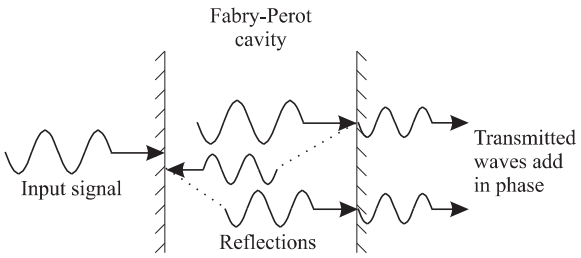

Figure 3.16 Principle of operation of a Fabry-Perot filter.

constants for the cladding modes are discussed in [Ven96b]. The amount of wavelength-dependent loss can be controlled during fabrication by controlling the UV exposure time. Complicated transmission spectra can be obtained by cascading multiple gratings with different center wavelengths and different exposures. The example shown in Figure 3.15 was obtained by cascading two such gratings [Ven96a]. These gratings are typically a few centimeters long.

3.3.5 Fabry-Perot Filters

A Fabry-Perot filter consists of the cavity formed by two highly reflective mirrors placed parallel to each other, as shown in Figure 3.16. This filter is also called a Fabry-Perot interferometer or etalon. The input light beam to the filter enters the first mirror at right angles to its surface. The output of the filter is the light beam leaving the second mirror.

This is a classical device that has been used widely in interferometric applications. Fabry-Perot filters have been used for WDM applications in several optical network testbeds. There are better filters today, such as the thin-film resonant multicavity filter that we will study in Section 3.3.6. These latter filters can be viewed as Fabry-Perot filters with wavelength-dependent mirror reflectivities. Thus the fundamental principle of operation of these filters is the same as that of the Fabry-Perot filter. The Fabry-Perot cavity is also used in lasers (see Section 3.5.1).

Compact Fabry-Perot filters are commercially available components. Their main advantage over some of the other devices is that they can be tuned to select different channels in a WDM system, as discussed later.

Principle of Operation

The principle of operation of the device is illustrated in Figure 3.16. The input signal is incident on the left surface of the cavity. After one pass through the cavity, as

shown in Figure 3.16, a part of the light leaves the cavity through the right facet and a part is reflected. A part of the reflected wave is again reflected by the left facet to the right facet. For those wavelengths for which the cavity length is an integral multiple of half the wavelength in the cavity—so that a round trip through the cavity is an integral multiple of the wavelength—all the light waves transmitted through the right facet *add in phase*. Such wavelengths are called the *resonant wavelengths* of the cavity. The determination of the resonant wavelengths of the cavity is discussed in Problem 3.7.

The power transfer function of a filter is the fraction of input light power that is transmitted by the filter as a function of optical frequency f, or wavelength. For the Fabry-Perot filter, this function is given by

$$T_{FP}(f) = \frac{\left(1 - \frac{A}{1-R}\right)^2}{\left(1 + \left(\frac{2\sqrt{R}}{1-R}\sin(2\pi f\tau)\right)^2\right)}. \tag{3.12}$$

This can also be expressed in terms of the optical free-space wavelength λ as

$$T_{FP}(\lambda) = \frac{\left(1 - \frac{A}{1-R}\right)^2}{\left(1 + \left(\frac{2\sqrt{R}}{1-R}\sin(2\pi nl/\lambda)\right)^2\right)}.$$

(By a slight abuse of notation, we use the same symbol for the power transfer function in both cases.) Here A denotes the absorption loss of each mirror, which is the fraction of incident light that is absorbed by the mirror. The quantity R denotes the *reflectivity* of each mirror (assumed to be identical), which is the fraction of incident light that is reflected by the mirror. The one-way propagation delay across the cavity is denoted by τ. The refractive index of the cavity is denoted by n and its length by l. Thus $\tau = nl/c$, where c is the velocity of light in vacuum. This transfer function can be derived by considering the sum of the waves transmitted by the filter after an odd number of passes through the cavity. This is left as an exercise (Problem 3.8).

The power transfer function of the Fabry-Perot filter is plotted in Figure 3.17 for $A = 0$ and $R = 0.75$, 0.9, and 0.99. Note that very high mirror reflectivities are required to obtain good isolation of adjacent channels.

The power transfer function $T_{FP}(f)$ is periodic in f, and the peaks, or *passbands,* of the transfer function occur at frequencies f that satisfy $f\tau = k/2$ for some positive integer k. Thus in a WDM system, even if the wavelengths are spaced sufficiently far apart compared to the width of each passband of the filter transfer function, several frequencies (or wavelengths) may be transmitted by the filter if

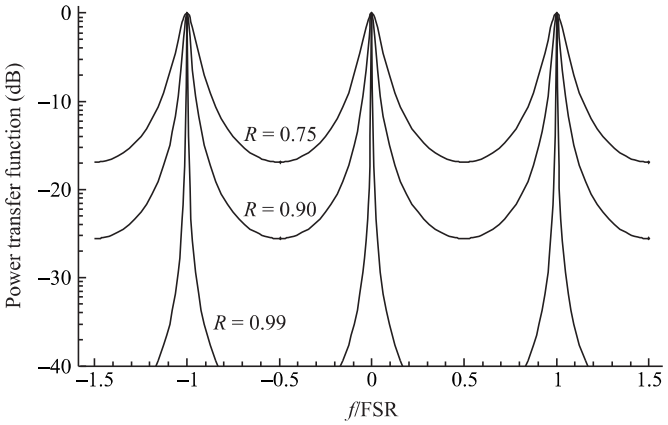

Figure 3.17 The transfer function of a Fabry-Perot filter. FSR denotes the free spectral range, f the frequency, and R the reflectivity.

they coincide with different passbands. The spectral range between two successive passbands of the filter is called the *free spectral range* (FSR). A measure of the width of each passband is its *full width* at the point where the transfer function is *half* of its *maximum* (FWHM). In WDM systems, the separation between two adjacent wavelengths must be at least a FWHM in order to minimize crosstalk. (More precisely, as the transfer function is periodic, adjacent wavelengths must be separated by a FWHM plus an integral multiple of the FSR.) Thus the ratio FSR/FWHM is an approximate (order-of-magnitude) measure of the number of wavelengths that can be accommodated by the system. This ratio is called the *finesse*, F, of the filter and is given by

$$F = \frac{\pi\sqrt{R}}{1 - R}. \tag{3.13}$$

This expression can be derived from (3.12) and is left as an exercise (Problem 3.9).

