Optical Communication Networks

EE654

Lectuer - 4

Spring 2016

Transmission System Engineering

4.1 System Model

Figure 5.1 shows a block diagram of the various components of a unidirectional WDM link. The transmitter consists of a set of DFB lasers, with or without external modulators, one for each wavelength. The signals at the different wavelengths are combined into a single fiber by means of an optical multiplexer. An optical power amplifier may be used to increase the transmission power. After some distance along the fiber, the signal is amplified by an optical in-line amplifier. Depending on the distance, bit rate, and type of fiber used, the signal may also be passed through a dispersion-compensating module, usually at each amplifier stage. At the receiving end, the signal may be amplified by an optical preamplifier before it is passed through a demultiplexer. Each wavelength is then received by a separate photodetector.

Throughout this chapter, we will be focusing on digital systems, although it is possible to transmit analog signals over fiber as well. The physical layer of the system must ensure that bits are transmitted from the source to their destination reliably. The measures of quality are the bit error rate (BER) and the additional power budget

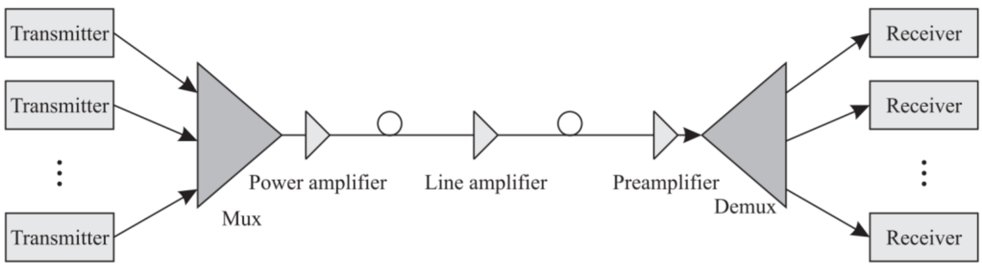

Figure 5.1 Components of a WDM link.

margin provided in the system. Usually the required bit error rates are of the order of 10^{-9} to 10^{-15}, typically 10^{-12}. The BER depends on the amount of noise as well as other impairments that are present in the system. Unless otherwise stated, we will assume that non-return-to-zero (NRZ) modulation is used. In some specific cases, such as chromatic dispersion, we consider both NRZ and return-to-zero (RZ) modulation.

The physical layer is also responsible for the link initialization and link take-down procedures, which are necessary to prevent exposure to potentially harmful laser radiation. This aspect is dealt with in Chapter 8.

We will look at the different components that are part of a system, including the transmitters, receivers, optical amplifiers, wavelength multiplexers, demultiplexers and switches, and the fiber itself, and we will discuss various forms of system impairments that arise from each of these components. Table B.1 in Appendix B summarizes the large number of parameters used in this chapter.

4.2 Power Penalty

The physical layer design must take into account the effect of a number of system impairments as previously discussed. Usually, each impairment results in a *power penalty* to the system. In the presence of an impairment, a higher signal power will be required at the receiver in order to maintain a desired bit error rate. One way to define the power penalty is as the increase in signal power required (in dB) to maintain the same bit error rate in the presence of impairments. Another way to define the power penalty is as the reduction in signal-to-noise ratio as quantified by the value of γ (the argument to the $Q(.)$ function as defined in Section 4.4.6) due to a specific impairment. We will be using the latter definition since it is easier to calculate and consistent with popular usage.

Let P_1 denote the optical power received during a 1 bit, and P_0 the power received during a 0 bit without any system impairments. The corresponding electrical currents are given by $\mathcal{R}P_1$ and $\mathcal{R}P_0$, respectively, where $\mathcal{R}$ is the responsivity of the photodetector.

Let σ_1 and σ_0 denote the noise standard deviations during a 1 bit and a 0 bit, respectively. Assume that the noise is Gaussian. The bit error rate, assuming equally likely 1s and 0s, is obtained from (4.14) as

$$\mathrm{BER} = Q\left(\frac{\mathcal{R}(P_1 - P_0)}{\sigma_1 + \sigma_0}\right). \tag{5.1}$$

This expression assumes that the receiver's decision threshold is set to the optimal value indicated by (4.12).

