1

Introduction

Optical Receiver and Transmitter. Figure 1.1 shows the block diagram of a typical optical receiver and transmitter. The optical signal from the fiber is received by a *photodetector* (PD), which produces a small output current proportional to the optical signal. This current is amplified and converted to a voltage by a *transimpedance amplifier* (TIA or TZA). The voltage signal is amplified further by either a *limiting amplifier* (LA) or an *automatic gain control amplifier* (AGC amplifier). The LA and AGC amplifier are collectively known as *main amplifiers* (MAs) or *post amplifiers.* The resulting signal, which is now several 100 mV strong, is fed into a *clock and data recovery circuit* (CDR), which extracts the clock signal and retimes the data signal. In high-speed receivers, a *demultiplexer* (DMUX) converts the fast serial data stream into n parallel lower-speed data streams that can be processed conveniently by the digital logic block. Some CDR designs (those with a parallel sampling architecture) perform the DMUX task as part of their functionality, and an explicit DMUX is not needed in this case [47]. The digital logic block descrambles or decodes the bits, performs error checks, extracts the payload data from the framing information, synchronizes to another clock domain, and so forth. The receiver just described also is known as a *3R receiver* because it performs signal *re-amplification* in the TIA (and the AGC amplifier, if present), signal *re-shaping* in the LA or CDR, and signal *re-timing* in the CDR.

On the transmitter side, the same process happens in reverse order. The parallel data from the digital logic block are merged into a single high-speed data stream using a *multiplexer* (MUX). To control the select lines of the MUX, a bit-rate (or half-rate) clock must be synthesized from the slower word clock. This task is performed by a *clock multiplication unit* (CMU). Finally, a *laser driver* or *modulator driver* drives

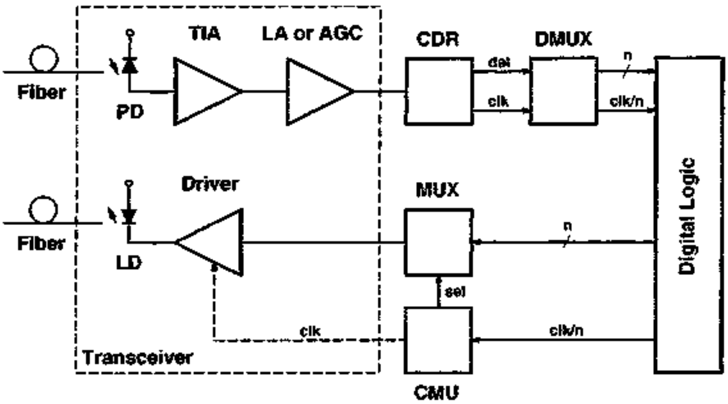

Fig. 1.1 Block diagram of an optical receiver (top) and transmitter (bottom).

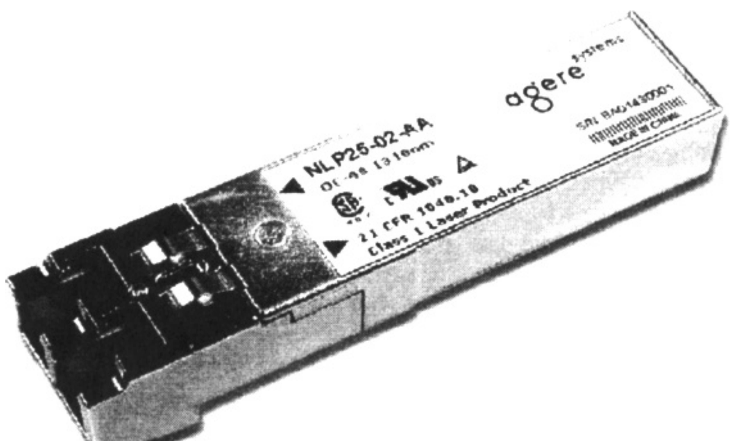

Fig. 1.2 A 2.5-Gb/s transceiver in a "small form-factor pluggable" (SFP) package (5.7 cm × 1.4 cm × 1.1 cm). The two fibers are plugged in from the left (LC connectors). Reprinted by permission from Agere Systems, Inc.

the corresponding optoelectronic device. The laser driver modulates the current of a *laser diode* (LD), whereas the modulator driver modulates the voltage across a *modulator*, which in turn modulates the light intensity from a *continuous wave* (CW) laser. Some laser/modulator drivers also perform data retiming, and thus require a bit-rate (or half-rate) clock from the CMU (dashed line in Fig. 1.1).