If the mirrors are highly reflective, won't virtually all the input light get reflected? Also, how does light get out of the cavity if the mirrors are highly reflective? To resolve this paradox, we must look at the light energy over all the frequencies. When we do this, we will see that only a small fraction of the input light is transmitted through the cavity because of the high reflectivities of the input and output facets, but at the right frequency, all the power is transmitted. This aspect is explored further in Problem 3.10.

Tunability

A Fabry-Perot filter can be tuned to select different wavelengths in one of several ways. The simplest approach is to change the cavity length. The same effect can be achieved by varying the refractive index within the cavity. Consider a WDM system, all of whose wavelengths lie within one FSR of the Fabry-Perot filter. The frequency f_0 that is selected by the filter satisfies $f_0\tau = k/2$ for some positive integer k. Thus f_0 can be changed by changing τ, which is the one-way propagation time for the light beam across the cavity. If we denote the length of the cavity by l and its refractive index by n, $\tau = ln/c$, where c is the speed of light in vacuum. Thus τ can be changed by changing either l or n.

Mechanical tuning of the filter can be effected by moving one of the mirrors so that the cavity length changes. This permits tunability only in times of the order of a few milliseconds. For a mechanically tuned Fabry-Perot filter, a precise mechanism is needed in order to keep the mirrors parallel to each other in spite of their relative movement. The reliability of mechanical tuning mechanisms is also relatively poor.

Another approach to tuning is to use a piezoelectric material within the cavity. A piezoelectric filter undergoes compression on the application of a voltage. Thus the length of the cavity filled with such a material can be changed by the application of a voltage, thereby effecting a change in the resonant frequency of the cavity. The piezo material, however, introduces undesirable effects such as thermal instability and hysteresis, making such a filter difficult to use in practical systems.

3.3.6 Multilayer Dielectric Thin-Film Filters

A thin-film resonant cavity filter (TFF) is a Fabry-Perot interferometer, or etalon (see Section 3.3.5), where the mirrors surrounding the cavity are realized by using multiple reflective dielectric thin-film layers (see Problem 3.13). This device acts as a bandpass filter, passing through a particular wavelength and reflecting all the other wavelengths. The wavelength that is passed through is determined by the cavity length.

A thin-film resonant multicavity filter (TFMF) consists of two or more cavities separated by reflective dielectric thin-film layers, as shown in Figure 3.18. The effect of having multiple cavities on the response of the filter is illustrated in Figure 3.19. As more cavities are added, the top of the passband becomes flatter and the skirts become sharper, both very desirable filter features.

In order to obtain a multiplexer or a demultiplexer, a number of these filters can be cascaded, as shown in Figure 3.20. Each filter passes a different wavelength and reflects all the others. When used as a demultiplexer, the first filter in the cascade

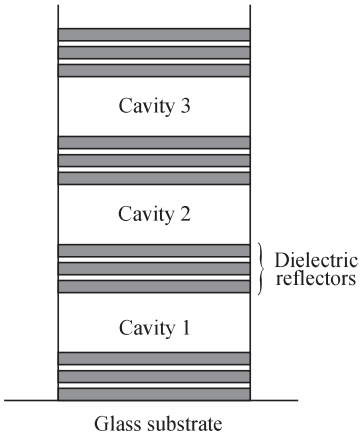

Figure 3.18 A three-cavity thin-film resonant dielectric thin-film filter. (After [SS96].)

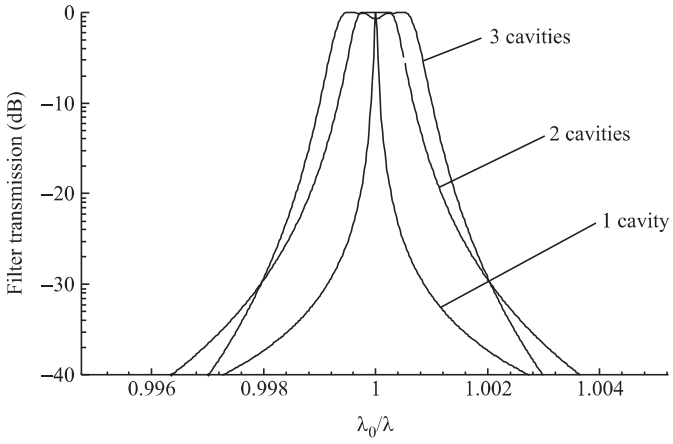

Figure 3.19 Transfer functions of single-cavity, two-cavity, and three-cavity dielectric thin-film filters. Note how the use of multiple cavities leads to a flatter passband and a sharper transition from the passband to the stop band.

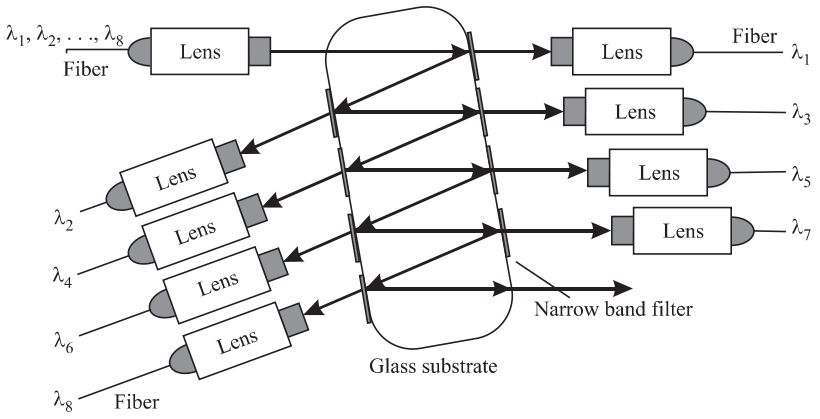

Figure 3.20 A wavelength multiplexer/demultiplexer using multilayer dielectric thin-film filters. (After [SS96].)

passes one wavelength and reflects all the others onto the second filter. The second filter passes another wavelength and reflects the remaining ones, and so on.

This device has many features that make it attractive for system applications. It is possible to have a very flat top on the passband and very sharp skirts. The device is extremely stable with regard to temperature variations, has low loss, and is insensitive to the polarization of the signal. Typical parameters for a 16-channel multiplexer are shown in Table 3.1. For these reasons, TFMFs are becoming widely used in commercial systems today. Understanding the principle of operation of these devices requires some knowledge of electromagnetic theory, and so we defer this to Appendix G.

3.3.7 Mach-Zehnder Interferometers

A Mach-Zehnder interferometer (MZI) is an interferometric device that makes use of two interfering paths of different lengths to resolve different wavelengths. Devices constructed on this principle have been around for some decades. Today, Mach-Zehnder interferometers are typically constructed in integrated optics and consist of two 3 dB directional couplers interconnected through two paths of differing lengths, as shown in Figure 3.21(a). The substrate is usually silicon, and the waveguide and cladding regions are silica (SiO_2).