In the presence of impairments, let P_1', P_0', σ_1', σ_0' denote the received powers and noise standard deviations, respectively. Assuming an optimized threshold setting, the power penalty is given by

$$\mathrm{PP} = -10\log\left(\frac{\frac{\mathcal{R}(P_1'-P_0')}{\sigma_1'+\sigma_0'}}{\frac{\mathcal{R}(P_1-P_0)}{\sigma_1+\sigma_0}}\right). \tag{5.2}$$

Calculating the power penalty in general for the simple AC-coupled receiver discussed in Section 4.4.6 is somewhat more complicated, but we will see that it is the same as the penalty for the optimized receiver for two important cases of interest.

The first case of interest is when the dominant noise component is receiver thermal noise, for which $\sigma_0 = \sigma_1 = \sigma_{th}$. This is usually the case in unamplified direct detection *pin* receivers. In this case, or in any situation where the noise is independent of the signal power, the power penalty is given by

$$\mathrm{PP}_{\text{sig-indep}} = -10 \log \left(\frac{P_1' - P_0'}{P_1 - P_0} \right) \tag{5.3}$$

and the best threshold setting corresponds to the setting of a simple AC-coupled receiver.

The other case of interest is amplified systems, or systems with APD receivers. In amplified systems, the dominant noise component is usually the amplifier signal-spontaneous beat noise (see Section 4.4.5). In APD receivers, the dominant noise component is the shot noise, which is enhanced because of the APD gain (see Section 3.6.1). In amplified systems, and in systems with APD receivers, we can assume that $\sigma_1 \propto \sqrt{P_1}$; that is, the noise variance depends on the signal power. Assume also that $P_0 \ll P_1$. In this case, we can assume that $\sigma_1 \gg \sigma_0$. Here an optimized receiver would set its threshold close to the 0 level, whereas the simple receiver would still set its threshold at the average received power and would have a somewhat higher bit error rate. However, the power penalties turn out to be the same in both cases. This penalty is given by

$$\mathrm{PP}_{\text{sig-dep}} = -5 \log \left(\frac{P_1'}{P_1} \right). \tag{5.4}$$

Finally, it must be kept in mind that polarization plays an important role in many system impairments where signals interfere with each other. The worst case is usually when the interfering signals have the same state of polarization. However, the state of polarization of each signal varies slowly with time in a random manner, and thus we can expect the power penalties to vary with time as well. The system must be designed, however, to accommodate the worst case, usually identical polarizations.

4.3 Transmitter

The key system design parameters related to the transmitter are its output power, rise-/fall-time, extinction ratio, modulation type, side-mode suppression ratio, relative intensity noise (RIN), and wavelength stability and accuracy.

Table 5.1 An example system design that allocates power penalties for various transmission impairments.

Impairment	Allocation (dB)
Ideal γ	17
Transmitter	1
Crosstalk	1
Chromatic dispersion	2
Nonlinearities	1
Polarization-dependent loss	3
Component aging	3
Margin	3
Required γ	31

The output power depends on the type of transmitter. DFB lasers put out about 1 mW (0 dBm) to 10 mW (10 dBm) of power. An optical power amplifier can be used to boost the power, typically to as much as 50 mW (17 dBm). The upper limits on power are dictated by nonlinearities (Section 5.8) and safety considerations (Section 8.7).

The extinction ratio is defined as the ratio of the power transmitted when sending a 1 bit, P_1, to the power transmitted when sending a 0 bit, P_0. Assuming that we are limited to an average transmitted power P, we would like to have $P_1 = 2P$ and $P_0 = 0$. This would correspond to an extinction ratio $r = \infty$. Practical transmitters, however, have extinction ratios between 10 and 20. With an extinction ratio r, we have

$$P_0 = \frac{2P}{r+1}$$

and

$$P_1 = \frac{2rP}{r+1}.$$

Reducing the extinction ratio reduces the difference between the 1 and 0 levels at the receiver and thus produces a penalty. The power penalty due to a nonideal extinction ratio in systems limited by signal-independent noise is obtained from (5.3) as

$$\mathrm{PP}_{\text{sig-indep}} = -10 \log \frac{r-1}{r+1}.$$

Note that on the one hand this penalty represents the decrease in signal-to-noise ratio performance of a system with a nonideal extinction ratio relative to a system with infinite extinction ratio, assuming the same *average* transmitted power for both systems. On the other hand, if we assume that the two systems have the same peak transmit power, that is, the same power for a 1 bit, then the penalty can be calculated to be

$$\text{PP}_{\text{sig-indep}} = -10 \log \frac{r-1}{r}.$$

Lasers tend to be physically limited by peak transmit power. Most nonlinear effects also set a limit on the peak transmit power. However, eye safety regulation limits (see Section 8.7.1), are stated in terms of average power. The formula to be used depends on which factor actually limits the power for a particular system.