A module containing a PD, TIA, MA, laser driver, and LD, that is, all the blocks shown inside the dashed box of Fig. 1.1, often is referred to as a *transceiver*.[1] See Fig. 1.2 for a so-called *small form-factor* transceiver module. A module that contains all the functionality of the transceiver plus a CDR, DMUX, CMU, and MUX frequently is called a *transponder*. In this book, we discuss the PD, TIA, MA, as well as the laser/modulator and their drivers in greater detail.

Modulation Schemes. The most commonly used modulation format in optical communication is the *non-return-to-zero* (NRZ) format shown in Fig. 1.3(a). This format is a form of *on-off keying* (OOK): the signal *is on* to transmit a one bit and is *off* to transmit a zero bit. When the signal (i.e., the laser light) is on, it stays on for the entire bit period. For example, when transmitting the periodic bit pattern "010101010..." at 10 Gb/s in NRZ format, a 5-GHz square wave with 50% duty cycle is produced.[2]

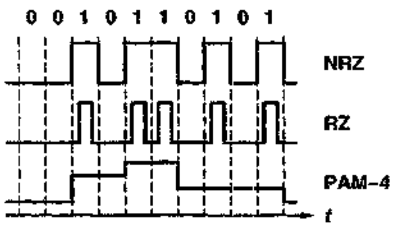

Fig. 1.3 Modulation schemes: (a) NRZ, (b) RZ, and (c) PAM-4.

In high-speed and long-haul transmission (e.g., fiber links between two continents), the *return-to-zero* (RZ) format, shown in Fig. 1.3(b), generally is preferred. In this format, the pulses, which represent the one bits, occupy only a fraction (e.g., 50%) of the bit period. Compared with the NRZ signal, the RZ signal requires less signal-to-noise ratio for reliable detection but occupies a larger bandwidth because of its shorter pulses.[3] An important advantage of this narrow-pulse format is that more pulse distortion and spreading can be tolerated without disturbing the adjacent bits.

[1] The term *transceiver* is a contraction of the words "*trans*mitter" and "re*ceiver*".

[2] In some standards, such as Fast Ethernet and FDDI, the *non-return-to-zero change-on-ones* (NRZ1 or NRZI) format is used. This format also is based on NRZ modulation, but before modulation, the bit stream is passed through a line coder that changes its (binary) output value when the bit to be transmitted is a one and leaves the output value unchanged when the bit is a zero.

[3] To receive data at a bit-error rate of 10^{-12}, we need a signal-to-noise ratio of 16.9 dB for NRZ modulation, 15.7 dB for 50%-RZ modulation, and 23.9 dB for PAM-4 modulation assuming additive Gaussian noise

Thus, this format is more immune to effects of fiber nonlinearity and polarization-mode dispersion. On the downside, faster, more expensive transceiver components (laser/modulator, photodetector, front-end electronics, etc.) are required to handle the shorter pulses. Furthermore, in optical multichannel systems, the wavelengths can be packed less densely because the RZ signal occupies a wider bandwidth than the NRZ signal for a given bit rate. Several variations of the RZ modulation, such as the *chirped return-to-zero* (CRZ) modulation, the *carrier-suppressed return-to-zero* (CS-RZ) modulation, and the *return-to-zero differential phase-shift keying* (RZ-DPSK) modulation also are used, but a discussion of these modulation schemes is beyond the scope of this book.

Since the late 1980s, the TV signals in *community-antenna television* (CATV) systems often are transported first *optically* from the distribution center to the neighborhood before they are distributed to the individual homes on conventional coaxial cable. This combination, called *hybrid fiber-coax* (HFC), has the advantage over an all-coax system in that it saves many electronic amplifiers (the loss in a fiber is much lower than the loss in a coax cable) and provides better signal quality (lower noise and distortions). In the optical part of the HFC system, the laser light is modulated with multiple radio-frequency (RF) carriers, so-called subcarriers, each one corresponding to a different TV channel. This method is known as *subcarrier multiplexing* (SCM).[4] Then, each subcarrier is modulated with a TV signal, for example, *amplitude modulation with vestigial sideband* (AM-VSB) is used for analog TV channels and *quadrature amplitude modulation* (QAM) is used for digital TV channels.

