filter demultiplexer shown in Figure 15.8, the final two channels have to pass through all seven filters, so they experience more loss than the other channels.

The differences between interleavers and other demultiplexers are important to remember. As you will see later, these differences are vital in designing hybrid WDM systems that take advantage of the strengths of two (or more) different approaches.

Bulk Diffraction Gratings


Diffraction gratings separate wavelengths by spreading out a spectrum.

I mentioned earlier that a diffraction grating—a series of parallel grooves or lines—diffracts light in a way that spreads out a spectrum. Interference between light waves scatters different wavelengths from the grating at different angles. The gory details of the optical physics aren't important here; what matters is that the wavelengths spread out, like a rainbow. You can see the same rainbow effect in a CD if you tilt it back and forth while looking at light reflected from it. The pits that store data on the CD are arranged in grooves that wind in a tight spiral around the disk, forming sets of parallel spots that act like a diffraction grating.

You can use the same effect to separate wavelengths, with suitable optics to focus the input light, collect the reflected light, and focus it into the output fibers.

Figure 15.15 shows a grating demultiplexer. This device uses a gradient-index (GRIN) rod lens in which the refractive index varies through a block of solid glass, producing the same focusing effect as a standard lens. (A GRIN lens is easier to align than a standard lens for this application.)

Input on three optical channels, λ_1 , λ_2 , and λ_3 , enters through the bottom fiber. The GRIN rod focuses the input light onto a diffraction grating at the back end of the rod, which is set at an angle to reflect the light at the right angle. The grating diffracts each wavelength at a different angle, and the GRIN rod focuses each wavelength onto an output fiber. When everything is properly aligned (which, of course, is a big part of the job),

**FIGURE 15.15**

A grating coupler with GRIN rod separates three wavelengths.

λ_1 emerges from the top output fiber, λ_2 emerges from the middle output fiber, and λ_3 emerges from the lowest output fiber.

Such simple bulk diffraction grating demultiplexers work well for separating a few wavelengths that are widely spaced, but they don't give high channel isolation between closely spaced wavelengths. However, the way that diffraction gratings spread out a continuous spectrum of wavelengths is an advantage for measurement instruments. If you want to measure the distribution of power as a function of wavelength, you usually want resolution finer than you need to look at a single channel. Scanning a continuous spectrum gives better resolution, and a diffraction grating makes this spectrum available. Thus measurement instruments are likely to use diffraction gratings to spread out the spectra that they measure. (The same is true for optical performance monitors, which measure the distribution of optical power across the spectrum in communication systems to check that all channels are operating properly.)

Special diffraction gratings called *echelle gratings* offer higher resolution than ordinary gratings, which makes them potentially attractive for use in DWDM. That technology and other grating multiplexers are still in development.


Arrayed Waveguide Gratings

Another way to demultiplex WDM signals is to pass them through an array of planar waveguides, as shown in Figure 15.16. Like a diffraction grating, an *arrayed waveguide grating (AWG)* diffracts light at angles that depend on the wavelength, so the technology also can be used for other applications including dynamic gain equalization, reconfigurable optical add/drop multiplexers, and wavelength-selectable lasers. You can think of an arrayed waveguide grating as a diffraction grating built using planar waveguide technology.

The central element is an array of narrow curved waveguides running beside each other between a pair of mixing regions or coupling zones. Input signals enter the first mixing region (at left in Figure 15.16), where they are coupled into the curved waveguides running to the second mixing region. The adjacent waveguides differ in length by an increment ΔL , which is much larger than the wavelength λ . As a result of this difference, light passing in adjacent waveguides between the two mixing regions has a phase shift of $n\Delta L/\lambda$, where n is the refractive index of the waveguide. This phase shift acts like a diffraction grating to disperse light across a range of angles in the second mixing region.

The second mixing region acts like a lens to focus the diffracted light onto a series of output ports. Interference effects combine with the refraction and diffraction to disperse a

An AWG is a planar waveguide version of a diffraction grating.

FIGURE 15.16*Array waveguide demultiplexer.*

spectrum of light across the output ports, so each port collects a limited range of wavelengths. In versions built for demultiplexing, each port collects wavelengths corresponding to one standard WDM optical channel. The number of output ports is the number of channels separated, with the free spectral range of the device equaling the channel spacing times the number of channels. The number of waveguides running between the two mixing regions is larger than the number of channels; for example, a 64-channel waveguide AWG has an array of 232 waveguides linking the mixing regions.

Arrayed waveguides can be made from silica, plastic, silicon, or III-V semiconductors such as indium phosphide. The devices are monolithic and can be integrated with other components such as optical switches. Like integrated electronic circuits, the master is expensive, but replication is relatively cheap. From a functional standpoint, an arrayed waveguide is a monolithic device that can separate many optical channels simultaneously. Filters require complex assemblies of many discrete components to do the same job.

Arrayed waveguides are used mainly in demultiplexing optical systems with high channel counts, where the cost per channel is much lower than with other WDM technologies. Laboratory versions have been made containing hundreds of channels spaced as little as 10 GHz apart, but typical commercial versions have 32 or 40 channels spaced at 50 or 100 GHz. AWGs can be used in reverse for multiplexing.

Standard arrayed waveguides have peak transmission in the center of the band, which declines gradually to the sides, as in Figure 15.3. Often some other optics are added to

●
AWG demultiplexers are best for high channel counts.

flatten the peak transmission and make the sides drop faster to prevent crosstalk. For example, an interleaver can split signals between a pair of AWGs so one receives the odd channels and the other the even channels.

One inherent limitation of arrayed waveguides is a high insertion loss arising from mismatches of transmission modes in the waveguides and in the mixing region. AWGs also are sensitive to temperature, and some materials—notably III-V semiconductors—are sensitive to polarization.

Array waveguide technology lends itself to other applications that require wavelength selectivity. For example, a pair of AWGs can be integrated with other devices to separate the wavelengths in an input signal, process the separated wavelengths, then recombine them in a multiplexer. This approach can be used to make a reconfigurable optical add/drop multiplexer, with switches placed in the path individual optical channels take between the demultiplexer and multiplexer. Variable filters could be put in the same position to make a dynamic gain equalization filter, which changes transmission to compensate for fluctuations in amplifier gain.

Building Multiplexers and Demultiplexers

So far you've learned the general principles of WDM optics. A few other considerations go into building actual multiplexers and demultiplexers.

You should realize that these technologies are building blocks. Actual systems may integrate two or more WDM technologies into a single multiplexer or demultiplexer. These hybrid WDMs can take advantage of the best features of the different technologies. For example, you could use a Mach-Zehnder interleaver as the first stage, to break up an 16-channel signal with 100-GHz spacing into two 8-channel signals with 200 GHz spacing. Then interference filters with 200-GHz resolution could break up the 8-channel signals. This could cut costs significantly because 200-GHz filters are much less expensive than 100-GHz filters. Similarly, two interleaving stages—one at 50 GHz, the other at 100 GHz—could break down a 32-channel signal with 50 GHz spacing to four separate 8-channel signals with 200-GHz channel spacing.

In talking about general principles, it's easiest to assume that channels are uniformly spaced, and that every slot is filled. This does not happen in general. The spacing of channel slots generally is uniform, but the channels may be grouped into blocks, and not all of the slots may be filled. For example, one or two channel slots may be left between 8-channel blocks in a system like the one shown in Figure 15.9, if the band-pass filters do not have sharp enough cutoffs. In practice, a gap of about 5 nm normally is left between the erbium fiber amplifier C-band at 1530 to 1565 nm and the L-band at 1570 to 1620 nm.

The ability to add more channels at reasonable cost is a major practical concern. Telecommunications carriers want room to increase their transmission capacity, but they don't want to pay for large amounts of equipment they can't use immediately. For example, they might like to build the 40-channel system shown in Figure 15.9 in increments of eight channels at a time. This has to be traded off against the possibly lower overall cost of installing a larger system all at once.

Finally, remember that not every possible channel slot has to be used. Many market analysts during the telecommunications bubble wrongly assumed that carriers were buying

Multiple technologies can be integrated in a single multiplexer or demultiplexer.

lots of WDM systems and filling all of the wavelength slots. That led to wildly inflated estimates of the growth of network capacity.

A prudent carrier may plan to use 16 or 32 channels on a single fiber eventually, but at first they are more likely to start with a couple of channels. They may install only the optics they need for two channels, then plan to add more modules later. Or if the price is right, they may go for an 8-channel optics package and fill the remaining slots later. Like spare fibers in a cable, the extra optical channels don't cost much until the transmitters, receivers, and optical amplifiers are installed. Network planners like to leave room for future growth. Like installing new electrical wiring in your house, you plan not just for the equipment you already own, but for what you expect to add in coming years.

What Have You Learned?

