1

Introduction

This chapter introduces the optical network. We begin with a survey of three generations of digital transport networks, followed by a discussion of the extraordinary capacity of optical fiber. The optical network marketplace is examined with a look at current and projected installations. Next, we examine the key nodes (machines) that make up the optical network, then we look inside a node to learn about its components. The chapter concludes with a general explanation of the attributes of optical fiber.

THREE GENERATIONS OF DIGITAL TRANSPORT NETWORKS

The focus of this book is on third generation digital transport networks, usually shorted to 3G, or 3rd generation, transport networks. The main characteristics of three generations of digital transport networks are provided in Table 1–1. The information in this table will be helpful as you read the remaining chapters in this book. Most of the terms in the table are self-explanatory, or, if not, are explained in this chapter.

The first column in the table is the name (or names) usually associated with the technology. The first generation systems are known as T1 or E1. The second generation systems are called SONET (for the Synchronous Optical Network) or SDH (for the Synchronous Digital Hierarchy). These terms are explained in more detail in later parts of this book. However, the

Table 1-1 Three Generations of Digital Transport (Carrier) Networks

Name	Family	Designed for	MUX/SW Schemes at Inception	Principal Media at Inception	Capacity	Typical Payload	Protocol Inter-Working?
T1/E1	First	Voice, Non-BOD, Static	TDM/E/E/E	Copper: (Early 1960s)	Mbit/s	Fixed Length	No
SONET/SDH	Second	Voice, Non-BOD, Static	TDM/O/E/O	Copper, Fiber: (Mid–1980s)	Gbit/s	Fixed Length	Somewhat: PPP, IP, ATM
OTN	Third	Voice, Video, Data, Tailored QOS, BOD, Dynamic	WDM/O/O/O	Fiber (Late 1990s to Early 2000s)	Tbit/s	Fixed or Variable Lengths	Yes: PPP, IP, ATM, MPLS

industry has not yet settled on a handle for the third generation digital carrier network, but the term Optical Transport Network (OTN) is widely used. The second column identifies the generation family.

The third column shows what kinds of user payloads the networks are designed to support. Although the first and second generation networks are designed to support voice traffic, they can and do transport data and video images. But they are not "optimized" for data and video traffic. In contrast, the 3G transport network is designed to support voice, video, or data payloads. When used with multiprotocol label switching (MPLS), the resource reservation protocol (RSVP), and Diff-Serv, as well as some of the new specifications dealing with optical bandwidth on demand, they are also designed to provide tailored quality-of-service (QOS) features for individual customers. The point will be made repeatedly in this book that the 3G transport network no longer consists of fixed, static "pipes" of capacity; it can dynamically change to meet the changing requirements of its users.

The third column also contains the notations of Non-BOD or BOD. The first and second generation systems are not designed to provide bandwidth of demand (BOD). The bandwidth is configured with crafting operations at each node. 3G systems are more dynamic and allow bandwidth to be requested on demand.

The fourth column lists the predominant multiplexing schemes: TDM or WDM. The fourth column also lists the manner in which the networks switch traffic when they were first deployed (at their inception). First generation systems were solely E/E/E operations: (a) they accepted electrical signals (the first E), (b) processed them (the second E), and (c) sent them to another node (the third E). Second generation systems are O/E/O operations: (a) they accept optical signals (the first O), (b) convert them to electrical signals for processing (the E), and (c) convert the electrical signals back to optical signals for transmission (the second O). Third generation systems are intended to be all optical (O/O/O), in that they process optical payloads, and do not need to convert the bits to electrical images for processing. Today, all three generations are mainly O/E/O oriented.

The fifth column lists the principal media used by the technologies at their inception, as well as the time that these networks were first introduced into the industry. All three generations now use a combination of copper, fiber, and wireless media.

The sixth column lists the typical capacity of the generation. It is evident that each succeeding family has increased its transport capacity by orders of magnitude.

The seventh column goes hand-in-hand with the third column ("Designed For"). The first and second generation networks were designed for fixed-length voice traffic, based on the 64 kbit/s payload, with a 125-µsec clocking increment. The third generation network supports this signal, but also supports variable-length payloads, an important capability for carrying data traffic. As well, the first and second generation networks can carry variable-length traffic, but they are not very efficient in how they go about transporting variable-length data traffic.