In contrast to NRZ and RZ modulation, which produce a two-level digital signal (laser light on or off), the AM-VSB and QAM modulation used in CATV applications produce continuous or multilevel analog signals. Figure 1.3(c) illustrates a multilevel signal produced by *pulse amplitude modulation* (PAM), a modulation scheme that is related to QAM. In this example, groups of two successive bits are encoded with one of four signal levels, and hence this format is known as PAM-4. Compared with the NRZ signal, the PAM-4 signal requires a higher signal-to-noise ratio for reliable detection but occupies a narrower bandwidth because its symbol rate is only half the bit rate. Similarly, the analog signals distributed over CATV/HFC systems require a high signal-to-noise ratio and transceivers with a good linearity to minimize distortions.

In the remainder of this book, we always assume that we are dealing with NRZ modulation, except if stated otherwise, as, for example, in the sections on analog receiver and transmitter circuits.

Line Codes. Before data bits are modulated onto the optical carrier, they usually are preconditioned with a so-called *line code*. The line code provides the transmitted bit stream with the following desirable properties: *DC balance*, short *run lengths*,

(cf. Problems 4.4 and 4.6). The signal bandwidth measured from DC to the first null is B for NRZ modulation, $2B$ for 50%-RZ modulation, and $B/2$ for PAM-4 modulation, where B is the bit rate.

[4] In contrast to *discrete multitone* (DMT) modulation (or *orthogonal frequency division multiplexing* [OFDM], the RF modulated equivalent), which uses overlapping channel spectra, SCM keeps a frequency gap between the channels.

and a high *transition density*. A DC balanced bit stream contains the same number of zeros and ones on average. This is equivalent to saying that the average *mark density* (number of one bits divided by all bits) is 50%. A DC balanced signal has the nice property that its average value (the DC component) is always centered halfway between the zero and one levels (half of the peak-to-peak value). This property often permits the use of AC coupling between circuit blocks, simplifying their design. Furthermore, it is desirable to keep the number of successive zeros and ones, the run length, to a small value. This provision reduces the low-frequency content of the transmitted signal and limits the associated *baseline wander* (a.k.a. *DC wander*) when AC coupling is used. Finally, a high transition density is desirable to simplify the clock recovery process.

In practice, line coding is implemented as either *scrambling*, *block coding*, or a combination of the two:

- *Scrambling*. In this case, a *pseudorandom bit sequence* (PRBS) is generated with a feedback shift register and xor'ed with the data bit stream (see Fig. 1.4). Note that the data can be descrambled with the same arrangement, provided the descrambling PRBS generator is synchronized with the scrambling PRBS generator. Scrambling provides DC balance without adding overhead bits to the bit stream, thus preserving the bit rate. On the down side, the maximum run length is not strictly limited, that is, there is a small chance for very long runs of zeros or ones, which can be hazardous. In practice, runs up to 72 bits usually are expected. The scrambling method is used in the United States telecommunication system described in the SONET (synchronous optical network) standard [188] and the almost identical SDH (synchronous digital hierarchy) standard used in Europe and Japan.

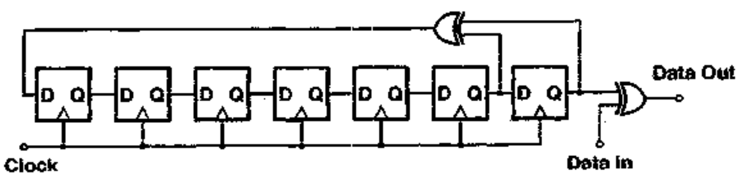

Fig. 1.4 Implementation of a SONET scrambler.

- *Block Coding*. In this case, a contiguous group of bits (a block) is replaced by another slightly larger group of bits such that the average mark density becomes 50% and DC balance is established. For example, in the 8B10B code, 8-bit groups are replaced with 10-bit patterns using a look-up table [198]. The 8B10B code increases the bit rate by 25%; however, the maximum run length is strictly limited to five zeros or ones in a row. The 8B10B code is used in the

Gigabit Ethernet (GbE, 1000Base-SX, 1000Base-LX) and Fiber Channel data communication systems.[5]

- *Combination.* In the serial 10-Gigabit Ethernet (10-GbE) system, DC balance is established first by scrambling the bit stream and then by applying a block-code (64B66B code) to it. This combination features low overhead ($\approx$ 3% increase in bit rate) and a run length that is strictly limited to 66 bits.

