Chapter 1

Introduction to Optical Communications

The rapidly-growing volumes of data in telecommunication networks have rekindled interest in high-speed optical and electronic devices and systems. With the proliferation of the Internet and the rise in the speed of microprocessors and memories, the transport of data continues to be the bottleneck, motivating work on faster communication channels.

The idea of using light as a carrier for signals has been around for more than a century, but it was not until the mid-1950s that researchers demonstrated the utility of the optical fiber as a medium for light propagation [1]. Even though early fibers suffered from a high loss, the prospect of guided transmission of light with a very wide modulation band ignited extensive research in the area of optical communications, leading to the practical realization of optical networks in the 1970s.

This chapter provides an overview of optical communications, helping the reader understand how the concepts introduced in subsequent chapters fit into the "big picture." We begin with a brief history and study a generic optical system, describing its principal functions. Next, we present the challenges in the design of modern optical transceivers. Finally, we review the state of the art and the trends in transceiver design.

1.1 Brief History

Attempts to "guide" light go back to the 1840s, when a French physicist named Jacque Babinet demonstrated that light could be "bent" along a jet of water. By the late 1800s, researchers had discovered that light could travel inside bent rods made of quartz. The "fiber" was thus born as a flexible, transparent rod of glass or plastic.

In 1954, Abraham van Heel of the Technical University of Delft (Holland) and Harold Hopkins and Narinder Kapany of the Imperial College (Britain) independently published the idea of using a bundle of fibers to transmit images. Around the same time, Brian O'Brien of the American Optical Company recognized that "bare" fibers lost energy to the surrounding air, motivating van Heel to enclose the fiber core in a coating and hence lower the loss. Fiber loss was still very high, about 1,000 dB/km, limiting the usage to endoscopy applications.

The introduction of the laser as an intense light source in the 1950s and 1960s played a crucial role in fiber optics. The broadband modulation capability of lasers offered great potential for carrying information, although no suitable propagation medium seemed available. In 1966, Charles Ko and Charles Hockem of the Standard Telecommunication Laboratory (Britain) proposed that the optical fiber could be utilized as a signal transmission medium if the loss was lowered to 20 dB/km. They also postulated that such a low loss would be obtained if the impurities in the fiber material were reduced substantially.

Four years later, Robert Mauer and two of his colleagues at Corning Glass Works demonstrated silica fibers having a loss of less than 20 dB/km. With advances in semiconductor industry, the art of reducing impurities and dislocations in fibers improved as well, leading to a loss of 4 dB/km in 1975 and 0.2 dB/km in 1979. The dream of carrying massive volumes of information over long distances was thus fulfilled: in 1977, AT&T and GTE deployed the first fiber optic telephone system.

The widespread usage of optical communication for the transport of high-speed data stems from (1) the large bandwidth of fibers (roughly 25 to 50 GHz) and (2) the low loss of fibers (0.15 to 0.2 dB/km). By comparison, the loss reaches 200 dB/km at 100 MHz for twisted-pair cables and 500 dB/km at 1 GHz for low-cost coaxial cables. Also, wireless propagation with carrier frequencies of several gigahertz incurs an attenuation of tens of decibels across a few meters while supporting data rates lower than 100 Mb/s.

The large (and free) bandwidth provided by fibers has led to another important development: the use of multiple wavelengths (frequencies) to carry several channels on a single fiber. For example, it has been demonstrated that 100 wavelengths, each carrying data at 10 Gb/s, allow communication at an overall rate of 1 Tb/s across 400 km.

1.2 Generic Optical System

The goal of an optical communication (OC) system is to carry large volumes of data across a long distance. For example, the telephone traffic in Europe is connected to that in the United States through a fiber system installed across the Atlantic Ocean.

Depicted in Fig. 1.1(a), a simple OC system consists of three components: (1) an electro-optical transducer (e.g., a laser diode), which converts the electrical data to optical form (i.e., it produces light for logical ONEs and remains off for logical ZEROs); (2) a fiber, which carries the light produced by the laser; and (3) a photodetector (e.g., a photodiode), which senses the light at the end of the fiber and converts it to an electrical signal. We call the transmit and receive sides the "near end" and the "far end," respectively. As explained in Chapter 3, lasers are driven by electrical currents, and photodiodes generate an output current.