The eighth column explains whether any of the generations were designed to interwork with and directly support other protocols. T1/E1 was not so designed; again, 1st generation transport systems were set up to support voice traffic. Any efforts to devise methods of carrying other payloads were an afterthought and in vendor-specific procedures. With the advent of 2nd generation systems with SONET/SDH, efforts were made by the standards groups to define procedures for carrying certain kinds of data traffic, and many manufacturers adapted these standards into their products.

3rd generation transport networks are geared toward supporting many kinds of payloads, and specifically the Internet, ATM, and MPLS protocol suites. As we shall see as we proceed though this book, extensive research has resulted in many specifications defining how MPLS contributes to the operations of the third generation digital (optical) transport network.

All Features Are Not Yet Available

Not all the features and attributes cited in Table 1–1 are available in 3G transport networks. In fact, third generation transport networks are just now appearing in the marketplace, and some capabilities that are touted for them are still in the lab. Nonetheless, many people think full-featured 3rd generation transport networks will be in the marketplace by around 2004. Certainly, pieces are emerging, such as bandwidth on demand, and of course, WDM and terabit networks. Other parts of 3G transport networks have yet to be implemented. For example, O/O/O operations are far from reaching commercial deployment on a mass scale.

Optical Fiber Capacity

To gain an appreciation of the transmission capacity of optical systems operating today, consider the facts in Table 1–2. Prior to the advent of optical fiber systems, a high-capacity network was capable

Table 1–2 Magnitudes and Meanings

Magnitude	Term	Initial	Meaning
1 000 000 000 000 000 000 = 10^{18}	exa	E	Quintillion
1 000 000 000 000 000 = 10^{15}	peta	P	Quadrillion
1 000 000 000 000 = 10^{12}	tera	T	Trillion
1 000 000 000 = 10^{9}	giga	G	Billion
1 000 000 = 10^{6}	mega	M	Million
1 000 = 10^{3}	kilo	k	Thousand
100 = 10^{2}	hecto	h	Hundred
10 = 10^{1}	deka	da	Ten

of operating (sending and receiving traffic) at several million bits per second (Mbit/s). These electrical/electromagnetic transmissions take place over some form of metallic medium such as copper wire or coaxial cable, or over wireless systems such as microwave. In contrast, optical fiber systems transmit light signals through a glass or plastic medium. These systems are many orders of magnitude "faster" than their predecessors, with the capability of operating in the terabits-per-second (Tbit/s) range.

As depicted in Figure 1–1, a terabit fiber carries 10^{12} bits per second. At this rate, the fiber can transport just over 35 million data connections at 28.8 kbit/s, or about 17 million digital voice channels, or just under 500,000 compressed TV channels (or combinations of these channels).

Approximately:
- **35 million data connections at 28 kbit/s or**
- **17 million digital voice telephony channels or**
- **1/2 million compressed TV channels**

Figure 1–1 Capacity of one fiber with a 1 Tbit/s rate.

Even the seasoned telecommunications professional pauses when thinking about the extraordinary capacity of optical fiber.

A logical question for a newcomer to optical networks is, why are they of much greater capacity than, say, a network built on copper wire, or coaxial cable? The answer is that optical signals used in optical networks operate in a very high position and range of the frequency spectrum, many orders of magnitude higher than electromagnetic signals. Thus, the use of the higher frequencies permits the sending of many more user payloads (voice, video, and data) onto the fiber medium.

Figure 1–2 shows the progress made in the transmission capacity of optical fiber technology since 1982 [CHRA99]. The top line represents experimental systems, and the bottom line represents commercial systems. The commercial results have lagged behind the experimental results by about six years. The dramatic growth in the experimental capacity was due to improved laboratory techniques and the progress made in dispersion management, a subject discussed later in this book. As the figure