1. WDM optics combine optical channels at the input end of a system and separate them after transmission through a fiber. Multiplexers combine channels; demultiplexers separate them.
2. An add-drop multiplexer goes in the middle of a system. It can both drop existing channels at an intermediate point and add new channels to a fiber carrying WDM signals. The added signal can replace the dropped channel.
3. Channel density depends on spacing between channels. Standard spacings for dense-WDM spacings are 200, 100, and 50 GHz. Coarse-WDM channel spacing is 20 nm.
4. Wavelength routers direct different wavelengths to different points.
5. Separating pump wavelengths from the outputs of fiber amplifiers also requires wavelength-division multiplexing.
6. WDM filters transmit selected wavelengths and reflect others.
7. An interference filter uses multiple thin layers to selectively transmit a narrow range of wavelengths; others are reflected. Interference filters can select a very narrow range of wavelengths.
8. Cutoff filters make a sharp transmission between transmitting and reflecting at a certain wavelength.
9. Equalizing filters compensate for the unequal gain of optical amplifiers across their operating ranges.
10. Most filters have fixed wavelength response, but some can be tuned to transmit different wavelengths. They include acousto-optic filters and Fabry-Perot interferometers.
11. Interference filters can select closely spaced optical channels. A cascaded series of interference filters can pick off one wavelength at a time to demultiplex optical channels. Each filter transmits one channel and reflects the rest.
12. Interference filters can select groups of optical channels as well as individual channels.

13. Fiber Bragg gratings reflect the selected wavelength and transmit other wavelengths. They must be used together with optical circulators for demultiplexing, but have high resolution in selecting optical channels.
14. Fused-fiber couplers can separate wavelengths by directing them out different ports, but their resolution is limited.
15. Mach-Zehnder interferometers interleave wavelengths, directing alternating channels out of each of two inputs. They are sometimes called interleavers, and can separate closely spaced optical channels. Interleavers separate channels in a series of stages.
16. Bulk diffraction gratings separate wavelengths by spreading out a spectrum. They are often used for measurement instruments.
17. Arrayed waveguides are planar waveguide devices that disperse light by its wavelength, like diffraction gratings. The arrays are monolithic, and can separate closely spaced optical channels. They are most economical for 32 or 40 optical channels.
18. WDM systems may combine multiple technologies for demultiplexing.

What's Next?

Chapter 16 will cover optical switches, optical modulators, and other active devices used in optical networks.

Further Reading

- J. Capmany et al. ed., special issue on “Arrayed Grating Routers/WDM Mux Demuxs and Related Applications/Uses,” *IEEE Journal Selected Topics in Quantum Electronics* 8, (November/December 2002)
- Kenneth O. Hill, “Fiber Bragg Gratings,” Chapter 9 in Michael Bass, ed., *Handbook of Optics Vol. IV: Fiber Optics & Nonlinear Optics* (McGraw-Hill, 2001)
- Rajiv Ramaswami and Kumar N. Sivarajan, *Optical Networks: A Practical Perspective* (Morgan Kaufmann, 2002)

Questions to Think About

1. An erbium-doped fiber amplifier can transmit signals at wavelengths between 1530 and 1565 nm. How many optical channels can you fit in this range with 200 GHz spacing? How many channels with 100 GHz spacing?
2. A transatlantic fiber-optic cable contains 100 optical amplifiers. It needs equalizing filters to balance the gain of the erbium-fiber amplifiers across their

operating ranges. If the receivers used on the system have a dynamic range of 20 dB, how closely do the equalizing filters have to balance gain? Assume all filters and amplifiers are identical.

3. A typical interference filter for demultiplexing 100-GHz channels has 0.5-dB loss on the reflected channels and 2.0-dB loss on the transmitted channels. How much loss does the signal suffer on the first channel picked off (λ_1) in Figure 15.8? What is the loss for the last channel of eight channels picked off (λ_8) in a similar arrangement? What channel in an 8-channel system suffers the highest total loss?
4. A typical fiber Bragg grating has 99.9% reflection (0.0043-dB loss) and 0.2-dB loss for transmitted wavelengths. Assume loss of 1 dB in the optical circulator. What is the loss for the first channel of eight channels picked off in a cascaded series of fiber Bragg gratings? What are losses for the seventh and eighth channels?
5. What should the difference in path lengths be in a Mach-Zehnder interferometer designed to interleave optical channels separated by 50 GHz? By 200 GHz? Assume the refractive index of the material is 1.5 and is uniform for both arms.
6. You want to separate 16 optical channels that are uniformly spaced 50 GHz apart with optical interleavers. How many interleavers do you need? How many interleavers does each optical channel pass through?
7. A 40-channel arrayed waveguide demultiplexer has average loss of 8 dB for each channel processed. How does this compare to the highest loss of the 40-channel interference-filter demultiplexer shown in Figure 15.9? You can use the results from Question 3 to give you the loss for the 8-channel demultiplexing boxes. What's the minimum loss?

Chapter Quiz

1. What is the broadest channel spacing that is considered “dense” WDM?
 - a. 400 GHz
 - b. 200 GHz
 - c. 100 GHz
 - d. 50 GHz
 - e. 0.8 nm
2. What does an add-drop multiplexer do?
 - a. converts all optical signals in a fiber to electronic form
 - b. amplifies optical signals after attenuation has reduced signal strength below $1 \mu\text{W}$ per optical channel
 - c. adds and drops optical channels at intermediate locations without interfering with other signals on the fiber

- d. adds and drops optical channels at an intermediate point while regenerating other channels on the fiber
 - e. switches signals between different wavelengths at an intermediate point in the system
- 3.** What selects the wavelengths transmitted by an interference filter?
- a. the thickness and composition of layers deposited on glass
 - b. the composition of the glass plate on which it is deposited
 - c. coloring dyes added to the layers in the interference filter
 - d. parallel ridges formed in the uppermost layer
 - e. only the refractive index of the surface layer
- 4.** You want to reflect light at wavelengths longer than 1567 nm and transmit light at shorter wavelengths. What type of filter do you want?
- a. color filter
 - b. cutoff filter
 - c. band-pass filter
 - d. line filter
 - e. attenuation filter
- 5.** What type of filter is tunable in wavelength?
- a. interference filter
 - b. cutoff filter
 - c. band-pass filter
 - d. Fabry-Perot interferometer
 - e. line filter
- 6.** An acousto-optic filter is
- a. a type of cutoff filter.
 - b. a tunable filter.
 - c. an interference filter.
 - d. a color filter.
 - e. impossible to build.
- 7.** Interference filters
- a. reflect the selected wavelength and absorb other wavelengths.
 - b. reflect the selected wavelength and transmit other light.
 - c. transmit the selected wavelength and reflect other wavelengths.
 - d. transmit the selected wavelength and absorb other light.
- 8.** What type of WDM system requires an optical circulator?
- a. interference filters
 - b. fiber Bragg gratings

- c. Mach-Zehnder interferometers or interleavers
 - d. bulk diffraction gratings
 - e. tunable optical filters
- 9.** Fiber Bragg gratings
- a. reflect the selected wavelengths and absorb other wavelengths.
 - b. reflect the selected wavelength and transmit other light.
 - c. transmit the selected wavelength and reflect other wavelengths.
 - d. transmit the selected wavelength and absorb other light.
- 10.** What type of technology is used in an interleaver?
- a. cutoff filter
 - b. interference filter
 - c. fiber Bragg grating
 - d. Mach-Zehnder interferometer
 - e. Fabry-Perot interferometer
- 11.** How many interference filters do you need to make an 8-channel demultiplexer that picks off one channel at a time?
- a. 4
 - b. 6
 - c. 7
 - d. 8
 - e. 9
- 12.** How many fiber gratings do you need to make an 8-channel demultiplexer that picks off one channel at a time?
- a. 4
 - b. 6
 - c. 7
 - d. 8
 - e. 9

Optical Switches, Modulators, and Other Active Components

About This Chapter

Fiber-optic systems are no longer passive pipes that merely carry optical signals from point to point. Signals may be modulated in intensity, switched between fibers, shifted in wavelength, and modified in other ways. The potential of “optical networking” was over-promoted during the telecommunications bubble, but the need to operate on signals is growing with the volume of traffic and the size of the global telecommunications network. The devices that perform these operations on signals are called *active components*.

You have already learned about several types of active components. Light sources, transmitters, receivers, optical amplifiers, and wavelength converters are all active components. In this chapter you’ll learn about optical modulators, optical switches, and other components that dynamically affect optical signals.

Defining Active Components

The distinction between active and passive components can be useful, but it also can be confusing, so it deserves a bit of explanation. Originally, active components were those that required power from an outside source while passive components drew no

●
Modulators and
switches are
active
components.

outside power. By this definition, a laser or optical amplifier is obviously an active component, and an optical fiber or attenuator is obviously passive. Yet the source of power doesn't necessarily say much about function.

Another approach is to consider how components affect a signal. In this sense, a passive component always performs the same operation on a signal, such as dividing it in half or attenuating it by a fixed amount. In contrast, an active device can change its effect on the signal, such as by modulating its intensity or switching it between fibers. This approach has the advantage of focusing on what a device does to the signal, but has its own limitations. If strictly applied, it might class a laser that generates a steady beam as passive because it does not affect the signal strength.