With long or low-cost fibers, the light experiences considerable attenuation as it travels from the near end to the far end. Thus, (1) the laser must produce a high light intensity, e.g., tens of milliwatts; (2) the photodiode must exhibit a high sensitivity to light; and (3) the electrical signal generated by the photodiode must be amplified with low noise. These observations lead to the more complete system shown in Fig. 1.1(b), where a "laser driver" delivers large currents to the laser and a "transimpedance amplifier" (TIA) amplifies the photodiode output with low noise and sufficient bandwidth, converting it to a voltage. For

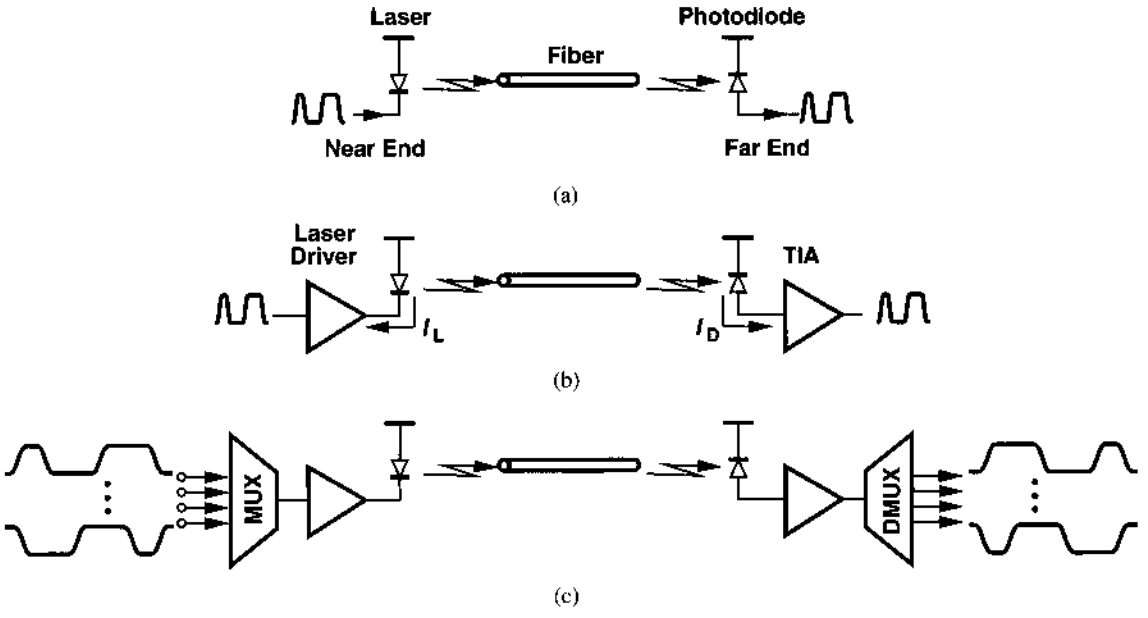

Figure 1.1 (a) Simple optical system, (b) addition of driver and amplifier, (c) addition of MUX and DMUX.

example, data at a rate of 10 Gb/s may be applied to the laser driver, modulate the laser light at a wavelength of 1.55 μm, and emerge at the output of the TIA with an amplitude of 10 mV.

The transmit and receive operations in Fig. 1.1(b) process high-speed "serial" data, e.g., a single stream of data at 10 Gb/s. However, the actual data provided to the transmitter (TX) is in the form of many low-speed channels ("parallel" data) because it is generated by multiple users. The task of parallel-to-serial conversion is performed by a "multiplexer" (MUX). Similarly, the receiver (RX) must incorporate a "demultiplexer" (DMUX) to reproduce the original parallel channels. The resulting system is shown in Fig. 1.1(c).

The topology of Fig. 1.1(c) is still incomplete. Let us first consider the transmit end. The multiplexer requires a number of clock frequencies with precise edge alignment. These clocks are generated by a phase-locked loop (PLL). Furthermore, in practice, the MUX output suffers from nonidealities such as "jitter" and "intersymbol interference" (ISI), mandating the use of a "clean-up" flipflop before the laser driver. These modifications lead to the transmitter illustrated in Fig. 1.2(a).

The receive end also requires additional functions. Since the TIA output swing may not be large enough to provide logical levels, a high-gain amplifier (called a "limiting amplifier") must follow the TIA. Moreover, since the received data may exhibit substantial noise, a clean-up flipflop (called a "decision circuit") is interposed between the limiting amplifier and the DMUX. The receiver thus appears as shown in Fig. 1.2(b).