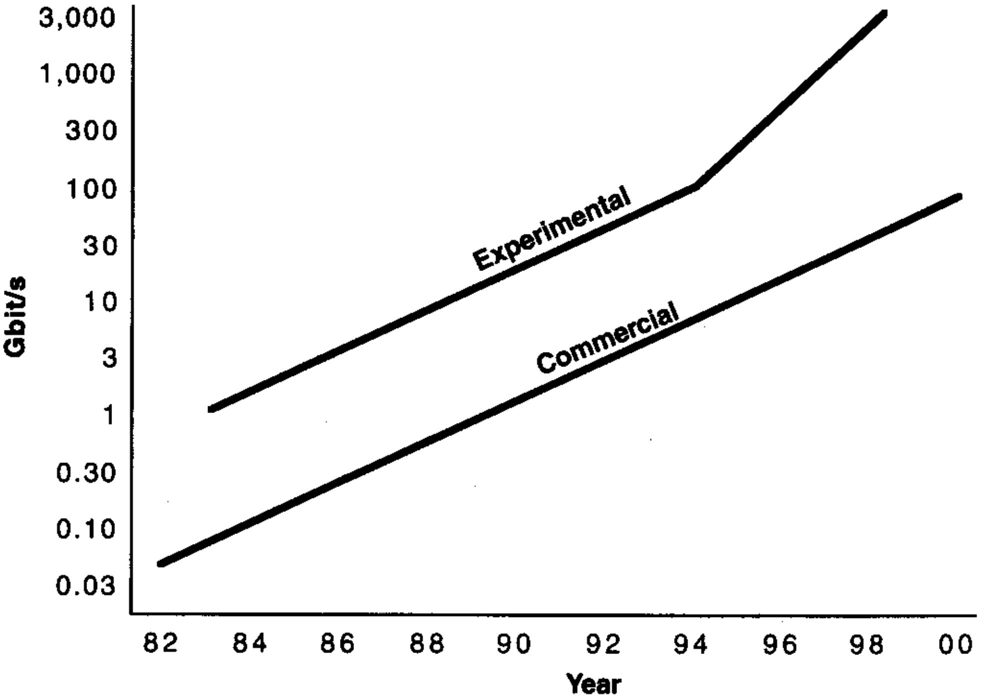

Figure 1–2 Transmission capacity as a function of year [CHRA99].

shows, the transmission capacity of optical fiber has been growing at an extraordinary rate since the inception of the technology.

A BRIEF INTRODUCTION TO WDM AND TDM

We would like to keep this introductory chapter free from technical detail as much as possible. However, it is necessary to introduce two terms before proceeding further: (a) wave division multiplexing (WDM), and (b) time division multiplexing (TDM). Later chapters will embellish on this introduction. Refer to Figure 1–3 during this discussion.

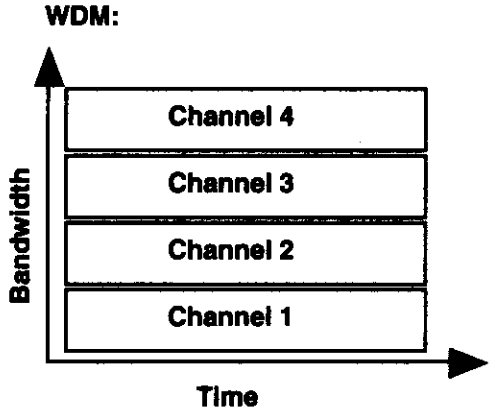

Channels are called wavelengths

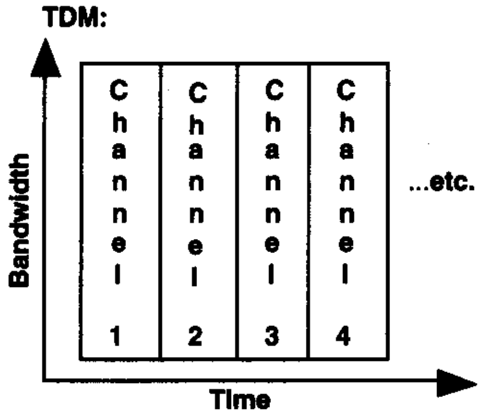

Channels are called time slots

Figure 1–3 WDM and TDM.

WDM is based on a well-known concept called frequency division multiplexing or FDM. With this technology, the bandwidth of a channel (its frequency domain) is divided into multiple channels, and each channel occupies a part of the larger frequency spectrum. In WDM networks, each channel is called a *wavelength*. This name is used because each channel operates at a different frequency and at a different optical wavelength (and the higher the frequency, the shorter the signal's wavelength). A common shorthand notation for wavelength is the Greek symbol lambda, shown as λ.