In practice, it's best to consider components as active if they either draw power from outside or modify the signal in changing ways. The most important active components in current optical systems—modulators and switches—do both. Some haziness in the definition is acceptable. The distinctions change as the technology evolves, and the main purpose here is not critical—it's a way of organizing this book into chapters of manageable size.

Modulators and Modulation

Light must be modulated to transmit a signal. As you learned in Chapters 9 and 10, the simplest modulation technique is to directly change the drive current passing through a laser or LED. Unfortunately, direct modulation runs into a number of limitations as speeds increase. The modulation rate, average output power, and difference between “off” and “on” states are all limited. Direct modulation also can distort analog signals and shift the output wavelength of any signal, an effect called *chirp*, which adds to chromatic dispersion.

The importance of these limitations depends on the system design. If the signals are transmitted through one optical channel in a DWDM system spanning long distances, problems may appear with direct modulation at data rates around 1 Gbit/s. If the signals are sent a kilometer or two in a campus network using CWDM or only a single wavelength per fiber, direct modulation can work at 10 Gbit/s. In either case, when direct modulation does not meet performance requirements, you can turn to external modulation, in which a separate device called a *modulator* changes the intensity of the light from a constant laser or LED source. External modulation can be faster than direct modulation of a laser or LED, and does not affect the source wavelength.

An optical modulator changes how much light it transmits in response to an external control signal. Many types have been developed for other applications, but fiber-optic systems are particularly demanding because they require modulation at gigabit rates, much faster than most modulation mechanisms. For example, liquid crystal devices cannot respond fast enough for fiber-optic modulators, but are fine for laptop computer displays, which operate much slower.

Modern fiber-optic systems use two main families of modulators. *Electro-optic modulators* rely on changes in the way certain planar waveguides carry light. *Electro-absorption modulators* are semiconductor diodes that in their internal structure resemble lasers, but are switched between states that transmit and absorb light. We will look at them separately.

●
External
modulation
improves system
performance.

Electro-Optic Modulators

Electro-optic modulation depends on the *electro-optic effect*, a change in the refractive index of certain materials when an electric field is applied to them. The change affects light passing through the material virtually instantaneously. The velocity of light in a material is the speed of light in a vacuum divided by the refractive index, so increasing the refractive index slows down the light; reducing it speeds up the light. The change is proportional to the voltage applied to the material.

When you look at a waveguide, you measure the effect of this change in refractive index as a shift in the phase of the light waves compared to what the phase would have been without the applied voltage. A shift of half a wavelength— 180° —would leave the shifted light completely out of phase with the unshifted wave. The phase shift normally is measured by this comparison:

$$\text{Phase shift } (\Delta\Phi) = 180^\circ \times \frac{V}{V_{180}}$$

where V is the voltage applied to the modulator and V_{180} is the voltage needed to shift the phase a half-wavelength, or 180° .

Merely delaying the light modulates its phase but not its intensity. To modulate the intensity, an electro-optic modulator splits the input light equally between a pair of parallel waveguides. In the example shown in Figure 16.1, a modulated voltage is applied to one waveguide, but not to the other. This modulates intensity of the light where the two waveguides merge at the right. If the waveguides are equal lengths and the voltage is zero, the light waves are in phase when they combine, so the waves add constructively, producing a signal. The light is “on.” However, if you apply the voltage needed to delay the signal by 180° , the

Electro-optic modulators rely on changes in refractive index caused by an electric field.

Delaying the light phase causes interference effects that modulate the output intensity.





FIGURE 16.1
Simple electro-optic waveguide modulator.

FIGURE 16.2

Variation in modulator output with voltage.



light in the two waveguides is out of phase when they merge. The two waves interfere destructively, canceling each other out, and the output intensity is nominally zero.

In practice, a little light remains when the signal is nominally zero, a quantity measured by the extinction ratio, which compares the output power P_{on} in the on state with that in the off state P_{off} .

$$\text{Extinction ratio (dB)} = -10 \log \left(\frac{P_{off}}{P_{on}} \right)$$

The same approach works for analog modulation, but in this case you adjust the voltage so the delay varies continuously between 0° and 180° . The result is a continuous variation in output intensity shown in Figure 16.2. (Note that the peak power is lower than the laser output by an amount that equals the insertion loss of the modulator, even when there is no voltage applied.)

Actual electro-optic modulators are more complex. Often voltages are applied across *both* waveguides, but with the opposite polarities, so the voltage delays the phase of one wave while speeding the phase of the other. In this case, a voltage of $+V_{180}/2$ is applied to one waveguide, and $-V_{180}/2$ is applied to the other, giving the same modulation with lower voltage. Typically the voltage signal applied to each channel is the sum of two signals, one a bias that sets the operating level, the other the modulating signal. For example, the bias may set the modulator to normally transmit a certain average power, with the variations in the modulation voltage changing the transmitted power above and below that level.

A further complication is that refractive index can vary with the polarization of light. In glass and many other materials, the refractive index is nearly identical for light of different polarizations, but in other materials it varies significantly with the orientation of the polarization relative to the crystal axes. Materials in which the refractive index differs significantly for vertically and horizontally polarized light are called *birefringent*.

Modulators are affected because the magnitude of the electro-optic effect also depends on polarization, so an applied voltage does not change the refractive index the same amount for vertically and horizontally polarized light. Thus an electric field that delays vertically polarized light by 180° may delay horizontally polarized light only 120° , so interference would not cancel out the horizontally polarized component of the light. One way to avoid this problem is by using a polarizing filter to block the undesired polarization before the light reaches an electro-optic modulator, so only one polarization is transmitted.

In theory, electro-optic modulators can be made of any material that displays the electro-optic effect and is transparent at the signal wavelength. In practice, the usual material for use at 1.3 and 1.55 μm is lithium niobate (LiNbO_3). Waveguides are made by diffusing titanium or hydrogen into the lithium niobate, raising the refractive index of a narrow stripe that forms a waveguide. One process raises the refractive index for one polarization but depresses it for the other, so only one polarization stays in the guide, while the other diffuses into the substrate. The goal is to eliminate the need to polarize light before sending it through the modulator.

Lithium niobate modulators are widely used today. The technology is well-developed and they can be modulated at rates to 40 Gbit/s for digital transmission. As waveguide devices, they can be integrated with some other optical devices. However, they cannot be integrated with light sources because lithium niobate is not a light emitter.

Electro-optic modulators are made of lithium niobate.

Electro-Absorption Semiconductor Modulators

An electro-absorption semiconductor modulator is a waveguide device based on different principles than the electro-optic modulator. The electro-absorption modulator has a structure similar to that of an edge-emitting semiconductor laser, and the two can be integrated on the same chip, as shown in Figure 16.3. In this arrangement, the laser and modulator share an active layer, so light generated within the laser stripe is coupled directly to the modulator waveguide.

Despite their common structure, an electro-absorption modulator operates differently than a semiconductor laser. The laser is forward-biased so current flows through it, causing current carriers to recombine and generate light. The modulator is operated with a reverse bias, like a *pin* photodetector. When the modulator is unbiased, no current flows and it is transparent to the laser wavelength. However, when the bias voltage is applied, the laser light can produce electron-hole pairs that are pulled in opposite directions by the bias voltage,

An electro-absorption modulator is a semiconductor diode that is reverse-biased so modulation makes it absorb rather than emit light.




FIGURE 16.3

An electro-absorption modulator integrated with a semiconductor laser.

causing a net absorption at the laser wavelength. Increasing the bias increases the absorption, blocking the beam.

The laser and modulator are electrically isolated from each other. A steady current drives the laser, so it generates a steady optical output. The input signal drives the modulator. For zero applied voltage, the optical output is at its highest level. Applying a higher voltage to the modulator increases light absorption. (Note that this means that high voltages generate no light output.)

Although the laser and modulator sections of the integrated structure have similar structures, they are not identical. The structure of the active layer also differs between the two. The laser may include a distributed-feedback grating in the cavity or a distributed Bragg grating in the waveguide. Thicknesses of the active layers differ, as does the doping that differentiates between the laser and modulator sections. Nonetheless, the two devices can be fabricated on the same substrate, forming a single, integrated light source and modulator.

Like electro-optic modulators, electro-absorption modulators are polarization sensitive, although integrating them on the laser chip simplifies packaging. They are made from InGaAsP semiconductors, so they can readily match laser wavelengths.

Variable Filters and Dynamic Gain Equalization

Variable filters change too slowly to modulate a light signal.

Other types of modulators are available, but they change the intensity of transmitted light too slowly to be used as external modulators. Some that operate slowly are considered variable attenuators, which were covered in Chapter 15. (I warned you the definitions could be hazy.) Others can be used as switches, covered later in the chapter. Some are used as dynamic filters, which can be adjusted using feedback loops to control system performance. Dynamic filters are used in dynamic gain equalization systems, which monitor power on all optical channels in a WDM system so they can equalize power levels at all wavelengths. Power equalization uses feedback from the optical channel monitor to control dynamic filters that selectively attenuate certain channels. Let's look at a few examples.