The penalty is higher when the system is limited by signal-dependent noise, which is typically the case in amplified systems (Section 4.4.5)—see Problem 5.10. This is due to the increased amount of noise present at the 0 level. Other forms of signal-dependent noise may arise in the system, such as *laser relative intensity noise,* which refers to intensity fluctuations in the laser output caused by reflections from fiber splices and connectors in the link.

The laser at the transmitter may be modulated directly, or a separate external modulator can be used. Direct modulation is cheaper but results in a broader spectral width due to chirp (Section 2.4). This will result in an added power penalty due to chromatic dispersion (see Section 2.4). Broader spectral width may also result in penalties when the signal is passed through optical filters, such as WDM muxes and demuxes. This penalty can be reduced by reducing the extinction ratio, which, in turn, reduces the chirp and, hence, the spectral width.

Wavelength stability of the transmitter is an important issue and is addressed in Sections 5.9 and 5.12.8.

4.4 Receiver

The key system parameters associated with a receiver are its *sensitivity* and *overload parameter*. The sensitivity is the average optical power required to achieve a certain bit error rate at a particular bit rate. It is usually measured at a bit error rate of 10^{-12} using a pseudorandom $2^{23} - 1$ bit sequence. The overload parameter is the maximum input power that the receiver can accept. Typical sensitivities of different types of receivers for a set of bit rates are shown in Table 5.2; a more detailed evaluation can be found in Section 4.4.6. APD receivers have higher sensitivities than pinFET receivers and are typically used in high-bit-rate systems operating at

Table 5.2 Typical sensitivities of different types of receivers in the 1.55 μm wavelength band. These receivers also operate in the 1.3 μm band, but the sensitivity may not be as good at 1.3 μm.

Bit Rate	Type	Sensitivity	Overload Parameter
155 Mb/s	pinFET	−36 dBm	−7 dBm
622 Mb/s	pinFET	−32 dBm	−7 dBm
2.5 Gb/s	pinFET	−23 dBm	−3 dBm
2.5 Gb/s	APD	−34 dBm	−8 dBm
10 Gb/s	pinFET	−18 dBm	−1 dBm
10 Gb/s	APD	−24 dBm	−6 dBm
40 Gb/s	pinFET	−7 dBm	3 dBm

and above 2.5 Gb/s. However, a pinFET receiver with an optical preamplifier has a sensitivity that is comparable to an APD receiver. The overload parameter defines the dynamic range of the receiver and can be as high as 0 dBm for 2.5 Gb/s receivers, regardless of the specific receiver type.

4.5 Optical Amplifiers

Optical amplifiers have become an essential component in transmission systems and networks to compensate for system losses. The most common optical amplifier today is the erbium-doped fiber amplifier (EDFA) operating in the C-band. L-band EDFAs and Raman amplifiers are also used. EDFAs are used in almost all amplified WDM systems, whereas Raman amplifiers are used in addition to EDFAs in many ultra-long-haul systems. These amplifiers are described in Section 3.4. In this section, we will focus mainly on EDFAs.

The EDFA has a gain bandwidth of about 35 nm in the 1.55 μm wavelength region. The great advantage of EDFAs is that they are capable of simultaneously amplifying many WDM channels. EDFAs spawned a new generation of transmission systems, and almost all optical fiber transmission systems installed over the last few years use EDFAs instead of repeaters. The newer L-band EDFAs are being installed today to increase the available bandwidth, and hence the number of wavelengths, in a single fiber.

Amplifiers are used in three different configurations, as shown in Figure 5.2. An optical *preamplifier* is used just in front of a receiver to improve its sensitivity. A *power amplifier* is used after a transmitter to increase the output power. A *line amplifier* is typically used in the middle of the link to compensate for link losses. The design of the amplifier depends on the configuration. A power amplifier is designed to provide the maximum possible output power. A preamplifier is designed to provide high gain and the highest possible sensitivity, that is, the least amount of additional noise. A line amplifier is designed to provide a combination of all of these.