Continuous Mode vs. Burst Mode. It is important to distinguish two types of transmission modes because they call for different circuit designs: *continuous mode* and *burst mode*. The signals corresponding to these two modes are shown schematically in Fig. 1.5.

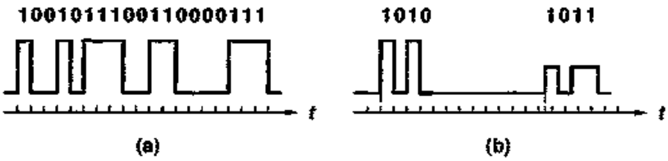

Fig. 1.5 (a) Continuous-mode vs. (b) burst-mode signals (schematically).

In continuous-mode transmission, a continuous, uninterrupted stream of bits is transmitted as shown in Fig. 1.5(a). The transmitted signal usually is DC balanced using one of the line codes described earlier. As a result, AC coupled circuits normally can be used. In burst-mode transmission, data are transmitted in short *bursts*, with the transmitter remaining silent (laser off) in between bursts. See Fig. 1.5(b) for an illustration, but note that practical bursts are much longer than those shown in the figure, typically longer than 400 bits.

Bursts can be fixed or can be variable in length. Bursts that encode ATM (asynchronous transfer mode) cells have a *fixed length*, they always contain 53 bytes plus a preamble (e.g., 3 bytes). Bursts that encode Ethernet frames have a *variable length* (70–1524 bytes). In either case, the bursts start out with a preamble (a.k.a. overhead) followed by the payload. The burst-mode receiver uses the preamble to establish the decision threshold level (slice level) and to synchronize the receiver clock. In passive optical network systems, to be discussed shortly, bursts arrive *asynchronously* and with strongly *varying power levels* (up to 30 dB); therefore, the clock signal must be synchronized and the slice level adjusted for every single burst (cf. Fig. 1.5(b)).

The average value (DC component) of a burst-mode signal varies with time, depending on the burst activity. If the activity is high, it may be close to the halfway point between the zero and one levels, as in a continuous mode system; if the activity is low, the average drifts arbitrarily close to the zero level. This means that the burst-mode signal is *not* DC balanced, and in general, AC coupling cannot be used because

[5] The 4B5B code used in Fast Ethernet (100Base-TX, 100Base-FX), FDDI, and so forth, does not achieve perfect DC balance; the worst-case unbalance is 10% [136].

it would lead to excessive baseline wander. (Note that the mark density *within a burst* may well be 50%, but the overall signal is still not DC balanced.) This lack of DC balance and the fact that bursts often arrive with varying amplitudes necessitate specialized amplifier and driver circuits for burst-mode applications. Furthermore, the asynchronous arrival of the bursts requires specialized fast-locking CDRs.

In the remainder of this book we always assume that we are dealing with DC-balanced, continuous-mode signals, except if stated otherwise, as, for example, in the sections on burst-mode circuits.

Optical Networks. We must distinguish two important types of optical networks: the simple *point-to-point connection* and the *point-to-multipoint network*. In the following, we discuss these networks and how continuous-mode and burst-mode transmissions are used with them.

An optical point-to-point connection between two *central offices* (CO) is illustrated schematically in Fig. 1.6(a). An example for such a connection is a SONET OC-192 link operating at 10 Gb/s (9.953 28 Gb/s to be precise), a bit rate that can carry about 130,000 voice calls. Point-to-point links are used over a wide range of bit rates and distances, from short computer-to-computer links to ultra-long-haul undersea lightwave systems.

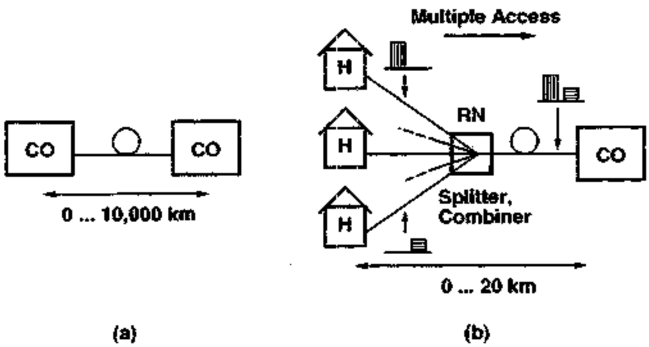

Fig. 1.6 Example of (a) a point-to-point link and (b) a point-to-multipoint network.