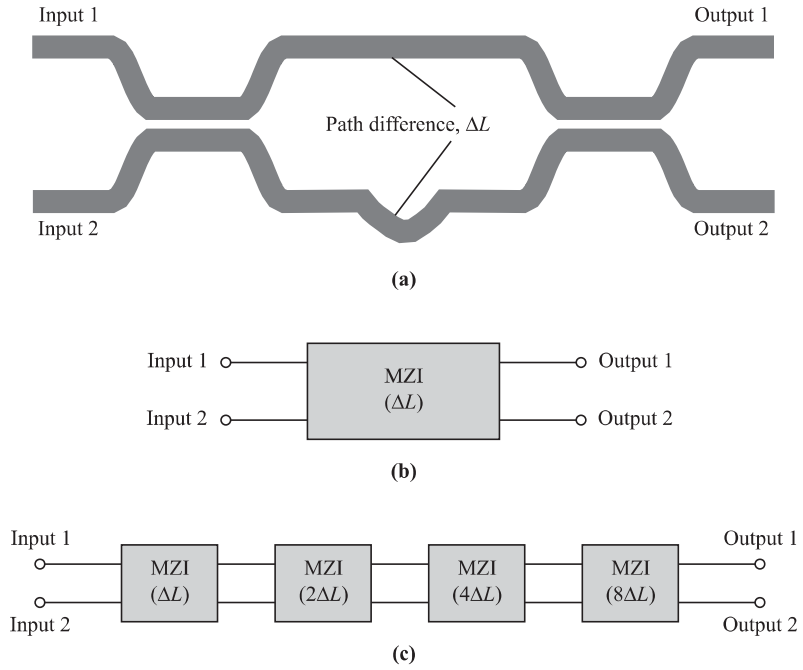

Figure 3.21 (a) An MZI constructed by interconnecting two 3 dB directional couplers. (b) A block diagram representation of the MZI in (a). ΔL denotes the path difference between the two arms. (c) A block diagram of a four-stage Mach-Zehnder interferometer, which uses different path length differences in each stage.

Mach-Zehnder interferometers are useful as both filters and (de)multiplexers. Even though there are better technologies for making narrow band filters, for example, dielectric multicavity thin-film filters, MZIs are still useful in realizing wide band filters. For example, MZIs can be used to separate the wavelengths in the 1.3 μm and 1.55 μm bands. Narrow band MZI filters are fabricated by cascading a number of stages, as we will see, and this leads to larger losses. In principle, very good crosstalk performance can be achieved using MZIs if the wavelengths are spaced such that the undesired wavelengths occur at, or close to, the nulls of the power transfer function. However, in practice, the wavelengths cannot be fixed precisely (for example, the wavelengths drift because of temperature variations or age). Moreover, the coupling ratio of the directional couplers is not 50:50 and could be wavelength dependent. As

a result, the crosstalk performance is far from the ideal situation. Also the passband of narrow band MZIs is not flat. In contrast, the dielectric multicavity thin-film filters can have flat passbands and good stop bands.

MZIs are useful as two-input, two-output multiplexers and demultiplexers. They can also be used as tunable filters, where the tuning is achieved by varying the temperature of one of the arms of the device. This causes the refractive index of that arm to change, which in turn affects the phase relationship between the two arms and causes a different wavelength to be coupled out. The tuning time required is of the order of several milliseconds. For higher channel-count multiplexers and demultiplexers, better technologies are available today. One example is the *arrayed waveguide grating* (AWG) described in the next section. Since understanding the MZI is essential to understanding the AWG, we will now describe the principle of operation of MZIs.

Principle of Operation

Consider the operation of the MZI as a demultiplexer; so only one input, say, input 1, has a signal (see Figure 3.21(a)). After the first directional coupler, the input signal power is divided equally between the two arms of the MZI, but the signal in one arm has a phase shift of $\pi/2$ with respect to the other. Specifically, the signal in the lower arm lags the one in the upper arm in phase by $\pi/2$, as discussed in Section 3.1. This is best understood from (3.1). Since there is a length difference of ΔL between the two arms, there is a further phase lag of $\beta\Delta L$ introduced in the signal in the lower arm. In the second directional coupler, the signal from the lower arm undergoes another phase delay of $\pi/2$ in going to the first output relative to the signal from the upper arm. Thus the total relative phase difference at the first or upper output between the two signals is $\pi/2 + \beta\Delta L + \pi/2$. At the output directional coupler, in going to the second output, the signal from the upper arm lags the signal from the lower arm in phase by $\pi/2$. Thus the total relative phase difference at the second or lower output between the two signals is $\pi/2 + \beta\Delta L - \pi/2 = \beta\Delta L$.

If $\beta\Delta L = k\pi$ and k is odd, the signals at the first output add in phase, whereas the signals at the second output add with opposite phases and thus cancel each other. Thus the wavelengths passed from the first input to the first output are those wavelengths for which $\beta\Delta L = k\pi$ and k is odd. The wavelengths passed from the first input to the second output are those wavelengths for which $\beta\Delta L = k\pi$ and k is even. This could have been easily deduced from the transfer function of the MZI in the following equation (3.14), but this detailed explanation will help in the understanding of the arrayed waveguide grating (Section 3.3.8).

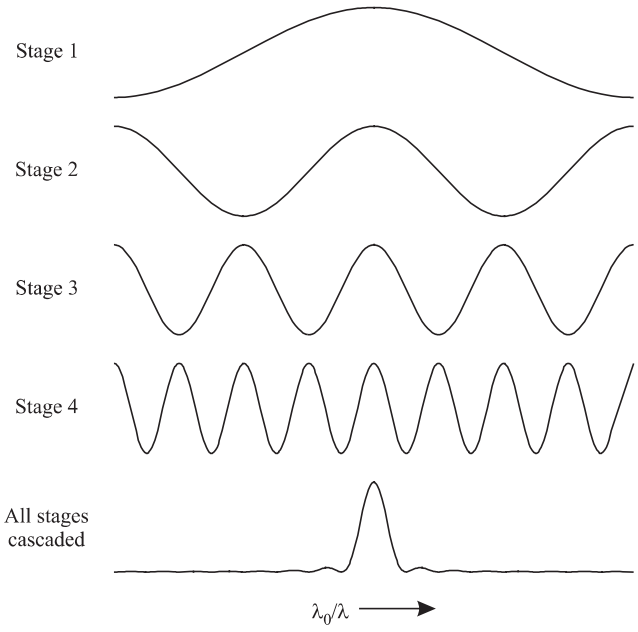

Figure 3.22 Transfer functions of each stage of a multistage MZI.