Unfortunately, the amplifier is not a perfect device. There are several major imperfections that system designers need to worry about when using amplifiers in a system. First, an amplifier introduces noise, in addition to providing gain. Second, the gain of the amplifier depends on the total input power. For high input powers, the EDFA tends to saturate and the gain drops. This can cause undesirable power transients in networks. Finally, although EDFAs are a particularly attractive choice for WDM systems, their gain is not flat over the entire passband. Thus some channels see more gain than others. This problem gets worse when a number of amplifiers are cascaded.

We have studied optically preamplified receivers in Section 4.4.5. In this section, we will study the effect of gain saturation, gain nonflatness, noise, and power transients in systems with cascades of optical amplifiers.

4.5.1 Gain Saturation in EDFAs

An important consideration in designing amplified systems is the saturation of the EDFA. Depending on the pump power and the amplifier design itself, the output power of the amplifier is limited. As a result, when the input signal power is increased, the amplifier gain drops. This behavior can be captured approximately by the following equation:

$$G = 1 + \frac{P^{\text{sat}}}{P_{\text{in}}} \ln \frac{G_{\max}}{G}. \tag{5.5}$$

Here, $G_{\max}$ is the unsaturated gain, and G the saturated gain of the amplifier, P^{sat} is the amplifier's internal saturation power, and P_{in} is the input signal power. Figure 5.3 plots the amplifier gain as a function of the input signal power for a typical EDFA.

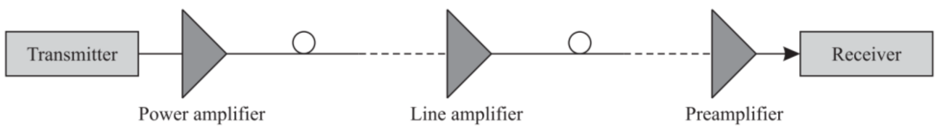

Figure 5.2 Power amplifiers, line amplifiers, and preamplifiers.

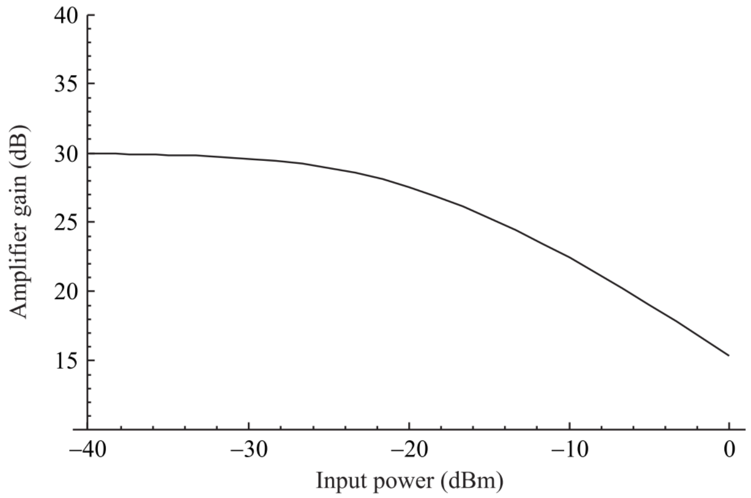

Figure 5.3 Gain saturation in an optical amplifier. Unsaturated gain $G_{\max} = 30$ dB and saturation power $P^{\text{sat}} = 10$ dBm.

For low input powers, the amplifier gain is at its unsaturated value, and at very high input powers, $G \to 1$ and the output power $P_{\text{out}} = P_{\text{in}}$. The output saturation power $P_{\text{out}}^{\text{sat}}$ is defined to be the output power at which the amplifier gain has dropped by 3 dB. Using (5.5) and the fact that $P_{\text{out}} = GP_{\text{in}}$, and assuming that $G \gg 1$, the output saturation power is given by

$$P_{\text{out}}^{\text{sat}} \approx P^{\text{sat}} \ln 2.$$

The saturation power is a function of the pump power and other amplifier parameters. It is quite common to have output saturation powers on the order of 10 to 100 mW (10 to 20 dBm).

There is no fundamental problem in operating an EDFA in saturation, and power amplifiers usually do operate in saturation. The only thing to keep in mind is that the saturated gain will be less than the unsaturated gain.