The wavelengths on the fiber are separated by unused spectrum. This practice keeps the wavelengths separated from each other and helps prevent their interfering with each other. This idea is called channel spacing, or simply spacing. It is similar to the idea of guardbands used in electrical systems. In Figure 1–3, the small gaps between each channel represent the spacing.

Time division multiplexing (TDM) provides a user the full channel capacity but divides the channel usage into time slots. Each user is given a slot and the slots are rotated among the users. A pure TDM system cyclically scans the input signals (incoming traffic) from the multiple incoming data sources (communications links, for example). Bits, bytes, or blocks of data are separated and interleaved together into slots on a single high-speed communications line.

Combining WDM and TDM

Most optical networks (or, for that matter, most networks in general) use a combination of WDM and TDM by time-division multiplexing fixed slots onto a specific wavelength, as shown in Figure 1–4. This concept is quite valuable because it allows multiple users to share one

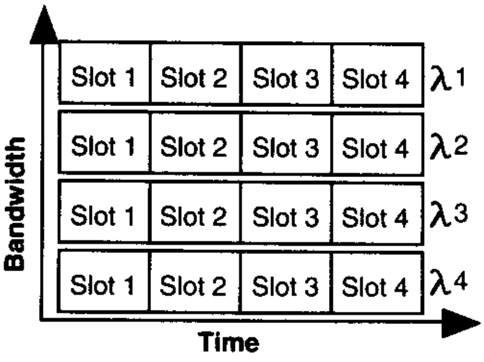

Figure 1–4 Combining WDM and TDM.

WDM wavelength's capacity. With some exceptions, the capacity of one wavelength exceeds an individual user's traffic capacity needs.

These introductory definitions should be sufficient for us to use them in this chapter. In later chapters, TDM and WDM are examined in considerable detail.

THE OPTICAL MARKETPLACE

The optical technology is a high-growth market. As Figure 1–5 shows, it is expected to more than double between 2000 and 2003. Most of the growth will be for terrestrial WDM and optical networks, and submarine cable systems. The growth of conventional TDM systems will continue (shown in Figure 1–5 as SONET/SDH, and explained in