In Chapter 15, we mentioned acousto-optic variable filters, in which acoustic waves create refractive-index variations in a crystal that deflect a fraction of the light passing through the crystal, reducing the transmitted intensity. The amount of reduction depends on the wavelength of the light and the frequency of the sound wave. Inducing multiple acoustic frequencies into the crystal can change attenuation over a range of wavelengths for use in dynamic gain equalization in a WDM system.

An alternative is thermo-optic modulation in planar waveguides. Heat changes the refractive index of the waveguide, so attaching heaters to one of a pair of waveguides can modulate the phase of light going through that guide. Combining light from the heated guide with light from the unheated guide modulates intensity the same way that electro-optic phase modulation does in an electro-optic modulator.

Liquid crystal devices also can be used as variable attenuators by varying their transmission, as in displays. LCDs respond much too slowly to modulate a laser beam with a signal, but they are adequate for variable attenuators, which don't require such high-speed response.

Another interesting approach is based on *micro-electro-mechanical systems* (MEMS). MEMS are arrays of micro-mechanical devices etched from silicon and integrated with electronic control circuits. For optical applications, MEMS are made with thin reflective

Diffraction MEMS change their reflectivity by shifting phase.

**FIGURE 16.4**

Diffractive MEMS device changes its reflectivity by shifting the phase of reflected light.

structures that move in response to electrical signals passing through circuits beneath them. Figure 16.4 shows a two-layer MEMS structure that can be used as a variable attenuator. One layer is an array of parallel reflective stripes suspended above the second layer, a refractive substrate. Moving the suspended upper layer a quarter wavelength relative to the substrate shifts the phase of the reflected light 180° relative to that of light from the substrate, shifting between constructive and destructive interference. Smaller variations in spacing adjust the intensity of the reflected light.

Switching in Optical Networks

Earlier you learned that the telecommunication network can be seen as an array of pipes and switches. The switches are the key difference between the old fiber-optic systems that merely piped signals from point to point, and the emerging optical network. *Optical switches* allow an optical network to process and direct light signals, as well as pipe them from place to place. This greatly enhances the functionality of an optical network, and makes optical switches very important.

Switches serve various functions in telecommunications, and you should understand a bit about these applications before looking at specific switches. The most familiar type of switching is directing signals from point to point. The telephone network and the Internet do this in different ways, but the end user does not see a big difference without looking closely. Users merely know that they dial a phone number or enter an Internet address.


Other switching functions are less obvious to users. Telecommunications companies install switches to route signals around failed equipment, so a single failure won't knock out an entire telecommunication system. This function is called *protection switching*, and it relies on having backup routes. One common approach is to connect several cities in a ring that has extra capacity reserved. A cable break triggers the protection switch to divert signals that would have gone through the broken fibers to pass through the reserve fibers, as shown in Figure 16.5.

Telecommunication companies are making increasing use of switches to change the services they provide to customers, a process called *provisioning*. Traditionally, carriers sent technicians to physically connect lines to customers, but now that the network is changing faster, it is more economical to install switches and make the changes remotely. A similar function is dynamically changing the transmission capacity of parts of the network to meet special needs. For example, a carrier might reconfigure the network around a sports stadium

Optical switches allow an optical network to process signals optically.

FIGURE 16.5

Protection switching. In case of a fiber break, switches at nodes 1 and 4 redirect traffic between those nodes over the spare fiber pair.



for a World Series or Superbowl. Optical switches are rarely used for *circuit switching*, making temporary connections such as those needed for a telephone call, because fibers handle much larger blocks of information.

Optical switching is still a young technology, and both hardware and applications are evolving. Most nonprotection switching is still electronic, but optical switching is coming, because it can manage higher-capacity transmission. Let's take a more careful look at the functional requirements for various types of switches.

Protection switching sends signals through a backup fiber.

Protection Switching

Protection switching is a simple but vital function. In case of a cable break or equipment failure, the network must redirect signals along a different path that will bypass the failure.

Ideally this should be done automatically in a small fraction of a second, so no telephone connections over the broken link are likely to be dropped. Some data will be lost, but protocols exist to retransmit the data.

Signals can be switched to a backup route at the transmitter output or at a node somewhere along the system. The equipment required is relatively simple. Something must detect the failure and command the switch to divert its signals along the backup fiber. Standard protection switches have two possible outputs, the main fiber and the backup. Their job is to sit and wait. Technicians may test them, but they're not used regularly and repeatedly. If there is a failure, the switch typically will be reset after repairs. This technology is common, and is part of an important telephone-industry standard called *SONET*, which you will learn about in Chapter 20.

Remote Provisioning and Reconfiguration

Provisioning is the changing of network configuration to alter the services delivered to customers, or to provide new services. Traditionally it has been done manually, but now there is much interest in *remote provisioning*. This is a robotic equivalent of sending a technician to a remote site to rearrange cables. The goal is to save money and help telecommunications companies manage their networks more efficiently. Remote provisioning is like going to your basement to set switches when you want to move a phone in your home. That analogy shows both the appeal and the difficulty—switching in new phone lines would be easier than stringing new wires, but you would have to install extra equipment to make it work.

Provisioning schedules are comparatively leisurely. If you want a new phone line, it doesn't have to be switched on in seconds. However, the operation usually is more complex than protection switching. Remote provisioning is still rare, but it's being designed into new equipment.

● Provisioning changes network configurations to deliver new services.


Cross-Connects and Circuit Switching

Directing signals among many possible users is a more complex task than protection switching. Figure 16.6 shows the basic idea. Signals must be directed from any of N possible inputs to any of M possible outputs. Switches that perform this task are called *cross-connects* or *switching fabrics*. They perform the same function as the old-fashioned telephone operator who sat at a switchboard plugging pairs of wires into sockets that led to different telephone lines. Monstrous banks of electro-magnetic switches did the same task a generation or two ago; now special-purpose electronic computers lie at the heart of telephone switching offices.

Optical cross-connects are just beginning to appear in the core of the telecommunication network, where they can transfer high-speed optical signals among input and output fibers. So far, most optical cross-connects can handle only a limited number of inputs and outputs, such as 8×8 switches, with 8 inputs and 8 separate outputs. Optical cross-connects have been demonstrated with up to 1000 inputs and outputs in the laboratory, and commercial versions are in development.

Although it is simplest to think of optical cross-connects as giant-scale versions of switchboards, that is not a very accurate view. Nobody today makes the equivalent of phone calls

● Cross-connects make connections among multiple inputs and outputs.

FIGURE 16.6*Optical cross-connect.*

that stream 10 Gbit/s between two end terminals. The only important uses of optical cross-connects are for load management, to deliver capacity where it is required.

WDM and Optical Switching

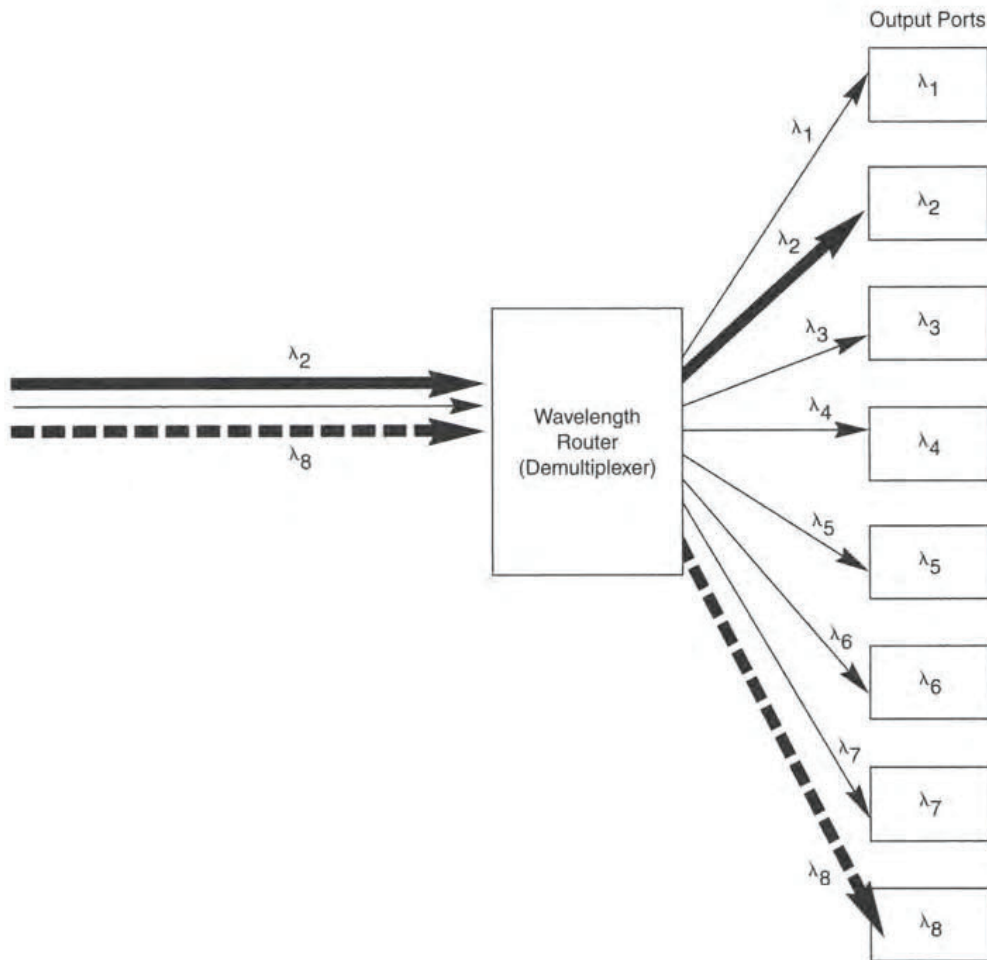
WDM channels may be switched together or separately.