Assume that the difference between these path lengths is ΔL and that only one input, say, input 1, is active. Then it can be shown (see Problem 3.14) that the power transfer function of the Mach-Zehnder interferometer is given by

$$\begin{pmatrix} T_{11}(f) \\ T_{12}(f) \end{pmatrix} = \begin{pmatrix} \sin^2(\beta \Delta L/2) \\ \cos^2(\beta \Delta L/2) \end{pmatrix}. \tag{3.14}$$

Thus the path difference between the two arms, ΔL, is the key parameter characterizing the transfer function of the MZI. We will represent the MZI of Figure 3.21(a) using the block diagram of Figure 3.21(b).

Now consider k MZIs interconnected, as shown in Figure 3.21(c) for $k = 4$. Such a device is termed a *multistage Mach-Zehnder interferometer.* The path length difference for the kth MZI in the cascade is assumed to be $2^{k-1}\Delta L$. The transfer function of each MZI in this multistage MZI together with the power transfer function of the entire filter is shown in Figure 3.22. The power transfer function of the multistage MZI is also shown on a decibel scale in Figure 3.23.

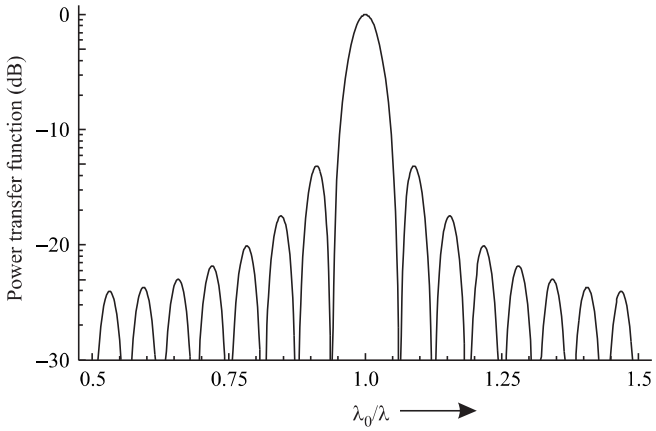

Figure 3.23 Transfer function of a multistage Mach-Zehnder interferometer.

We will now describe how an MZI can be used as a 1×2 demultiplexer. Since the device is reciprocal, it follows from the principles of electromagnetics that if the inputs and outputs are interchanged, it will act as a 2×1 multiplexer.

Consider a single MZI with a fixed value of the path difference ΔL. Let one of the inputs, say, input 1, be a wavelength division multiplexed signal with all the wavelengths chosen to coincide with the peaks or troughs of the transfer function. For concreteness, assume the propagation constant $\beta = 2\pi n_{\text{eff}}/\lambda$, where n_{eff} is the effective refractive index of the waveguide. The input wavelengths λ_i would have to be chosen such that $n_{\text{eff}}\Delta L/\lambda_i = m_i/2$ for some positive integer m_i. The wavelengths λ_i for which m is odd would then appear at the first output (since the transfer function is $\sin^2(m_i\pi/2)$), and the wavelengths for which m_i is even would appear at the second output (since the transfer function is $\cos^2(m_i\pi/2)$).

If there are only two wavelengths, one for which m_i is odd and the other for which m_i is even, we have a 1×2 demultiplexer. The construction of a $1 \times n$ demultiplexer when n is a power of two, using $n - 1$ MZIs, is left as an exercise (Problem 3.15). But there is a better method of constructing higher channel count demultiplexers, which we describe next.

3.3.8 Arrayed Waveguide Grating

An *arrayed waveguide grating* (AWG) is a generalization of the Mach-Zehnder interferometer. This device is illustrated in Figure 3.24. It consists of two multiport

couplers interconnected by an array of waveguides. The MZI can be viewed as a device where *two* copies of the same signal, but shifted in phase by different amounts, are added together. The AWG is a device where *several* copies of the same signal, but shifted in phase by different amounts, are added together.

The AWG has several uses. It can be used as an $n \times 1$ *wavelength multiplexer.* In this capacity, it is an n-input, 1-output device where the n inputs are signals at different wavelengths that are combined onto the single output. The inverse of this function, namely, $1 \times n$ *wavelength demultiplexing,* can also be performed using an AWG. Although these wavelength multiplexers and demultiplexers can also be built using MZIs interconnected in a suitable fashion, it is preferable to use an AWG. Relative to an MZI chain, an AWG has lower loss and flatter passband, and is easier to realize on an integrated-optic substrate. The input and output waveguides, the multiport couplers, and the arrayed waveguides are all fabricated on a single substrate. The substrate material is usually silicon, and the waveguides are silica, Ge-doped silica, or SiO_2-Ta_2O_5. Thirty-two–channel AWGs are commercially available, and smaller AWGs are being used in WDM transmission systems. Their temperature coefficient (0.01 nm/°C) is not as low as those of some other competing technologies such as fiber gratings and multilayer thin-film filters. So we will need to use active temperature control for these devices.

Another way to understand the working of the AWG as a demultiplexer is to think of the multiport couplers as lenses and the array of waveguides as a prism. The input coupler collimates the light from an input waveguide to the array of waveguides. The array of waveguides acts like a prism, providing a wavelength-dependent phase shift, and the output coupler focuses different wavelengths on different output waveguides.

The AWG can also be used as a static wavelength crossconnect. However, this wavelength crossconnect is not capable of achieving an arbitrary routing pattern. Although several interconnection patterns can be achieved by a suitable choice of

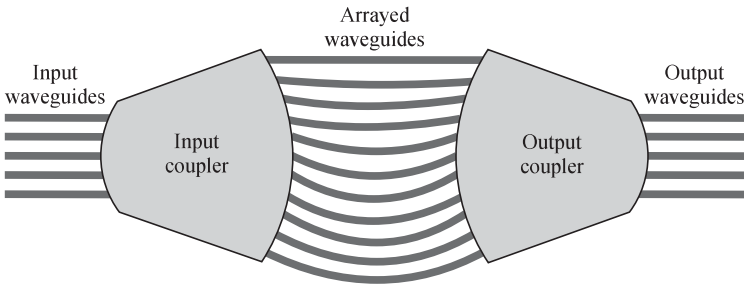

Figure 3.24 An arrayed waveguide grating.