Point-to-point connections can be assembled into more complex structures such as *ring networks* and *active star networks*. Examples for ring networks are provided by SONET/SDH rings and FDDI token rings. An active star is formed, for example, by Gigabit Ethernet links converging into a hub. It is important to realize that each individual optical connection of the star has a transceiver on both ends and therefore forms an optical point-to-point link. This is in contrast to a *passive star network* or an optical point-to-multipoint network, where multiple optical fibers are coupled with a passive optical device. We discuss the latter network type below.

Continuous-mode transmission is used on almost all point-to-point connections. One exception occurs in half-duplex systems, in which bidirectional communication is implemented by periodically reversing the direction of traffic following a pingpong pattern, so-called *time compression multiplexing* (TCM; a.k.a. *time division duplexing*). Such systems require burst-mode transmitters and receivers. However, for bandwidth efficiency reasons, TCM systems are limited to relatively short links (e.g., home networking applications) and are not widely used. In all other cases of bidirectional transmission, for example, with two fibers, so-called *space division multiplexing* (SDM), or two wavelengths, so-called *wavelength division multiplexing* (WDM), continuous-mode transmission is used.

Passive Optical Network. A *passive optical network* (PON) is illustrated schematically in Figure 1.6(b). A feeder fiber from the central office (CO) runs to a *remote node* (RN), which houses a passive optical power splitter/combiner. From there, around 32 fibers branch out to the subscribers. If these fibers extend all the way to the *homes* (H), as shown in Fig 1.6(b), this system is known as a *fiber-to-the-home* (FTTH) system. Alternatively, if the fibers terminate at the curb, the system is known as a *fiber-to-the-curb* (FTTC) system. The final distribution from the curb to the homes is accomplished, for example, by twisted-pair copper wires or radio. All systems that bring the fiber relatively close to the subscriber are collectively known as FTTx systems.

In a traditional telephony access network, the connection between the CO and the remote node is a digital, possibly optical line. The final distribution from the remote node to the subscribers, however, is accomplished with analog signals over twisted-pair copper wires. Thus, the remote node must be *active*; that is, it needs to be powered to perform the conversion from the high-speed digital signal to the analog signals. In contrast, a PON system is all optical and *passive*. Because a PON does not require outside power supplies, it is low in cost, easy to maintain, and reliable.

A PON is a point-to-multipoint network because the optical medium is *shared* among the subscribers. Information transmitted downstream, from the CO to the subscriber, is received by all subscribers, and information transmitted upstream, from the subscribers to the CO, is superimposed at the passive combiner before it is received at the CO. To avoid data collisions in the *upstream direction*, the subscriber data must be buffered and transmitted in short bursts. The CO must coordinate which subscriber can send a burst at which point in time. This method is known as *time division multiple access* (TDMA) and requires *burst-mode transmission*. The *downstream direction* is more straightforward: the CO tags the data with addresses and broadcasts it to all subscribers in sequential order. Each subscriber simply selects the information with the appropriate address tag. This method is known as *time division multiplexing* (TDM), and conventional continuous-mode transmission can be used. Upstream and downstream transmissions usually are separated by using two different wavelengths (WDM bidirectional transmission).

The most promising PON systems are (i) BPON (broadband passive optical network), which carries the data in ATM cells and hence also is known as ATM-PON [30, 52], and (ii) EPON (Ethernet passive optical network), which carries the data in Ethernet frames, as the name implies [48]. In general, PON FTTx networks are limited to relatively small distances (<20 km) and currently are operated at modest bit rates (50 Mb/s–1.25 Gb/s). In a typical BPON FTTH scenario, 16 to 32 homes share a bit rate of 155 Mb/s, giving each subscriber an average speed of 5 to 10 Mb/s. This is sufficient for fast Internet access, telephone service, and video on demand. Sometimes, an all-optical CATV service is provided over the PON infrastructure by means of a third wavelength.