4.5.2 Gain Equalization in EDFAs

The flatness of the EDFA passband becomes a critical issue in WDM systems with cascaded amplifiers. The amplifier gain is not exactly the same at each wavelength. Small variations in gain between channels in a stage can cause large variations in the power difference between channels at the output of the chain. For example, if the gain variation between the worst channel and the best channel is 1 dB at each stage, after 10 stages it will be 10 dB, and the worst channel will have a much poorer signal-to-noise ratio than the best channel. This effect is shown in Figure 5.4(a). Building amplifiers with flat gain spectra is therefore very important (see Section 3.4.3) and is the best way to solve this problem. In practice, it is possible to design EDFAs to be inherently flat in the 1545–1560 nm wavelength region, and this is where many early WDM systems operate. However, systems with a larger number of channels will need to use the 1530–1545 nm wavelength range, where the gain of the EDFA is not flat.

The gain spectrum of L-band EDFAs is relatively flat over the L-band from about 1565 nm to about 1625 nm so that gain flattening over this band is not a significant issue.

At the system level, a few approaches have been proposed to overcome this lack of gain flatness. The first approach is to use preequalization, or preemphasis, as shown in Figure 5.4(b). Based on the overall gain shape of the cascade, the transmitted power per channel can be set such that the channels that see low gain are launched with higher powers. The goal of preequalization is to ensure that all channels are received with approximately the same signal-to-noise ratios at the receiver and fall within the receiver's dynamic range. However, the amount of equalization that can be done is limited, and other techniques may be needed to provide further equalization. Also this technique is difficult to implement in a network, as opposed to a point-to-point link.

The second approach is to introduce equalization at each amplifier stage, as shown in Figure 5.4(c). After each stage, the channel powers are equalized. This equalization can be done in many ways. One way is to demultiplex the channels, attenuate each channel differently, and then multiplex them back together. This approach involves using a considerable amount of hardware. It adds wavelength tolerance penalties due to the added muxes and demuxes (see Section 5.6.6). For these reasons, such an approach is impractical. Another approach is to use a multichannel filter, such as an acousto-optic tunable filter (AOTF). In an AOTF, each channel can be attenuated differently by applying a set of RF signals with different frequencies. Each RF signal controls the attenuation of a particular center wavelength, and by controlling the RF powers of each signal, it is possible to equalize the channel powers. However, an AOTF requires a large amount of RF drive power (on the order of 1 W) to equalize more than a few (2–4) channels. Both approaches introduce several decibels of additional loss and some power penalties due to crosstalk. The preferred solution today is to add an optical filter within the amplifier with a carefully designed passband to compensate for the gain spectrum of the amplifier so as to obtain a flat spectrum at its output. Both dielectric thin-film filters (Section 3.3.6) and long-period fiber gratings (Section 3.3.4) are good candidates for this purpose.

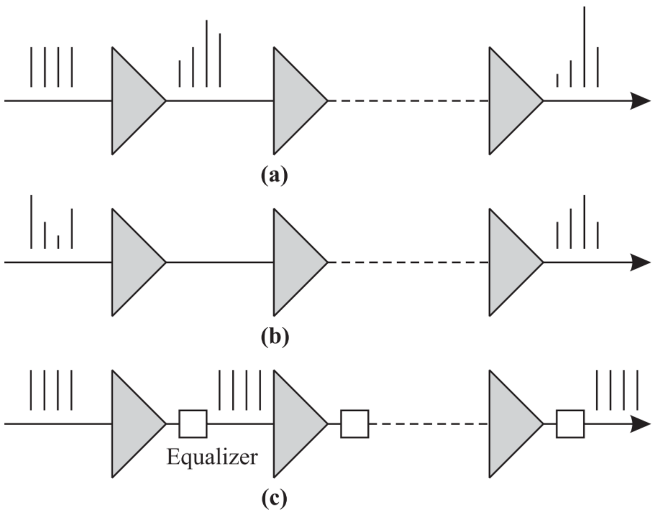

Figure 5.4 Effect of unequal amplifier gains at different wavelengths. (a) A set of channels with equal powers at the input to a cascaded system of amplifiers will have vastly different powers and signal-to-noise ratios at the output. (b) This effect can be reduced by preequalizing the channel powers. (c) Another way to reduce this effect is to introduce equalization at each amplifier stage. The equalization can be done using a filter inside the amplifier as well.