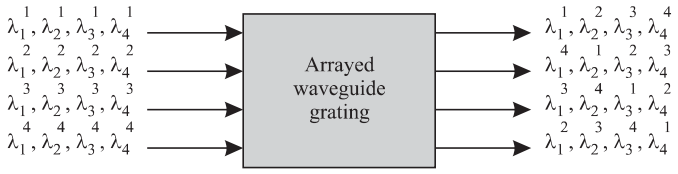

Figure 3.25 The crossconnect pattern of a static wavelength crossconnect constructed from an arrayed waveguide grating. The device routes signals from an input to an output based on their wavelength.

the wavelengths and the FSR of the device, the most useful one is illustrated in Figure 3.25. This figure shows a 4×4 static wavelength crossconnect using four wavelengths with one wavelength routed from each of the inputs to each of the outputs.

In order to achieve this interconnection pattern, the operating wavelengths and the FSR of the AWG must be chosen suitably. The FSR of the AWG is derived in Problem 3.17. Given the FSR, we leave the determination of the wavelengths to be used to achieve this interconnection pattern as another exercise (Problem 3.18).

Principle of Operation

Consider the AWG shown in Figure 3.24. Let the number of inputs and outputs of the AWG be denoted by n. Let the couplers at the input and output be $n \times m$ and $m \times n$ in size, respectively. Thus the couplers are interconnected by m waveguides. We will call these waveguides *arrayed waveguides* to distinguish them from the input and output waveguides. The lengths of these waveguides are chosen such that the difference in length between consecutive waveguides is a constant denoted by ΔL. The MZI is a special case of the AWG, where $n = m = 2$. We will now determine which wavelengths will be transmitted from a given input to a given output. The first coupler splits the signal into m parts. The relative phases of these parts are determined by the distances traveled in the coupler from the input waveguides to the arrayed waveguides. Denote the differences in the distances traveled (relative to any one of the input waveguides and any one of the arrayed waveguides) between input waveguide i and arrayed waveguide k by d_{ik}^{in}. Assume that arrayed waveguide k has a path length larger than arrayed waveguide $k-1$ by ΔL. Similarly, denote the differences in the distances traveled (relative to any one of the arrayed waveguides and any one of the output waveguides) between arrayed waveguide k and output

waveguide j by d_{kj}^{out}. Then, the relative phases of the signals from input i to output j traversing the m different paths between them are given by

$$\phi_{ijk} = \frac{2\pi}{\lambda}(n_1 d_{ik}^{\text{in}} + n_2 k\Delta L + n_1 d_{kj}^{\text{out}}), \qquad k = 1, \ldots, m. \tag{3.15}$$

Here, n_1 is the refractive index in the input and output directional couplers, and n_2 is the refractive index in the arrayed waveguides. From input i, those wavelengths λ, for which ϕ_{ijk}, $k = 1, \ldots, m$, differ by a multiple of 2π will add in phase at output j. The question is, Are there any such wavelengths?

If the input and output couplers are designed such that $d_{ik}^{\text{in}} = d_i^{\text{in}} + k\delta_i^{\text{in}}$ and $d_{kj}^{\text{out}} = d_j^{\text{out}} + k\delta_j^{\text{out}}$, then (3.15) can be written as

$$\begin{aligned}\phi_{ijk} &= \frac{2\pi}{\lambda}(n_1 d_i^{\text{in}} + n_1 d_j^{\text{out}}) \\ &\quad + \frac{2\pi k}{\lambda}(n_1\delta_i^{\text{in}} + n_2\Delta L + n_1\delta_j^{\text{out}}), \qquad k = 1, \ldots, m.\end{aligned} \tag{3.16}$$

Such a construction is possible and is called the *Rowland circle construction*. It is illustrated in Figure 3.26 and discussed further in Problem 3.16. Thus wavelengths λ that are present at input i and that satisfy $n_1\delta_i^{\text{in}} + n_2\Delta L + n_1\delta_j^{\text{out}} = p\lambda$ for some integer p add in phase at output j.

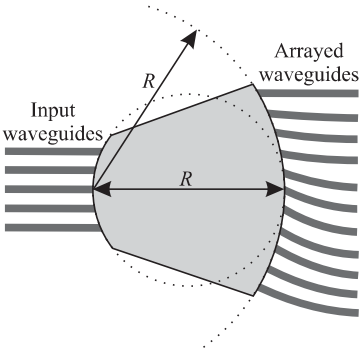

Figure 3.26 The Rowland circle construction for the couplers used in the AWG. The arrayed waveguides are located on the arc of a circle, called the *grating circle*, whose center is at the end of the central input (output) waveguide. Let the *radius* of this circle be denoted by R. The other input (output) waveguides are located on the arc of a circle whose *diameter* is equal to R; this circle is called the *Rowland circle*. The vertical spacing between the arrayed waveguides is chosen to be constant.

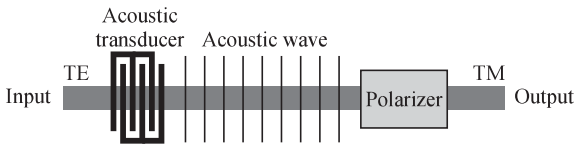

Figure 3.27 A simple AOTF. An acoustic wave introduces a grating whose pitch depends on the frequency of the acoustic wave. The grating couples energy from one polarization mode to another at a wavelength that satisfies the Bragg condition.

For use as a demultiplexer, all the wavelengths are present at the same input, say, input i. Therefore, if the wavelengths, $\lambda_1, \lambda_2, \ldots, \lambda_n$ in the WDM system satisfy $n_1\delta_i^{\text{in}} + n_2\Delta L + n_1\delta_j^{\text{out}} = p\lambda_j$ for some integer p, we infer from (3.16) that these wavelengths are demultiplexed by the AWG. Note that though δ_i^{in} and ΔL are necessary to define the precise set of wavelengths that are demultiplexed, the (minimum) spacing between them is independent of δ_i^{in} and ΔL, and determined primarily by δ_j^{out}.

Note in the preceding example that if wavelength λ'_j satisfies $n_1\delta_i^{\text{in}} + n_2\Delta L + n_1\delta_j^{\text{out}} = (p+1)\lambda'_j$, then both λ_j and λ'_j are "demultiplexed" to output j from input i. Thus like many of the other filter and multiplexer/demultiplexer structures we have studied, the AWG has a periodic response (in frequency), and all the wavelengths must lie within one FSR. The derivation of an expression for this FSR is left as an exercise (Problem 3.17).