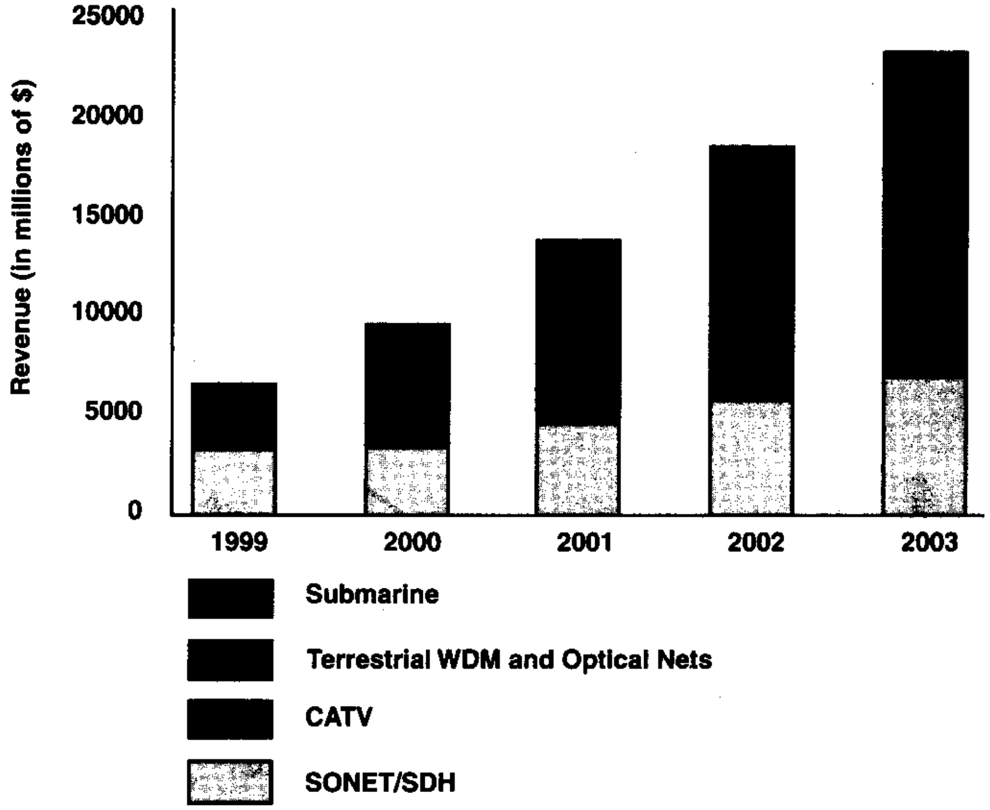

Figure 1–5 Worldwide market for optical components.

Chapter 5), but at a lesser rate than the other areas. Optical systems are being incorporated into CATV networks, but this growth will be modest because the coaxial cable plant to the homes cannot be re-wired in a cost-effective manner. A residence does not need the bandwidth of fiber (at least not for the foreseeable future).

Figure 1–6 shows another study comparing the projected transmission capacity and the demand through 2004 [STRI01]. The figures pertain to the national backbone in the United States and not to the access loops. This study holds that the building-out of optical networks discussed earlier will provide excess capacity for the early part of this decade, and if one examines the gap between demand and capacity, it is reasonable to expect that capacity will exceed demand well beyond 2004.

There are those in the industry who disagree. They state that the upcoming applications will require huge amounts of bandwidth, and that this supposed excess capacity will be consumed by these applications. There is no question that some applications do indeed require a lot of bandwidth. A prime example is interactive high-quality Web traffic, exhibiting the integration of high-resolution, real-time voice, video, and data.

The Local Loop Bottleneck Must Be Solved

It is my view that the upcoming applications, and their demand for large chunks of bandwidth, are not going to be realized to any large

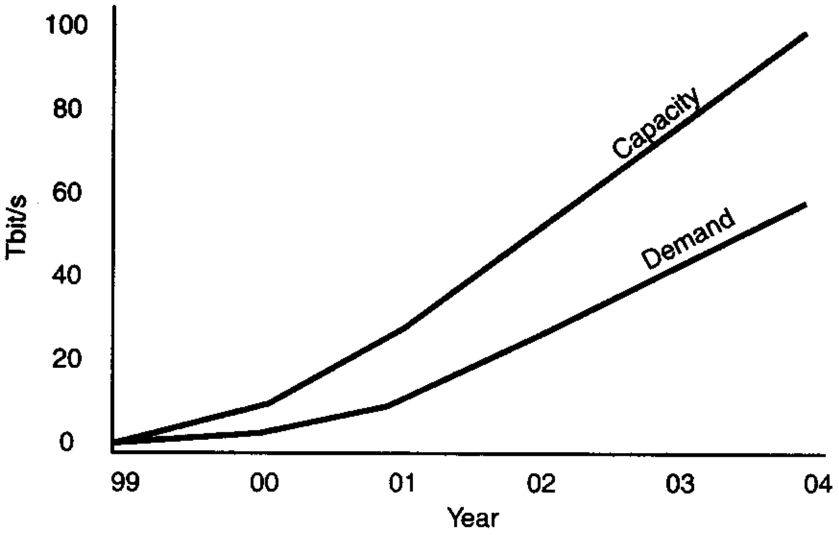

Figure 1–6 Demand vs. capacity [STRI01].

degree until bandwidth is available on the access line (the local loop) to the end user. It does little good to download a Web response to a user at a terabit rate within the network when the vast majority of access lines are restricted to V.90 speeds (56 kbit/s). Certainly, some businesses can afford to purchase large bandwidths from the business to the terabit backbones. But many businesses cannot afford to purchase this bandwidth, nor can the majority of residential users. In addition, broadband access loops are not available to most residences anyway.

The situation in the United States is interesting, and quite frustrating to many customers, because many of them are limited to very low capacity links to the Internet. What is the incentive for the local access providers (the telephone local exchange carriers (LECs), CATV operators, and wireless providers) to expand their local access plant to the megabit or terabit rate? After all, with some minor exceptions (and in spite of the 1984 and 1996 legislative efforts), these companies have a lock on their market. Some people believe that these companies do not have a lot of incentive to invest in the upgrading of their plants.