The receiver of Fig. 1.2(b) lacks a means of generating the clock necessary for the de-

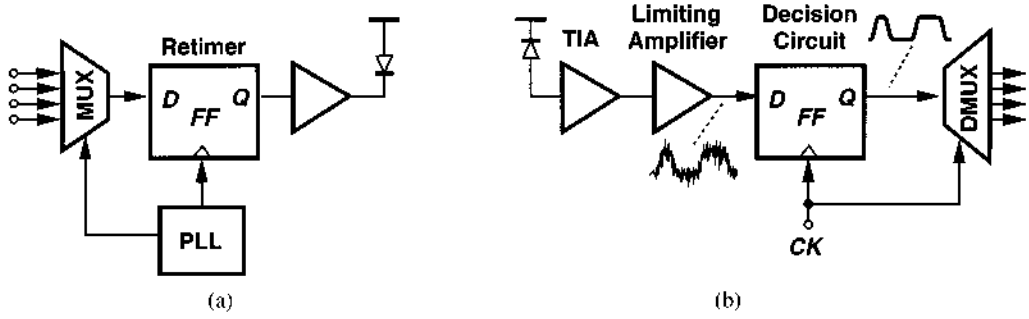

Figure 1.2 Modified (a) transmitter and (b) receiver.

cision circuit and the DMUX. This clock must bear a well-defined phase relationship with respect to the received data so that the flipflop samples the high and low levels "optimally," i.e., at the midpoint of each bit. The task of generating such a clock from the incoming data is called "clock recovery." The overall operation of clock recovery and data cleanup is called "clock and data recovery" (CDR). Figure 1.3 shows the complete system. Note that the laser driver incorporates power control (Chapter 10) and the TIA employs automatic gain control (AGC) (Chapter 4).

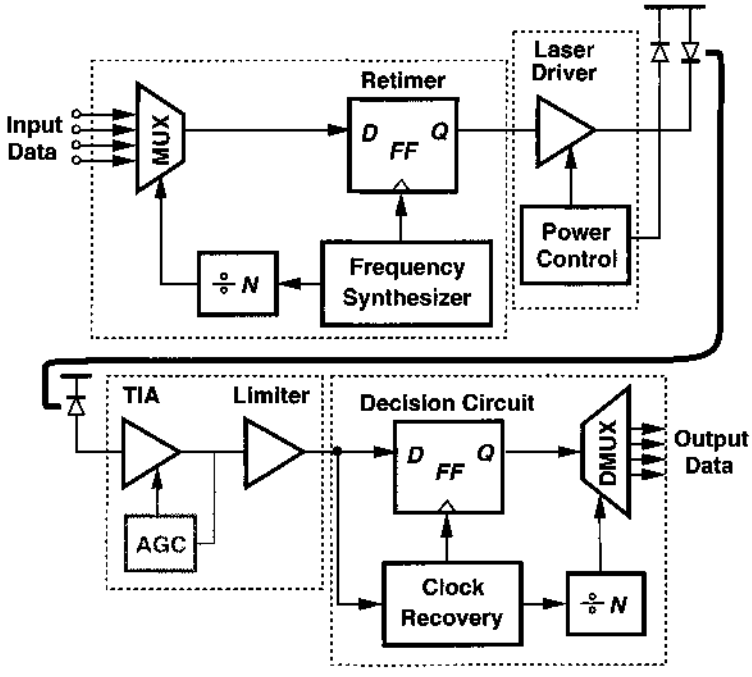

Figure 1.3 Complete system.

1.3 Design Challenges

While the system topology of Fig. 1.3 has not changed much over the past several decades, the design of its building blocks and the levels of integration have. Motivated by the evolution and affordability of IC technologies as well as the demand for higher performance, this change has created new challenges, necessitating new circuit and architecture techniques. We review some of the challenges here.

The transmitter of Fig. 1.3 entails several issues that manifest themselves at high speeds and/or in scaled IC technologies. Since the jitter of the transmitted data is determined primarily by that of the PLL, a robust, low-noise design with high supply and substrate rejection becomes essential. Furthermore, the design of skew-free multiplexers proves difficult at high data rates.

Another critical challenge arises from the laser driver, a circuit that must deliver tens of milliamperes of current with very short rise and fall times. Since laser diodes may experience large voltage swings between on and off states, the driver design becomes more difficult as scaled technologies impose lower supply voltages. The package parasitics also severely limit the speed with which such high currents can be switched to the laser [2].

The optical components in Fig. 1.3, namely, the laser diode, the fiber, and the photodiode, introduce their own nonidealities, requiring close interaction between electronic and optical design. Effects such as chirp, dispersion, attenuation, and efficiency play a major role in the overall link budget.

The receiver of Fig. 1.3 also presents many problems. The noise, gain, and bandwidth of the TIA and the limiter directly impact both the sensitivity and the speed of the overall system, raising additional issues as the supply voltage scales down. Moreover, the clock and data recovery functions must provide a high speed, tolerate long runs (sequences of identical bits), and satisfy stringent jitter and bandwidth requirements.