3.3.9 *Acousto-Optic Tunable Filter*

The acousto-optic tunable filter is a versatile device. It is probably the only known *tunable* filter that is capable of selecting several wavelengths simultaneously. This capability can be used to construct a wavelength crossconnect, as we will explain later in this section.

The acousto-optic tunable filter (AOTF) is one example of several optical devices whose construction is based on the interaction of sound and light. Basically, an acoustic wave is used to create a Bragg grating in a waveguide, which is then used to perform the wavelength selection. Figure 3.27 shows a simple version of the AOTF. We will see that the operation of this AOTF is dependent on the state of polarization of the input signal. Figure 3.28 shows a more realistic polarization-independent implementation in integrated optics.

Principle of Operation

Consider the device shown in Figure 3.27. It consists of a waveguide constructed from a birefringent material and supporting only the lowest-order TE and TM modes (see Section 2.3.4). We assume that the input light energy is entirely in the TE mode. A *polarizer*, which selects only the light energy in the TM mode, is placed at the other end of the channel waveguide. If, somehow, the light energy in a narrow spectral range around the wavelength to be selected is converted to the TM mode, while the rest of the light energy remains in the TE mode, we have a wavelength-selective filter. This conversion is effected in an AOTF by launching an acoustic wave along, or opposite to, the direction of propagation of the light wave.

As a result of the propagation of the acoustic wave, the density of the medium varies in a periodic manner. The period of this density variation is equal to the wavelength of the acoustic wave. This periodic density variation acts as a Bragg grating. From the discussion of such gratings in Section 3.3.3, it follows that if the refractive indices n_{TE} and n_{TM} of the TE and TM modes satisfy the Bragg condition

$$\frac{n_{TM}}{\lambda} = \frac{n_{TE}}{\lambda} \pm \frac{1}{\Lambda}, \tag{3.17}$$

then light couples from one mode to the other. Thus light energy in a narrow spectral range around the wavelength λ that satisfies (3.17) undergoes TE to TM mode conversion. Thus the device acts as a narrow bandwidth filter when only light energy in the TE mode is input and only the light energy in the TM mode is selected at the output, as shown in Figure 3.27.

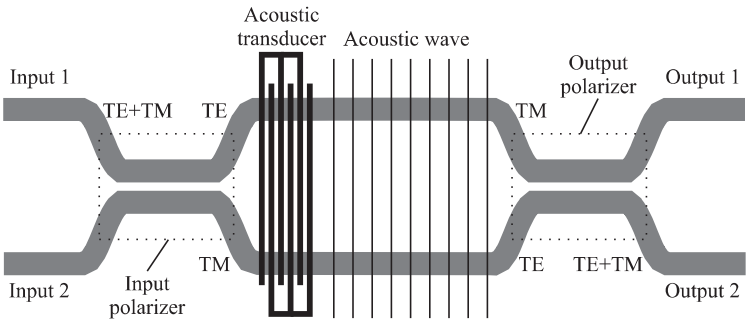

Figure 3.28 A polarization-independent integrated-optics AOTF. A polarizer splits the input signal into its constituent polarization modes, and each mode is converted in two separate arms, before being recombined at the output.

In $LiNbO_3$, the TE and TM modes have refractive indices n_{TE} and n_{TM} that differ by about 0.07. If we denote this refractive index difference by (Δn), the Bragg condition (3.17) can be written as

$$\lambda = \Lambda(\Delta n). \tag{3.18}$$

The wavelength that undergoes mode conversion and thus lies in the passband of the AOTF can be selected, or tuned, by suitably choosing the acoustic wavelength Λ. In order to select a wavelength of 1.55 μm, for $(\Delta n) = 0.07$, using (3.18), the acoustic wavelength is about 22 μm. Since the velocity of sound in $LiNbO_3$ is about 3.75 km/s, the corresponding RF frequency is about 170 MHz. Since the RF frequency is easily tuned, the wavelength selected by the filter can also be easily tuned. The typical insertion loss is about 4 dB.

The AOTF considered here is a polarization-dependent device since the input light energy is assumed to be entirely in the TE mode. A polarization-independent AOTF, shown in Figure 3.28, can be realized in exactly the same manner as a polarization-independent isolator by decomposing the input light signal into its TE and TM constituents and sending each constituent separately through the AOTF and recombining them at the output.

Transfer Function

Whereas the Bragg condition determines the wavelength that is selected, the width of the filter passband is determined by the length of the acousto-optic interaction. The longer this interaction, and hence the device, the narrower the passband. It can be shown that the wavelength dependence of the fraction of the power transmitted by the AOTF is given by

$$T(\lambda) = \frac{\sin^2\left((\pi/2)\sqrt{1 + (2\Delta\lambda/\Delta)^2}\right)}{1 + (2\Delta\lambda/\Delta)^2}.$$

This is plotted in Figure 3.29. Here $\Delta\lambda = \lambda - \lambda_0$, where λ_0 is the optical wavelength that satisfies the Bragg condition, and $\Delta = \lambda_0^2/l\Delta n$ is a measure of the filter passband width. Here, l is the length of the device (or, more correctly, the length of the acousto-optic interaction). It can be shown that the full width at half-maximum (FWHM) bandwidth of the filter is $\approx 0.8\Delta$ (Problem 3.20). This equation clearly shows that the longer the device, the narrower the passband. However, there is a trade-off here: the tuning speed is inversely proportional to l. This is because the tuning speed is essentially determined by the time it takes for a sound wave to travel the length of the filter.

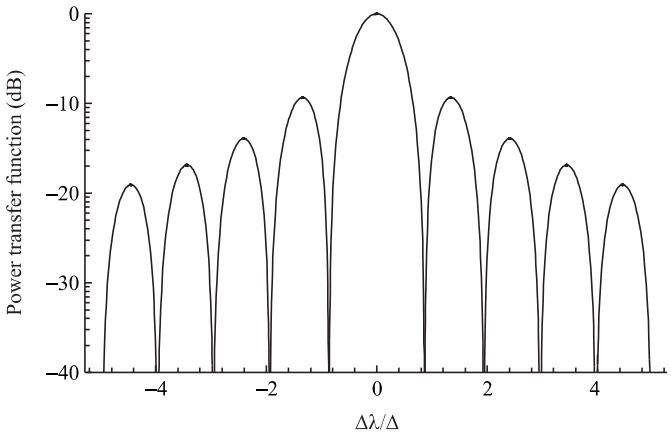

Figure 3.29 The power transfer function of the acousto-optic tunable filter.