Maybe so, but the local access providers will expand their plant if they think there is sufficient demand to enable them to make money on their investment. So, beyond the issues of government-sponsored monopolies, is there really that much demand for the deployment of high-capacity systems into the mass marketplace?[1] Most Internet users use the Internet for email or simple text-oriented Web retrievals, and many have been conditioned to the slow response time in their interactions with their networks.

The present situation can be illustrated with a diagram shown in Figure 1–7.[2] The circle in this figure illustrates the relationships of: (a) user applications' requirements for bandwidth (labeled "Applications" in the figure), (b) the capacity of the user or network computers (labeled "CPU" in the figure), and (c) the capacity of the communications media to support traffic (labeled "Bandwidth" in the figure).

Historically, the bottleneck in this circle has varied. At times, it has been the lack of CPU (and memory) capacity in the user's computer. At

---

[1] By mass marketplace, I mean deployment into residences on a large scale, well beyond the 15–20% penetration rate for the current efforts of telephone company and the cable company.

[2] I call this illustration the "eternal circle," because it shows a seemingly never-ending dependency-relationship between the three components.

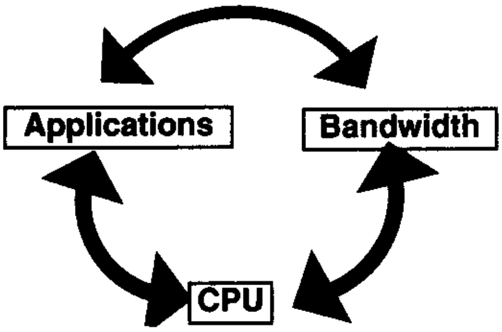

Figure 1–7 The eternal circle.

other times, it has been the lack of capacity in the network to support the capacity requirements of the users' applications.

Figure 1–7 shows the relationships with two-way arrows, which suggests that these three operations are interrelated and dependent upon each other. But which comes first? Does the application's requirement for more capacity lead to faster CPUs and/or the expansion of network bandwidth? Or does the introduction of more bandwidth encourage the development of faster computers and more powerful applications? There are no clear answers to these questions. Sometimes one pushes the other, and at other times the opposite occurs.

However, at this time in the telecommunications industry, we can state the following:

- Within the optical network (the backbone or core network) the bottleneck is the "CPU," because its electrical-based architecture (in switches, routers, and bridges) cannot handle a large number of connected optical fiber WDM links that operate in the terabit range. Thus, the creation of all-optical photonic switches (PXCs) is a high priority in the industry.
- At the edge of the network, and to the end user, the bottleneck is the "bandwidth," but not because of the optical fiber. The bottleneck is due to the continued use of the telephone-based copper plant, and the mobile phone links (and the very slow process of getting it upgraded).
- The "applications" part of the Eternal Circle is a question mark to some people. If the network operators finally provide the bandwidth all the way to the mass market (the residence), will sophisticated three-dimensional, voice/video/data applications be developed to

take advantage of the increased capacity? I believe the answer is a resounding yes, assuming the network operators can keep the price affordable to most households.

Expansion of Network Capacity

One of the more interesting changes occurring in the long-distance carrier industry in the United States is the extraordinary growth of bandwidth capacity. This growth is occurring due to the maturation of the WDM technology, and its wide-scale deployment. It is also occurring due to the aggressive deployment of fiber networks by the "non-traditional" carriers; that is, those carriers who have come into the industry in the last few years.

Figure 1–8 shows the growth of long-distance capacity since 1996, and projections through 2001. The shaded bars show total mileage, and the white bars show total capacity, in terabits per second.