So far, we have not considered how many optical channels are transmitted through each fiber. Wavelength-division multiplexing is an important issue in optical switching because different applications require different treatment of optical channels. Some applications require switching all optical channels carried by a fiber in the same way; others require that the optical channels be separated and switched independently.

All optical channels carried by the fiber need to be redirected for protection switching; the standard approach is to simultaneously switch them all to a backup fiber. Switching of all optical channels also may be needed when directing large volumes of traffic; for example, transmitting a large volume of traffic through a series of major switching nodes in a long-distance network. Network managers may organize transmission so one fiber carries signals from New York to Cleveland, which are then switched to another fiber for transmission from Cleveland to Chicago. A separate fiber from New York may carry signals that the Cleveland switch directs to Detroit. Simultaneously switching multiple channels in the same fiber simplifies switch operation.

On the other hand, traffic management often requires redirecting optical channels that arrive through the same fiber into several different fibers. For example, another fiber may carry signals from New York for distribution to other cities in Ohio; that is, one wavelength may go to Toledo, another to Akron, a third to Columbus, and a fourth to Cincinnati. In that case, a demultiplexer would separate the WDM signals from New York, then the switch in Cleveland would process them separately.

Redirecting individual wavelengths requires first separating the wavelengths. Depending on the configuration, this may require either isolating one wavelength with an add-drop multiplexer, or completely demultiplexing all the wavelengths.

**FIGURE 16.7**

Wavelength router directs input signals by their wavelength.

Wavelength Routers

One type of WDM switch deserves special mention—the *wavelength router*. Essentially a wavelength router is a special-purpose demultiplexer that directs optical channels to different destinations, depending on their destination. You can think of it as a conventional WDM demultiplexer with fixed output ports. It gets its name from the fact that it routes input signals to their destinations based on their wavelength, as shown in Figure 16.7. Any input signal at λ_2 is routed to port 2, while any input signal at λ_8 is routed to port 8. If you switch the wavelength, you switch the output port. Although present applications are limited by the need for wavelength converters, wavelength routers can provide a distinct function.

Note that wavelength routers are distinct from Internet routers, as described below.

Switches and Routers

The difference between switches and routers is an important one in telecommunications, but is easy for newcomers to misunderstand. Although both switches and routers direct signals, they do so in different ways and operate on different kinds of signals.

● A wavelength router separates signals by wavelength.

● Switches connect circuits. Routers direct data packets based on their headers.

●
Circuit-switched systems reserve dedicated channels.

So far this section has concentrated on switches and switching. Originally, switches made physical connections between electrical circuits, like a wall switch connects a light fixture to an electric power line, turning on the light. Old-fashioned electro-mechanical switches made physical connections between the wires running from your telephone and the wires running to your neighbor's phone.

Today, most switching is electronic, with solid-state circuits making connections. Once calls are digitized, your call does not have a whole wire (or fiber) to itself, but it does have a fixed time slot in the series of pulses being transmitted. Engineers still call this connection a *circuit* (or sometimes a *virtual circuit*), although it is not a set of wires dedicated to your conversation. Such circuit-switched systems reserve a guaranteed capacity for each call. It's functionally the same as having your own dedicated pair of wires during your entire conversation, always available whether or not you are talking. Your entire conversation follows the same route.

An alternative approach is called *packet switching*. Instead of holding a dedicated channel open for you all the time, you share the system with many other users. The signals you send are divided into data packets, with headers added to indicate their destination. Devices called *routers* read the headers, then decide where to send the packet based on that information and network conditions at the moment. You can think of them as drivers of parcel delivery trucks who read the label (the header) at your door, then decide the best route to take the package to its destination. The Internet is the most familiar example of packet switching.

Note that there are important functional differences between switches and routers. Switches set up a circuit and leave it alone as long as it's carrying signals. When the connection is finished, the switch hangs up and waits for another call. Switches don't pay any attention to the content of the call beyond the initial information needed to make the connection, and monitoring to see that the line is still in use.

Routers have a more complex job. They must read the headers of each and every packet, then direct it to one of many other routers partway to the packet's destination. The packet is likely to go through a series of routers. Each router in sequence reads the header and sends the packet closer to its destination. Like mail sorters, routers may bundle together packets that are going in the same direction, to be sorted and redistributed at their destination. In addition to reading the headers, routers monitor network conditions to establish the best routes for sending data packets.

It's important to remember that circuit switching and packet routing are different operations, with distinct requirements and hardware. Electronics can do both. So far, optical circuit switches are available, but true optical routers are in the research stage.

Transparent versus Opaque Switches

●
Transparent optical switches let light go straight through; opaque switches do not.

Optical switches can be divided into two broad categories: *transparent* and *opaque*. The names imply the key difference. Optical signals go straight through a transparent switch without being converted into any other form. One example is a mirror that moves back and forth, directing incident light into one of two possible outputs. The same optical signal that enters the switch is reflected from the mirror, and goes out one of the two possible outputs.

In an opaque switch, the signal is converted into some other form before switching. A simple example is a switch that converts the optical signals into electronic form, processes them electronically, then sends the output signals in one of two possible directions. This

sort of switch is considered opaque because the light signal does not go straight through it. Even though both input and output signals are in the form of light, the light is converted into some other form in between.

Free-space Optical Switching

Unlike electrical signals, optical signals can travel freely through the air or empty space. This means that optical switches can have internal gaps, unlike electrical switches that must have continuous physical connections so electrons can flow through them.

Free-space optical switching simply means sending signals between points through empty space instead of through optical fibers. For example, a mirror might be tilted to one of several positions, each one aiming a beam striking the mirror in a different direction.

Signals may pass through free space inside an optical switch.

Optical Switching Technologies

Several technologies can be used for optical switching, and more are in development. The essential idea is to move the beam from one point to another. This may be done mechanically by moving an optical component, or in other ways that shift or deflect light without moving parts. The technologies differ in how fast they can redirect a beam, and in how many different directions they can point it. Some can switch a beam between two directions, while others can aim it over a range of angles.

Opto-Mechanical Switches

Opto-mechanical switches redirect signals by moving fibers or optical components so they transfer light into different fibers. (They are considered distinct from the MEMS switches considered below.) Figure 16.8 shows a simple example. The input signal comes through the fiber on the left. A mechanical slider moves that fiber up and down, latching into one of three positions. Each position directs light from the input fiber into a different output fiber. In this design, the slider flexes a short length of the input fiber.

Many other designs are possible. Instead of moving a fiber, an opto-mechanical switch could move a mirror or lens to focus light into different fibers. The switch could be toggled mechanically or electronically. With precise optics, you can make an optical cross-connect

Opto-mechanical switches move fibers or optics to redirect signals.




FIGURE 16.8

An opto-mechanical switch.

that focuses light from one of several input fibers onto one of several output fibers. Collimating and collecting optics can focus the beam into the core of the output fiber.

The common element of all opto-mechanical switches is that their operation involves mechanical motion of an optical component. Precise motion is important because fiber alignment tolerances are tight, although large collecting optics can ease requirements. Although opto-mechanical switches are simple in concept, they are far more demanding in practice than electrical switches. Another disadvantage is that telecommunication companies generally prefer to avoid moving parts in our solid-state age.

Nonetheless, opto-mechanical switches have come into wide use because they are the simplest and cheapest optical switches available. They are used mainly for protection switching and other applications where it is vital to be able to switch signals when necessary, but where you hope it isn't necessary very often. They also are used in some instruments.


MEMS Switches


MEMS switches redirect light using tiny moving micromirrors.

Micro-Electro-Mechanical Systems (MEMS) are tiny mechanical structures made by depositing and etching a substrate material in a series of steps. We mentioned MEMS earlier, but they deserve more attention because they are potentially important for optical switching. MEMS structures can be coated with a reflective layer and moved back and forth to deflect light. Although their operation sounds electro-mechanical, that term normally is used only for larger-scale devices, and MEMS are considered a distinct technology.