AOTF as a Wavelength Crossconnect

The polarization-independent AOTF illustrated in Figure 3.28 can be used as a two-input, two-output dynamic wavelength crossconnect. We studied the operation of this device as a filter earlier; in this case, only one of the inputs was active. We leave it as an exercise (Problem 3.21) to show that when the second input is also active, the energy at the wavelength λ satisfying the Bragg phase-matching condition (3.18) is *exchanged* between the two ports. This is illustrated in Figure 3.30(a), where the wavelength λ_1 satisfies the Bragg condition and is exchanged between the ports.

Now the AOTF has one remarkable property that is not shared by any other tunable filter structure we know. By launching multiple acoustic waves *simultaneously*, the Bragg condition (3.18) can be satisfied for multiple optical wavelengths simultaneously. Thus multiple wavelength exchanges can be accomplished simultaneously between two ports with a single device of the form shown in Figure 3.28. This is illustrated in Figure 3.30(b), where the wavelengths λ_1 and λ_4 are exchanged between the ports. Thus this device performs the same routing function as the static crossconnect of Figure 3.7. However, the AOTF is a completely general two-input, two-output *dynamic* crossconnect since the routing pattern, or the set of wavelengths to be exchanged, can be changed easily by varying the frequencies of the acoustic waves launched in the device. In principle, larger dimensional dynamic crossconnects (with more input and output ports) can be built by suitably cascading 2×2 crossconnects. We will see in Section 3.7, however, that there are better ways of building large-scale crossconnects.

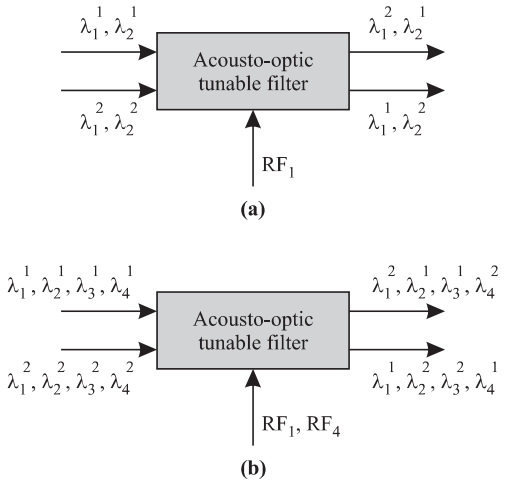

Figure 3.30 Wavelength crossconnects constructed from acousto-optic tunable filters. (a) The wavelength λ_1 is exchanged between the two ports. (b) The wavelengths λ_1 and λ_4 are simultaneously exchanged between the two ports by the simultaneous launching of two appropriate acoustic waves.

As of this writing, the AOTF has not yet lived up to its promise either as a versatile tunable filter or a wavelength crossconnect. One reason for this is the high level of crosstalk that is present in the device. As can be seen from Figure 3.29, the first side lobe in its power transfer function is not even 10 dB below the peak transmission. This problem can be alleviated to some extent by cascading two such filters. In fact, the cascade can even be built on a single substrate. But even then the first side lobe would be less than 20 dB below the peak transmission. It is harder to cascade more such devices without facing other problems such as an unacceptably high transmission loss. Another reason for the comparative failure of the AOTF today is that the passband width is fairly large (100 GHz or more) even when the acousto-optic interaction length is around 1 inch (Problem 3.22). This makes it unsuitable for use in dense WDM systems where channel spacings are now down to 50 GHz. Devices with larger interaction lengths are more difficult to fabricate. However, some recent theoretical work [Son95] indicates that some of these problems, particularly crosstalk, may be solvable. The crosstalk problems that arise in AOTFs when used as wavelength crossconnects are discussed in detail in [Jac96].

3.3.10 High Channel Count Multiplexer Architectures

With the number of wavelengths continuously increasing, designing multiplexers and demultiplexers to handle large numbers of wavelengths has become an important problem. The desired attributes of these devices are the same as what we saw at the beginning of Section 3.3. Our discussion will be based on demultiplexers, but these demultiplexers can all be used as multiplexers as well. In fact, in bidirectional applications, where some wavelengths are transmitted in one direction over a fiber and others in the opposite direction over the same fiber, the same device acts as a multiplexer for some wavelengths and a demultiplexer for others. We describe several architectural approaches to construct high channel count demultiplexers below.

Serial

In this approach, the demultiplexing is done one wavelength at a time. The demultiplexer consists of W filter stages in series, one for each of the W wavelengths. Each filter stage demultiplexes a wavelength and allows the other wavelengths to pass through. The architecture of the dielectric thin-film demultiplexer shown in Figure 3.20 is an example. One advantage of this architecture is that the filter stages can potentially be added one at a time, as more wavelengths are added. This allows a "pay as you grow" approach.

Serial approaches work for demultiplexing relatively small numbers of channels but do not scale to handle a large number of channels. This is because the insertion loss (in decibels) of the demultiplexer increases almost linearly with the number of channels to be demultiplexed. Moreover, different channels see different insertion losses based on the order in which the wavelengths are demultiplexed, which is not a desirable feature.

Single Stage

Here, all the wavelengths are demultiplexed together in a single stage. The AWG shown in Figure 3.24 is an example of such an architecture. This approach provides relatively lower losses and better loss uniformity, compared to the serial approach. However, the number of channels that can be demultiplexed is limited by the maximum number of channels that can be handled by a single device, typically around 40 channels in commercially available AWGs today.

Multistage Banding

Going to larger channel counts requires the use of multiple demultiplexing stages, due to the limitations of the serial and single-stage approaches discussed above. A popular approach used today is to divide the wavelengths into *bands*. For example,

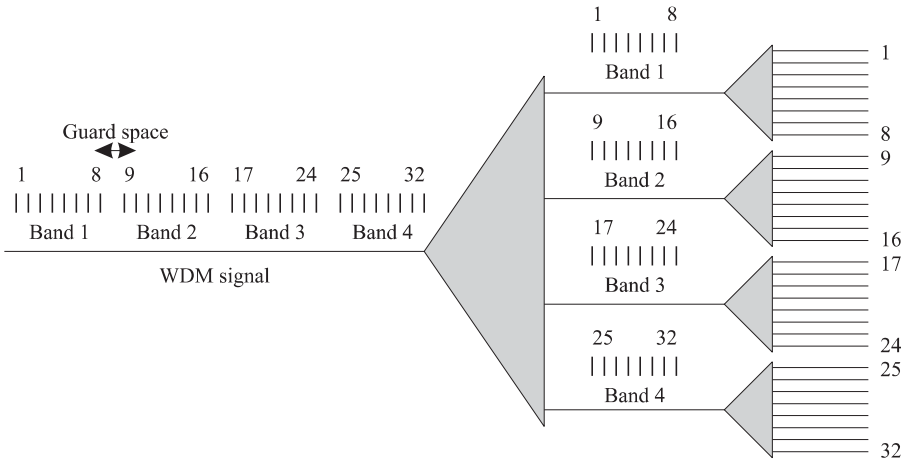

Figure 3.31 A two-stage demultiplexing approach using bands. A 32-channel demultiplexer is realized using four bands of 8 channels each.

a total of 32 wavelengths may be divided into four bands, each with 8 wavelengths. The demultiplexing is done in two stages, as shown in Figure 3.31. In the first the set of wavelengths is demultiplexed into bands. In the second stage, the bands are demultiplexed, and individual wavelengths are extracted. The scheme can be extended to more than two stages as well. It is also modular in that the demultiplexers in the second stage (or last stage in a multistage scheme) can be populated one band at a time.