Full integration of the transceiver shown in Fig. 1.3 on a single chip raises a number of concerns. The high-speed digital signals in the MUX and DMUX may corrupt the receiver input or the oscillators used in the PLL and the CDR circuit. The high slew rates produced by the laser driver may lead to similar corruptions and also desensitize the TIA. Finally, since the oscillators in the transmit PLL and the receive CDR circuit operate at slightly different frequencies (with the difference given by the mismatch between the crystal frequencies in two communicating transceivers), they may "pull" each other, generating substantial jitter.

The above issues have resulted in multichip solutions that integrate the noisy and sensitive functions on different substrates. The dashed boxes in Fig. 1.3 indicate a typical partitioning, suggesting the following single-chip blocks: the PLL/MUX circuit (also called the "serializer"), the laser driver along with its power control circuitry, the TIA/limiter combination, and the CDR/MUX circuit (also called the "deserializer"). Recent work has integrated the serializer and deserializer (producing a "SERDES") but the TX and RX amplifiers may remain in isolation.

1.4 State of the Art

The new optical revolution is reminiscent of the monumental change that radio-frequency (RF) design began to experience in the early 1990s. This resurgence entails three important trends: (1) Modular, general-purpose building blocks are gradually replaced by end-to-end solutions that benefit from device/circuit/architecture codesign. (2) Greater levels of integration on a single chip provide higher performance and lower cost. (3) Mainstream VLSI technologies such as CMOS and BiCMOS continue to take over the territories thus far claimed by GaAs and InP devices. Modern OC transceiver applications continue to challenge designers along many dimensions.

Realization in CMOS Technology The cost and integration advantages of CMOS technology have motivated extensive work on high-speed CMOS design. Issues such as noise, speed, voltage headroom, and substrate coupling pose many difficulties in the design of CMOS transceivers. Research on 10-Gb/s CMOS CDR circuits yielded results in 2000 [3], and CMOS serializers and deserializers operating at this rate were reported in 2002 [4, 5].

Speed With the increasing volume of data transported in the backplane of the Internet, optical communication at rates of 40 Gb/s has become attractive. Such high speeds emerge as a new territory for IC design because prior work at these frequencies ("millimeter-wave frequencies") has been limited to narrowband, low-complexity circuits for wireless applications. Pushing bipolar and, preferably, CMOS technologies to such speeds, designers must cope with broadband characterization of active and passive devices, transmission-line behavior of on-chip interconnects, and high-speed packaging issues. A 40-Gb/s SiGe CDR circuit has been reported in [6].

Level of Integration Integrating a complete SERDES on a single CMOS chip serves as the first step toward much greater sophistication in OC transceiver design. Two important trends particularly suited to CMOS technology are: (1) integration of the SERDES along with the large digital processor that interfaces with the network (the "framer"); such integration eliminates a large number of high-speed printed-circuit board (PCB) lines between the two, simplifying the package design and saving substantial power. (2) integration of multiple SERDES on one chip; since the total data rate can be increased through the use of multiple light *wavelengths* on a single fiber, an important thrust is to integrate several SERDES on the same substrate, thereby increasing the "port density."

Power Dissipation At high speeds and/or high port densities, the power dissipation of optical transceivers becomes critical as it determines the type and size of the package in which the entire module is housed. Today's 10-Gb/s SERDES consume about 1 W of power, leading to serious packaging issues if four must be integrated on one chip. Interestingly, the low supply voltage required for deep-submicron CMOS technologies does reduce the overall power dissipation (e.g., in the output buffers) while making circuit design more difficult.

References

1. D. G. Goff, *Fiber Optic Reference Guide,* Boston: Focal Press, 1999.
2. H.-M. Rein and M. Moller, "Design Considerations for Very High Speed Si Bipolar ICs Operating up to 50 Gb/s," *IEEE Journal of Solid-State Circuits,* vol. 31, pp. 1076–1090, August 1996.
3. J. Savoj and B. Razavi, "A 10-Gb/s CMOS Clock and Data Recovery Circuit," *Symp. on VLSI Circuits Dig. of Tech. Papers,* pp. 136–139, June 2000.
4. M. M. Green et al., "OC-192 Transmitter in Standard 0.18-μm CMOS," *ISSCC Dig. of Tech. Papers,* pp. 186–187, Feb. 2002.
5. J. Cao et al., "OC-192 Receiver in Standard 0.18-μm CMOS," *ISSCC Dig. of Tech. Papers,* pp. 187–188, Feb. 2002.
6. M. Reinhold et al., "A Fully Integrated 40-Gb/s Clock and Data Recovery IC with 1:4 DMUX in SiGe Technology," *IEEE Journal of Solid-State Circuits,* vol. 36, pp. 1937–1945, Dec. 2001.