4.5.3 Amplifier Cascades

Consider a system of total length L with amplifiers spaced l km apart (see Figure 5.5). The loss between two stages is $e^{-\alpha l}$, where α is the fiber attenuation. Each amplifier adds some spontaneous emission noise. Thus the optical signal-to-noise ratio, OSNR (see Section 4.4.6 for the definition), gradually degrades along the chain.

The amplifier gain must be at least large enough to compensate for the loss between amplifier stages; otherwise, the signal (and hence the OSNR) will degrade rapidly with the number of stages. Consider what happens when we choose the unsaturated amplifier gain to be larger than the loss between stages. For the first few stages, the total input power (signal plus noise from the previous stages) to a stage increases with the number of stages. Consequently, the amplifiers begin to saturate and their gains drop. Farther along the chain, a spatial steady-state condition is reached where the amplifier output power and gain remains the same from stage to stage. These values, $\overline{P}_{\text{out}}$ and $\overline{G}$, respectively, can be computed by observing that

$$(\overline{P}_{\text{out}}e^{-\alpha l})\overline{G} + 2P_n B_o(\overline{G} - 1) = \overline{P}_{\text{out}}. \tag{5.6}$$

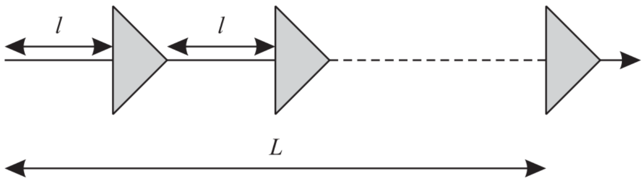

Figure 5.5 A system with cascaded optical amplifiers.

Here $\overline{P}_{\text{out}}e^{-\alpha l}$ is the total input power to the amplifier stage, and the second term, from (4.5), is the spontaneous emission noise added at this stage. Also from (5.5) we must have

$$\overline{G} = 1 + \frac{P^{\text{sat}}}{\overline{P}_{\text{out}}e^{-\alpha l}} \ln \frac{G_{\max}}{\overline{G}}. \tag{5.7}$$

Equations (5.6) and (5.7) can be solved simultaneously to compute the values of $\overline{P}_{\text{out}}$ and $\overline{G}$ (Problem 5.11). Observe from (5.6) that $\overline{G}e^{-\alpha l} < 1$; that is, the steady-state gain will be slightly smaller than the loss between stages, due to the added noise at each stage. Thus in designing a cascade, we must try to choose the saturated gain G to be as close to the loss between stages as possible.

Let us consider a simplified model of an amplifier cascade where we assume the saturated gain $G = e^{\alpha l}$. With L/l amplifiers in the system, the total noise power at the output, using (4.5), is

$$P_{\text{noise}}^{\text{tot}} = 2P_n B_o(G-1)L/l = 2P_n B_o(e^{\alpha l}-1)L/l. \tag{5.8}$$

Given a desired OSNR, the launched power P must satisfy

$$P \geq (\text{OSNR})P_{\text{noise}}^{\text{tot}} = (\text{OSNR})2P_n B_o(e^{\alpha l}-1)L/l.$$

Figure 5.6 plots the required power P versus amplifier spacing l. If we don't worry about nonlinearities, we would try to maximize l subject to limitations on transmit power and amplifier output power. The story changes in the presence of nonlinearities, as we will see in Section 5.8.

4.6 CrossTalk

Crosstalk is the general term given to the effect of other signals on the desired signal. Almost every component in a WDM system introduces crosstalk of some form or another. The components include filters, wavelength multiplexers/demultiplexers, switches, semiconductor optical amplifiers, and the fiber itself (by way of nonlinearities). Two forms of crosstalk arise in WDM systems: *interchannel crosstalk* and *intrachannel crosstalk*. The first case is when the crosstalk signal is at a wavelength sufficiently different from the desired signal's wavelength that the difference is larger than the receiver's electrical bandwidth. This form of crosstalk is called interchannel crosstalk. Interchannel crosstalk can also occur through more indirect interactions, for example, if one channel affects the gain seen by another channel, as with nonlinearities (Section 5.8). The second case is when the crosstalk signal is at the same wavelength as that of the desired signal or sufficiently close to it that the difference in wavelengths is within the receiver's electrical bandwidth. This form of crosstalk is called intrachannel crosstalk or, sometimes, *coherent crosstalk*. Intrachannel crosstalk effects can be much more severe than interchannel crosstalk, as we will see. In both cases, crosstalk results in a power penalty.