MEMS technology is adapted from the photolithographic methods of making integrated electronic circuits. Doping and deposition build up a series of patterned layers on a semiconductor substrate—in practice, on silicon. Then some of the material is etched away to leave mirrors supported by posts, as shown in Figure 16.9. Circuits deposited on layers below the suspended mirrors can carry currents, which generate electromagnetic forces that can pull on the mirrors, tilting them. The tilting mirrors scan reflected light across space and can direct beams to output ports. They require about 10 volts to activate, can switch position in microseconds, and can operate for hundreds of millions of cycles.

FIGURE 16.9
MEMS mirrors tilt back and forth.



**FIGURE 16.10**

Two-axis tilting mirror. The center mirror pivots on two axes defined by the two surrounding rings. (Courtesy of Lucent Technologies)

Originally developed for use in displays, optical MEMS have been adapted for switching. Arrays of mirrors are fabricated on silicon substrates. With careful design, complex mirrors can be made to tilt back and forth in two dimensions, so they could scan both vertically and horizontally. Figure 16.10 shows an example of such a mirror, encircled by a pair of rings that can tilt it in two dimensions.

Such tilting mirror structures can scan over a range of angles, so they can collect light from many distinct input ports and direct it to any of many distinct outputs. This gives tremendous flexibility, but it makes accuracy essential. If the mirrors drift from their assigned positions, the output can go to the wrong port.

An alternative design moves mirrors between two distinct positions, where they latch in place. You can think of them as being in either the “up” or “down” position. If they are down, light goes through their position unchanged. If the mirror is up, it reflects light in a single alternative direction. Because the mirror latches into place, it always reflects light in the same direction, easing the need for alignment.

Latching structures are somewhat more complex than tilting mirrors, but moving between fixed positions is attractive for some applications. Advocates sometimes call these latching mirrors “digital” MEMS because they have only two positions, the equivalent of “off” and “on.” This is a careful choice of words, because it implies tilting mirrors are “analog” and thus imprecise and obsolete.

MEMS technology enjoyed a boom during the telecommunications bubble, but suffered badly in the aftermath. Many problems came from overpromotion of business prospects, but technical issues also remain to be resolved, including the long-term performance of MEMS optical switches.


Bubble Switches

Bubbles that move back and forth in liquid guides also can serve as an optical switch. In these devices, two grids of planar optical waveguides cross each other on a substrate, as shown in Figure 16.11. These waveguides have a higher refractive index than the substrate, so they guide light across the flat device. Channels containing a liquid with the same refractive index as the waveguides cross through the junction points of the waveguides.

As long as liquid fills the channel at the junction point, the light sees a continuous waveguide and goes straight through the junction. This changes when a bubble moves into the junction point. The refractive index in the bubble is much lower than that in the liquid,

Some MEMS switches latch into two distinct positions.

Bubbles moving back and forth in liquid guides can switch light by total internal reflection.

FIGURE 16.11*Bubble switch.*

and the waveguides cross at a sharp angle, which is beyond the critical angle for total internal reflection. When the bubble is in place, it causes total internal reflection, diverting the light down the other waveguide, as shown in Figure 16.11.


The same techniques used to control ink-jet printers can move bubbles back and forth in the channels, which are sealed to keep the liquid from escaping. The bubbles can be formed by vaporizing small amounts of the liquid, with expansion regions left in the liquid channels to allow for changes in volume. Bubble switches have no mechanical moving parts, although the liquid does move. The technology is considered promising for optical cross-connects but is still young.

Electro-Optical Switches

The electro-optic waveguide technology used for the electro-optic modulators described earlier in this chapter also can be used to make a solid-state optical switch with no solid or liquid moving parts. To make a switch, the single input waveguide is replaced by a pair of input guides that meet in a 2×2 coupler connected to the active section. Then the single output of the two parallel guides in the active section is replaced by a 2×2 coupler splitting the signal between a pair of output guides. Figure 16.12 shows the idea.

As in the modulator, operation depends on applying a voltage to one or both of the parallel electro-optic guides in the active section. This changes the refractive index, delaying the phase of light in one waveguide relative to the other. Changing the relative phase of the output signal from the two guides by 180° switches it from one output port to the other. This can switch a single input from one output channel to the other. If separate input signals are entering the top and bottom ports, a 180° phase shift can swap the two between different outputs, as shown in Figure 16.12. If you merely want to switch a signal off and on, you can switch it off by directing the light to a port that goes nowhere. The result is a solid-state switch with no moving parts and a very quick response time. The only changes are in the drive voltage. Normally electro-optical switches are made of lithium niobate, because applying a voltage across the waveguide causes a large change in refractive index.

● Voltages applied to planar waveguide channels switch signals in electro-optical switches.

**FIGURE 16.12**

A 2 × 2 electro-optic switch.

Switching Signal	Connections
V_1	A-C, B-D
V_2	A-D, B-C

This approach works well for switching one or two ports, but more complex switching configurations require cascading a series of waveguide switches. A 1×4 waveguide switch requires a cascade of three 1×2 switches, with the first switch providing the inputs for the other two, which have a total of four outputs. An 8×8 switch would require a total of 64 switching elements. Although this sounds cumbersome, it has been effective in meeting actual requirements for optical switching, which proved far more modest than market analysts had expected. Waveguide electro-optic switches enjoy two major advantages—a well established technology and all solid-state operation, with no moving parts.

Thermo-Optic Switches

Thermo-optic switches work by using interference effects produced by changing the refractive indexes of a pair of parallel arms coupled at both ends. They can be made either as waveguides or as discrete devices. In thermo-optic devices, heating changes the refractive index of one arm, shifting the relative phase of the light to switch the output between two ports. Thermo-optic switches use different materials than electro-optic switches, but the optical principles are the same.

Thermo-optic switches are slower than electro-optic switches, with response time of milliseconds. However, this is adequate for most types of optical switching, and they are widely used. They are also solid-state devices, and are less expensive than electro-optic switches.

Liquid-Crystal Switches

Another switching technology, long used in optical displays, is liquid crystals. Liquid crystals get their name because their large molecules tend to orient themselves in the liquid phase, although they do not form a fixed lattice like a solid crystal. This molecular alignment can polarize transmitted light. The types of liquid crystals used in displays have another important property—applying an electric field can change their orientation, and thus change their effect on the polarization of transmitted light.

Thermo-optic switches change refractive index by heating.

Liquid-crystal switches work by changing light polarization.

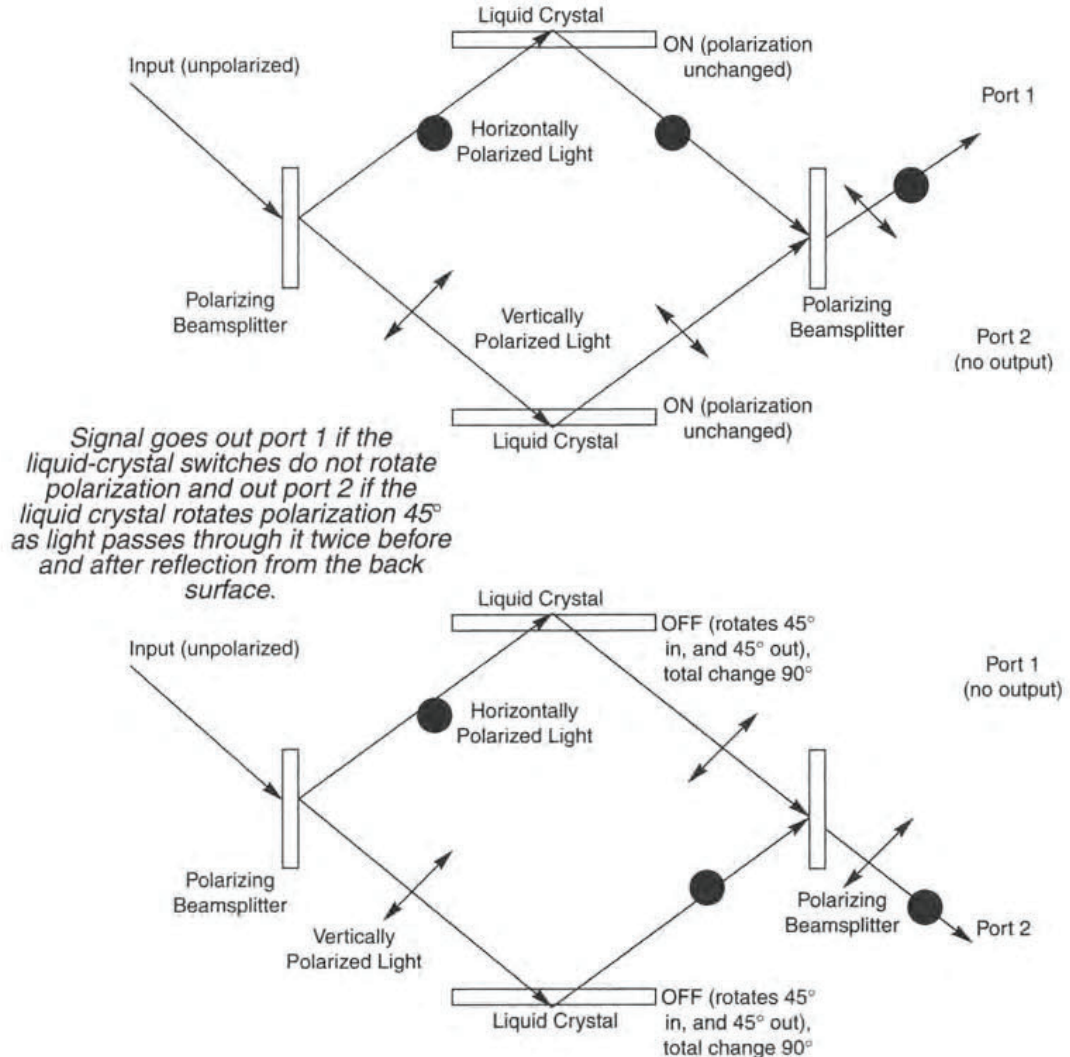
For displays or switches, liquid crystals are sandwiched in a thin layer between two parallel glass plates with electrodes applying a voltage across the liquid. The voltage switches between two states, typically one that rotates the polarization, and one that leaves the polarization unchanged.