Besides the TDM/TDMA approach outlined above, there are several other types of PON systems. For example, the WDM-PON system, where a different wavelength is assigned to each subscriber, has been studied extensively. In WDM-PON systems, data collisions are avoided without the need for burst-mode transmission. However, the optical WDM components required for such a system currently are too expensive, making WDM-PON uneconomical. For more information on PON systems see [40, 69].

Book Outline. In Chapter 2, we introduce the communication channel presented by the optical fiber. Its loss, bandwidth, and various forms of dispersion are described briefly. The relationship between pulse spreading and transmitter linewidth is given and is used later in Chapter 7.

Chapter 3 is the first of four chapters dealing with the receiver. We start by studying the responsivity and noise properties of three photodetector types: the p-i-n photodetector, the avalanche photodetector, and the optically preamplified p-i-n detector. In Chapter 4, we present the receiver at the system level. This chapter introduces terminology and concepts that simplify the discussion in later chapters. First, we analyze how noise in the receiver causes bit errors. This leads to the definition of the receiver sensitivity. Next, after introducing the concept of power penalty, we study the impact of the receiver's bandwidth and frequency response on its performance. The adaptive equalizer, used to mitigate distortions in the received signal, is covered briefly. We then turn to other receiver impairments, such as nonlinearity (in analog receivers), jitter, decision threshold offset, and sampling time offset. We conclude with a brief description of forward error correction, a technique that can improve the receiver's performance dramatically. In Chapter 5, we discuss the transimpedance amplifier. We start by introducing the main specifications. Next, we discuss circuit concepts in a general, and as much as possible, technology independent manner. This includes the shunt-feedback architecture and variations thereof, noise optimization procedures, and special techniques for burst-mode and analog TIAs. Then, we illustrate these concepts with practical implementations in a broad range of technologies. We conclude with a brief overview of product examples and current research topics. In Chapter 6, we discuss the main amplifier, that is, the limiting and AGC amplifier. As in the previous chapter, we proceed from specifications to circuit concepts to implementation examples. The circuit concepts covered include the multistage architecture, techniques for broadband stages, offset compensation, and automatic gain control.

Chapter 7 is the first of two chapters dealing with the transmitter. We start by studying various types of lasers and optical modulators. In Chapter 8, we discuss the driver circuits for directly modulated lasers as well as external modulators. As in Chapters 5 and 6, we proceed from specifications to circuit concepts to implementation examples. The circuit concepts covered include the current-steering output stage with and without back termination, the predriver with pulse-width control, data retiming, automatic power control, and special techniques for burst-mode and analog drivers.

Appendices A through D provide additional material on eye diagrams, differential circuits, S parameters, transistors, and technologies. Answers to all problems are given in Appendix E. The notation for currents, voltages, and noise quantities used in this book is explained in Appendix F. For reference, Appendix G defines all symbols and Appendix H defines all acronyms used in this book.

About Numerical Examples. In the following chapters, we make extensive use of numerical examples. Frequently when we introduce a new quantity or relationship, we illustrate it with so-called *typical values*. In my own learning experience, this approach is most helpful: it makes the subject more concrete and promotes a feeling for the numerical values. However, specialists tend to be quite critical about such values because they are never quite right. Typical values may change over time as the technology advances or they may depend on several conditions that may or may not be met in a particular case. It therefore is important to take the subsequent typical values only as an illustration and not as the basis for your next design project!

Further Reading. The book *High-Speed Circuits for Lightwave Communications* edited by K. Wang contains an interesting collection of papers discussing many insightful case studies [197]. The textbook *Design of Integrated Circuits for Optical Communications* by B. Razavi presents an overview of circuits for optical communication, including phase-locked loops and clock and data recovery circuits [141]. The latest developments in integrated circuits for optical communication are published in several magazines and at various conferences, for example, the *IEEE Journal of Solid-State Circuits* and the *IEEE International Solid-State Circuits Conference* are excellent sources. Information about optical fibers, lasers, detectors, optical amplifiers, as well as optical networks can be found in the textbooks *Fiber-Optic Communication Systems* by G. Agrawal [5] and *Optical Networks: A Practical Perspective* by R. Ramaswami et al. [136]. The continually updated book series *Optical Fiber Telecommunications* (so far Volumes I–IV have been published) covers the latest developments in optical/optoelectronic components and systems [58, 59, 60, 61, 93]. Finally, I recommend *City of Light: The Story of Fiber Optics* by Jeff Hecht for an entertaining and informative historical account on fiber optics [43].