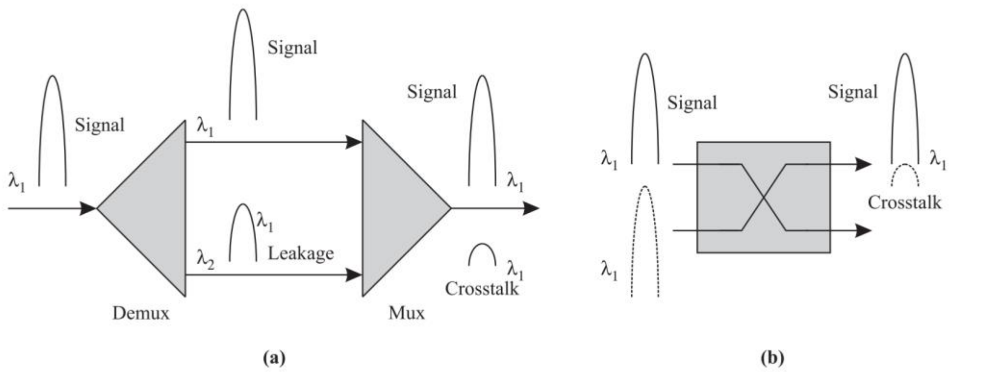

Figure 5.9 Sources of intrachannel crosstalk. (a) A cascaded wavelength demultiplexer and a multiplexer, and (b) an optical switch.

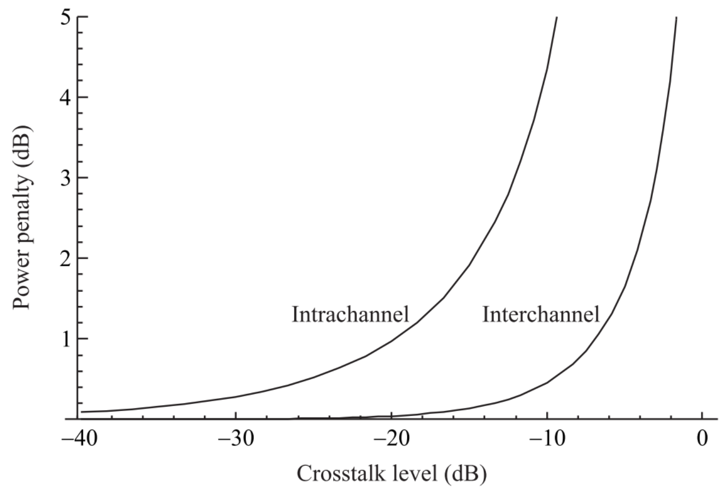

Figure 5.10 Intrachannel and interchannel crosstalk power penalties that are limited by thermal noise are shown as a function of crosstalk level, $10 \log \epsilon$.

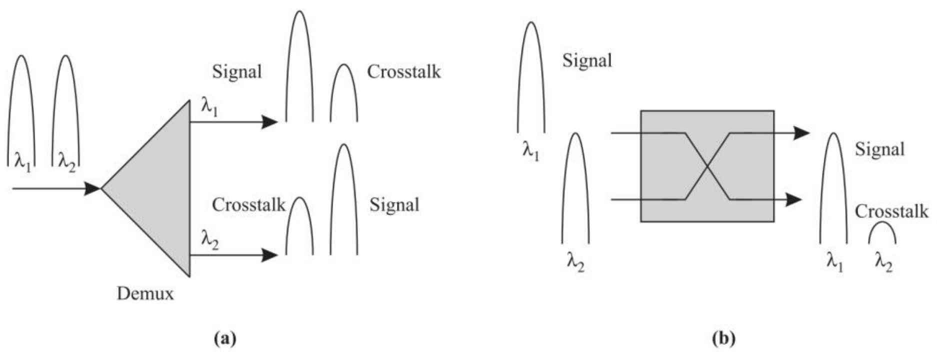

Figure 5.11 Sources of interchannel crosstalk. (a) An optical demultiplexer, and (b) an optical switch with inputs at different wavelengths.