One drawback of the banding approach is that we will usually need to leave a "guard" space between bands, as shown in Figure 3.31. This guard space allows the first-stage filters to be designed to provide adequate crosstalk suppression while retaining a low insertion loss.

Multistage Interleaving

Interleaving provides another approach to realizing large channel count demultiplexers. A two-stage *interleaver* is shown in Figure 3.32. In this approach the first stage separates the wavelengths into two groups. The first group consists of wavelengths 1, 3, 5, . . . and the second group consists of wavelengths 2, 4, 6, ... The second stage extracts the individual wavelengths. This approach is also modular in the sense that the last stage of demultiplexers can be populated as needed. More than two stages can be used if needed as well.

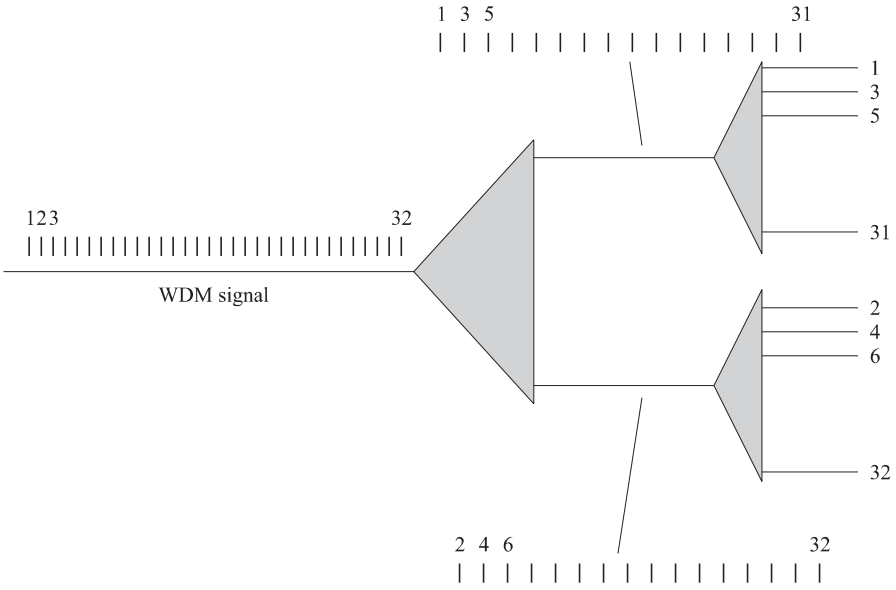

Figure 3.32 A two-stage multiplexing approach using interleaving. In this 32-channel demultiplexer, the first stage picks out every alternate wavelength, and the second stage extracts the individual wavelength.

A significant benefit of this approach is that the filters in the last stage can be much wider than the channel width. As an example, suppose we want to demultiplex a set of 32 channels spaced 50 GHz apart. After the first stage of demultiplexing, the channels are spaced 100 GHz apart, as shown in Figure 3.32. So demultiplexers with a broader passband suitable for demultiplexing 100 GHz spaced channels can be used in the second stage. In contrast, the single-stage or serial approach would require the use of demultiplexers capable of demultiplexing 50 GHz spaced channels, which are much more difficult to build. Carrying this example further, the second stage itself can in turn be made up of two stages. The first stage extracts every other 100 GHz channel, leading to a 200 GHz interchannel spacing after this stage. The final stage can then use even broader filters to extract the individual channels. Another advantage of this approach is that no guard bands are required in the channel plan.

The challenges with the interleaving approach lie in realizing the demultiplexers that perform the interleaving at all the levels except the last level. In principle,

any periodic filter can be used as an interleaver by matching its period to the desired channel spacing. For example, a fiber-based Mach-Zehnder interferometer is a common choice. These devices are now commercially available, and interleaving is becoming a popular approach toward realizing high channel count multiplexers and demultiplexers.

3.4 Optical Amplifiers

In an optical communication system, the optical signals from the transmitter are attenuated by the optical fiber as they propagate through it. Other optical components, such as multiplexers and couplers, also add loss. After some distance, the cumulative loss of signal strength causes the signal to become too weak to be detected. Before this happens, the signal strength has to be restored. Prior to the advent of optical amplifiers over the last decade, the only option was to regenerate the signal, that is, receive the signal and retransmit it. This process is accomplished by *regenerators*. A regenerator converts the optical signal to an electrical signal, cleans it up, and converts it back into an optical signal for onward transmission.

Optical amplifiers offer several advantages over regenerators. On one hand, regenerators are specific to the bit rate and modulation format used by the communication system. On the other hand, optical amplifiers are insensitive to the bit rate or signal formats. Thus a system using optical amplifiers can be more easily upgraded, for example, to a higher bit rate, without replacing the amplifiers. In contrast, in a system using regenerators, such an upgrade would require all the regenerators to be replaced. Furthermore, optical amplifiers have fairly large gain bandwidths, and as a consequence, a single amplifier can simultaneously amplify several WDM signals. In contrast, we would need a regenerator for each wavelength. Thus optical amplifiers have become essential components in high-performance optical communication systems.

Amplifiers, however, are not perfect devices. They introduce additional noise, and this noise accumulates as the signal passes through multiple amplifiers along its path due to the analog nature of the amplifier. The spectral shape of the gain, the output power, and the transient behavior of the amplifier are also important considerations for system applications. Ideally, we would like to have a sufficiently high output power to meet the needs of the network application. We would also like the gain to be flat over the operating wavelength range and to be insensitive to variations in input power of the signal. We will study the impact of optical amplifiers on the physical layer design of the system in Chapters 4 and 5. Here we explore their principle of operation.