For example, suppose a vertical polarizer is put on the top of a liquid crystal device, so light passes through the polarizer before entering the liquid crystal. Applying a voltage then rotates polarization 45° as the light passes through the liquid crystal layer when it is reflected by the rear surface, and rotated another 45° as it passes back through the liquid crystal. Thus the light exiting the liquid crystal has rotated 90° to be horizontally polarized, and is blocked by the polarizer. Those regions would look dark on a liquid crystal display.

The switch shown in Figure 16.13 works in a similar way. Input light is separated into its two polarizations and reflected off a liquid crystal plate. In this case, the liquid crys-

FIGURE 16.13

Liquid-crystal switch.



tal rotates polarization 45° when the voltage is off, but does not affect polarization when it is on. As a result, the reflected light has different polarizations when the voltage is off or on. Polarizing optics deflect that light in different directions, depending on its polarization, so the light emerges from different ports depending on the voltage applied to the liquid crystal.

Wavelength Switching and Conversion

So far we have assumed that all the switching is between fibers, with any WDM signals all being switched together. Although that's usually the case for most optical switching (other than at the ends of a system), the switching process can get considerably more complex in some situations.

One case is a variation on the optical add-drop multiplexer described in Chapter 15. Sometimes called a *reconfigurable optical add-drop multiplexer*, this device is actually a switch that selects one or more wavelengths to drop at a particular location. Actually two steps are involved—selecting the wavelength, then switching it. In practice, this can be done by demultiplexing the signal and routing the wavelengths separately to switches that select which optical channels will be dropped, as shown in Figure 16.14. Other approaches in development would switch a single wavelength without having to demultiplex the entire signal.

Switching of an individual optical channel also can create the need for converting the signal on that optical channel to another wavelength for transmission through a different WDM fiber system. You can think of this conversion process as shifting the signal between different lanes on the optical information highway.




FIGURE 16.14
Switching at a reconfigurable add-drop multiplexer.

You learned about wavelength conversion in Chapter 12. There are two main approaches. One is opto-electro-optical (OEO) conversion, in which the optical signal is converted to electronic form to drive a laser transmitter emitting at the desired wavelength. That's a brute-force approach, but it's easy to implement, and it's done quietly today within electronic switches at the end points of transmission lines. OEO conversion also can be made quite flexible by using tunable lasers in the transmitter, so the output wavelength can be selected. All-optical wavelength conversion is possible, but the technology is still in development and is not widely used.

Both wavelength switching and wavelength conversion are expected to become important technologies as the telecommunications network evolves toward a future all-optical network. We aren't there yet, however, and today those functions are largely handled by converting optical signals into electronic format.

Integrated Optics

●
Integrated optics
combine functions
in a single
monolithic device.

In Chapter 6 you learned about passive planar waveguides that guide light in ways similar to optical fibers but are rectangular in cross-section. Planar waveguides can be integrated on the same monolithic substrate with other flat components—such as semiconductor lasers—to make *integrated optics* that can perform multiple functions. The concept has been around since the late 1960s, when it was proposed as an optical counterpart to integrated electronic circuits. However, optical integration has proved considerably more difficult than electronic integration, and so far has found only limited applications.

Some planar waveguide components are already well developed, including the electro-optic modulators and switches described earlier in this chapter. Edge-emitting semiconductor lasers and semiconductor optical amplifiers also are planar waveguide devices, because light is guided through a narrow stripe in their thin active layer. Yet the degree of integration is modest compared to electronics. One example is the combination of an electro-absorption modulator with a semiconductor laser on the same substrate. Another is the arrayed waveguide grating described in Chapter 15.

●
Integrated optics
is widespread in
the laboratory.

Integrated optics activity is widespread in the laboratory, pushed by efforts both to achieve high performance and to reduce costs. Arrayed waveguide gratings have been integrated with other components so that WDM signals can be demultiplexed, individual optical channels modified, and the signals then multiplexed for further transmission. Many advanced functions such as wavelength conversion have been demonstrated with integrated optics. Mass production of simple integrated optics is promising for reducing system and component costs.

A separate class of devices, sometimes called *integrated optoelectronics*, is being developed to reduce costs, particularly of transmitters and receivers. As the name suggests, they integrate electronic as well as optical functions. Examples include combining detectors and amplification circuits, or drive circuits with light sources. Some of these devices are not true monolithic circuits made on a single substrate, but hybrid circuits that bond together two chips made of different semiconductors, one for the optics and the other for the electronics.

What Have You Learned?

1. Active components change signals. They typically draw power from an external source. Modulators and switches are important examples.
2. In high-performance transmitters, an external modulator produces the output signal by modulating the intensity of a continuous beam from the laser source.
3. Electro-optic modulators rely on electric fields to change the refractive index of lithium niobate in a planar waveguide. This delays the phase of light in one of two parallel guides, causing interference, which modulates light intensity.
4. Electro-absorption semiconductor modulators are reverse-biased diodes. Applying a drive current causes them to absorb light at the laser wavelength. Laser and modulator can be integrated on the same chip.
5. Optical networks require optical switches to direct and process signals.
6. Protection switching sends signals through a backup fiber in case of failure. Provisioning changes the services delivered over telecommunication lines.
7. Cross-connects can connect any of multiple inputs to any of multiple outputs. A telephone switch is a good example.
8. The WDM channels carried in a single fiber may be switched separately or collectively.
9. A wavelength router directs input signals according to their wavelength. The same wavelength always goes to the same destination.
10. Switches connect circuits or reserved channels. Routers direct data packets based on their headers. The standard telephone network is circuit-switched. Internet data is transmitted by packet switching and directed by routers.
11. Opto-mechanical switches move fibers or bulk optics to redirect signals.
12. MEMS are micro-electro-mechanical systems made by etching tiny mechanical structures from a semiconductor. They include tiny micromirrors, which can be moved to switch optical signals. MEMS mirrors may tilt over a continuous range, or latch into distinct positions.
13. Bubble switches direct signals by moving bubbles back and forth in liquid guides to the junction points of planar waveguides. The bubbles redirect signals by total internal reflection.
14. Electro-optical and thermo-optical switches change the refractive index of planar waveguides, causing interference, which switches the signal between a pair of output ports.
15. Liquid crystal switches affect the polarization of light; they are used together with polarizers to switch light.
16. Variable or dynamic filters change attenuation too slowly to modulate an optical signal.

17. A reconfigurable add-drop multiplexer switches individual optical channels in a WDM system. Wavelength conversion may also be needed in WDM systems.
18. Integrated optics combine multiple functions. Only modest integration is used in commercial components, but more extensive integration is in development.

What's Next?

Chapter 17 describes optical measurement techniques.

Further Reading

- J. Capmany et al., eds., special issue on “Arrayed Grating Routers/WDM Mux Demuxs and Related Applications/Uses,” *IEEE Journal Selected Topics in Quantum Electronics* 8, (November/December 2002)
- K. Okamoto, *Fundamentals of Optical Waveguides* (Academic Press, 2000).
- Rajiv Ramaswami and Kumar N. Sivarajan, *Optical Networks: A Practical Perspective* (Morgan Kaufmann, 2002)

Questions to Think About

1. An external modulator with a 20-dB extinction ratio modulates the output of a 1-mW laser. The signal then passes through 20 km of fiber with loss of 0.5 dB/km. Neglecting other losses, what are the power levels at the detector when the light is off and when it is on?
2. What are the power levels in Question 1 if the external modulator has insertion loss of 3 dB?
3. What are the power levels for Question 1 if the external modulator has 3-dB insertion loss and an extinction ratio of 10 dB?
4. One important issue in switch design is the number of elements required to switch the signals. Suppose you have a simple cross-connect switch such as the one shown in Figure 16.5, with a single switch element at each node, which either transmits or reflects the beam. If you have 10 inputs and 10 outputs, how many switching elements do you need? What if you have 100 inputs and 100 outputs?
5. A tilting-mirror switch can reflect the light input from a single input port to any of N output ports. With this design, how many switching elements do you need for a 10×10 switch? A 100×100 switch? Assume the tilting mirror can point the beam at as many output ports as needed.
6. How does the bubble switch shown in Figure 16.11 scale? What number of bubble-waveguide intersections do you need for an $N \times N$ optical cross-connect?

Chapter Quiz

1. What phase shift do you need to cause destructive interference between two coherent beams of light, canceling them out?
 - a. 0°
 - b. 45°
 - c. 90°
 - d. 180°
 - e. 360°
2. What material is used in electro-optic modulators and switches?
 - a. lithium niobate
 - b. gallium arsenide
 - c. indium phosphide
 - d. silica on silicon
 - e. any of the above
3. How should an electro-absorption modulator be biased to block light transmission?
 - a. No bias is required; it is normally opaque.
 - b. reverse bias
 - c. forward bias
 - d. It must be biased in the same direction as the integrated laser light source.
4. Telecommunications customers use provisioning switches for
 - a. backup during repairs of a failed transmission line.
 - b. making temporary circuit connections to direct telephone calls.
 - c. changing services provided to customers.
 - d. routing data packets over the Internet.
5. Operation of an optical cross-connect is analogous to
 - a. fuses that block electrical power transmission if a circuit overloads.
 - b. a telephone switchboard that makes connections between callers.
 - c. a fleet of trucks delivering parcels over the best available routes.
 - d. municipal water services that pipe water to all homes and businesses.
6. A wavelength router
 - a. directs incoming signals to outputs according to their wavelengths.
 - b. is an Internet router able to process WDM signals at multiple wavelengths.
 - c. is an Internet router that can process optical signals at only one wavelength.
 - d. is an optical cross-connect that converts optical signals to different wavelengths.

7. The difference between switches and routers is
 - a. switches are mechanical and routers are electronic.
 - b. switches are optical and routers are mechanical.
 - c. switches connect circuits and routers direct packets.
 - d. switches direct packets and routers reserve channels.
 - e. just a difference in marketing buzzwords.
8. What kind of switch converts an optical signal to electronic form, then uses the electronic signal to drive another optical transmitter?
 - a. transparent
 - b. opaque
 - c. opto-mechanical
 - d. electro-optical
 - e. bubble
9. An opto-mechanical switch
 - a. uses light to mechanically move an electrical switch.
 - b. mechanically moves a fiber, mirror, or lens to redirect optical signals.
 - c. mechanically moves an electronic switch to redirect optical signals.
 - d. uses light to mechanically move an optical switch.
 - e. none of the above
10. What type of optical switch can direct light across a range of angles to any of the many possible output ports?
 - a. an electro-optical switch
 - b. a pop-up MEMS switch that latches in one of two positions
 - c. a tilting-mirror MEMS switch
 - d. a bubble switch
 - e. a liquid-crystal switch
11. How could you assemble a 1×8 electro-optical switch?
 - a. by dividing one input waveguide into eight optical waveguides
 - b. by moving one input fiber to connect with one of eight output fibers
 - c. by tilting a mirror to one of eight possible positions directing light to different outputs
 - d. by cascading a series of seven 2×2 switches, with the two outputs of the first going to inputs of two separate second-stage switches, and the four outputs of those switches going to inputs of four final-stage switches
 - e. It's impossible.

- 12.** What is the key difference between a MEMS switch and an opto-mechanical type?
- The MEMS switch has no moving parts.
 - Opto-mechanical switches have only two possible outputs.
 - MEMS switches are miniature monolithic devices with movable elements.
 - None of the above
- 13.** Which type of switch operates by changing the polarization of light?
- MEMS
 - opto-mechanical
 - bubble
 - liquid-crystal
 - electro-optical
- 14.** Which of the following is not an application of optical switches?
- protection switching around a damaged fiber
 - changing transmission lines to serve a new customer
 - redirecting signals at the terminal point of a cable
 - balancing transmission load among several possible routes for a telecommunications carrier
 - routing Internet packets

Fiber-Optic Measurements

About This Chapter

Fiber-optic technology has its own distinct set of measurements, based largely on a mixture of optical and electronic techniques. This chapter covers the important optical aspects of fiber-optic measurements, assuming you know something about electronics. It covers measurement units, the quantities measured, and the types of measurements performed, pointing out differences between optical and electronic measurements.

The emphasis here is on basic optical concepts that you need to know when working with fiber optics. The next chapter covers fiber-optic test equipment, along with its use in troubleshooting fiber-optic systems.

Basics of Optical Power Measurement

Most important fiber-optic measurements involve light, in the same way that important electronic measurements involve electric fields and currents. There are some exceptions, such as the length and diameter of optical fibers and cables, the sizes of other components, and the electrical characteristics of transmitter and receiver components. Because this is a book about fiber optics, I will mention such measurements only in passing. However, I will talk about measuring some things other than light, because you cannot quantify the properties of optical fibers if you consider only light. For example, to measure the dispersion of light pulses traveling through an optical fiber, you must observe how light intensity varies as a function of time, which requires measuring time as well as light.

When you're working with light, you need to know what can be measured. The most obvious quantity is optical power, which like electrical voltage is a fundamental measuring stick. However, power alone is rarely enough; it usually must be measured as a function

Fiber-optic measurements involve light and other quantities, such as the variation of light with time.

Specialized terminology makes fine distinctions about quantities related to optical power.

of other things, such as time, position, and wavelength. Wavelength itself is important because optical properties of optical components, materials, light sources, and detectors all depend on wavelength. Other quantities that are sometimes important are the phase and polarization of the light wave. You need to learn a little more about these concepts before getting into more detail on measurement types and procedures.

Optical Power and Energy

People have an intuitive feeling for the idea of optical power (measured in watts) as the intensity of light. However, a closer look shows that optical power and light intensity are rather complex quantities and that you need to be careful what you talk about. Table 17.1 lists the most important quantities, which are described in more detail later.

Each photon or quantum of light carries a characteristic *energy*, as you learned in Chapter 2. Energy is often denoted by E , but the symbol Q is often used in optics. The amount is a function of the wavelength or frequency of the electromagnetic wave. *Photon energy* is easiest to express as the frequency of the wave (ν) times Planck's constant h , which equals 6.63×10^{-34} joule-second, or 4.14×10^{-15} electron-volt-second.

$$\text{Energy} = h\nu$$

When working in wavelength units, the formula for photon energy (in joules) is

$$\text{Energy (J)} = \frac{hc}{\lambda}$$

where c is the speed of light (approximately 300,000 km/s) and λ the wavelength in meters. If the wavelength is expressed in micrometers, the formula for photon energy (in electron volts) becomes

$$\text{Energy (eV)} = \frac{1.2406}{\lambda(\text{in } \mu\text{m})}$$

Table 17.1 Measurable quantities related to optical power

Quantity and Symbol	Meaning	Units
Energy (Q or E)	Amount of light energy	joules
Optical power (P or ϕ) (or radiant flux)	Flow of light energy past a point at a particular time (dQ/dt)	watts
Intensity (I)	Power per unit solid angle	watts per steradian
Irradiance (E)	Power incident per unit area	W/cm ²
Radiance (L)	Power per unit solid angle per unit projected area	W/steradian-m ²
Average power	Power averaged over time	watts
Peak power	Peak power in a pulse	watts

Photon energy is the energy carried by a single photon. Normally a pulse of light contains many photons, and the total pulse energy carried by the pulse is the sum of all the photon energies. You can think of energy as the sum of the energies of all the photons that arrive at a destination. However, only knowing the total energy does not tell you if the energy arrived in a single pulse lasting a tiny fraction of a second, or in a slow but steady trickle that took all day.

Power (P or ϕ) measures the rate of energy transfer per unit time. You can think of it as the rate at which photons (or, equivalently, electromagnetic waves) carry energy through the system or arrive at their destination. The rate of energy transfer can vary with time, so power is a function of time. Mathematically, power P is expressed as a rate of change or derivative with respect to energy (Q):

$$\text{Power} = \frac{d(\text{energy})}{d(\text{time})}$$

or


$$P = \frac{dQ}{dt}$$

Sometimes called *radiant flux*, optical power is measured in watts. One *watt* is defined as the flow of one joule of energy per second. Watts are units used to measure the transfer rates of all types of energy, including electrical energy or heat. That is, a watt of light delivers the same amount of energy per second as a watt of electricity. (Note, however, that the power ratings of light bulbs in watts measure how much electrical power they use, not the amount of visible light they radiate, which is much lower.)

Optical and Electrical Power

Both optical and electrical power measure more fundamental quantities. For light and other types of electromagnetic radiation, the power is proportional to the square of the amplitude of the electromagnetic wave, shown as A in Figure 17.1, as well as to the number of photons received per second. The wave amplitude measures the strength of the electrical field in the wave.

Optical power is proportional to the square of the light wave amplitude.



Frequency (ν) =
Number of Waves per Second

FIGURE 17.1

Properties of an electromagnetic wave.