Some people question if this bandwidth will be used. Others see it as an opportunity to discount excess capacity, at the expense of the traditional carriers, who are enjoying healthy profits from their long-distance

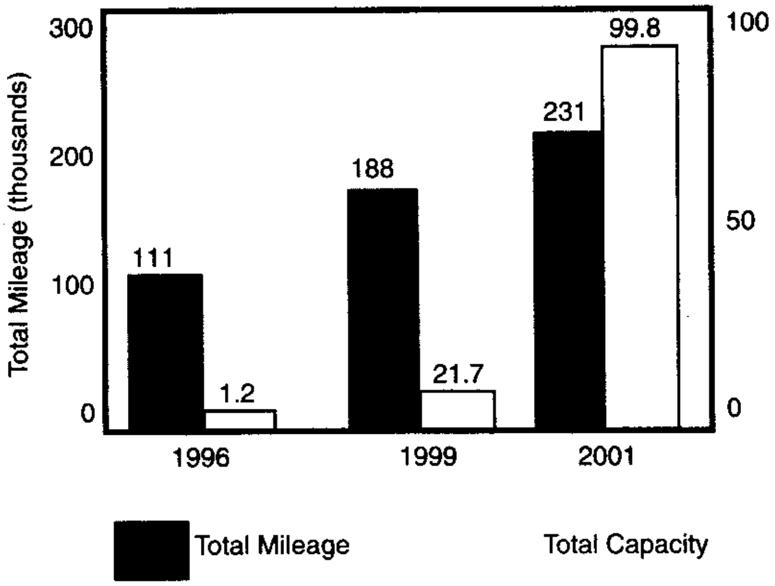

Figure 1–8 Long-distance growth in the United States.

revenues. It will be interesting to see how the scenarios develop over the next few years. Some marketing forecasts state that this situation will lead to a "fire sale" of DS1 and DS3 lines.

WIRELESS OPTICAL SYSTEMS

Another entry into the optical technology is a multichannel optical wireless system from [BELL99]. Figure 1–9 shows a system that operates at 10 Gbit/s using four wavelengths. It can transmit over 2.7 miles of free space. Each wavelength operates at 2.5 Gbit/s. It requires a line-of-sight topology. Other vendors are offering this system, including Nortel Networks.

The systems uses WDM with custom-built telescopes, and standard optical transmitters and receivers. Light signals are sent from a transmitting telescope to a receiving telescope and are focused onto the core of an optical fiber using coupling optics within the second telescope.

The system is attractive in situations where the deployment of fiber cable is not feasible, for example, across restricted-access terrain, or bodies of water. It can be deployed much more quickly than fiber cable systems. It can also offer a cost-effective solution to line-of-sight channels in conference and convention centers.

In most countries, optical wireless requires no governmental licensing or frequency allocation schemes.

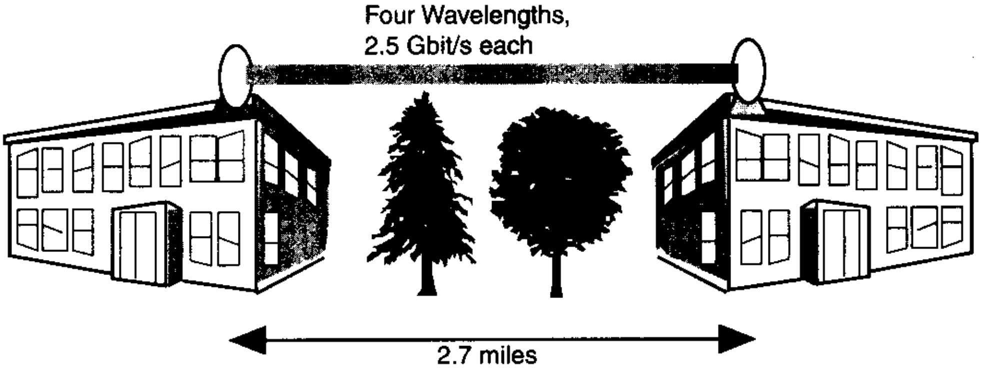

Figure 1–9 Wireless optical systems.

KEY OPTICAL NODES

We now leave the subject of the optical network marketplace, and focus our attention on the nodes (machines) that comprise the network. Figure 1–10 shows the key nodes in an optical network. The topology is a ring, but the topology can be set up either as a ring, a point-to-point, multipoint, or meshed system. In most large networks, the ring is a dual ring, operating with two or more optical fibers. The structure of the dual ring topology permits the network to recover automatically from failures on the optical links and in the link/node interfaces. This is known as a *self-healing ring* and is explained in later chapters.

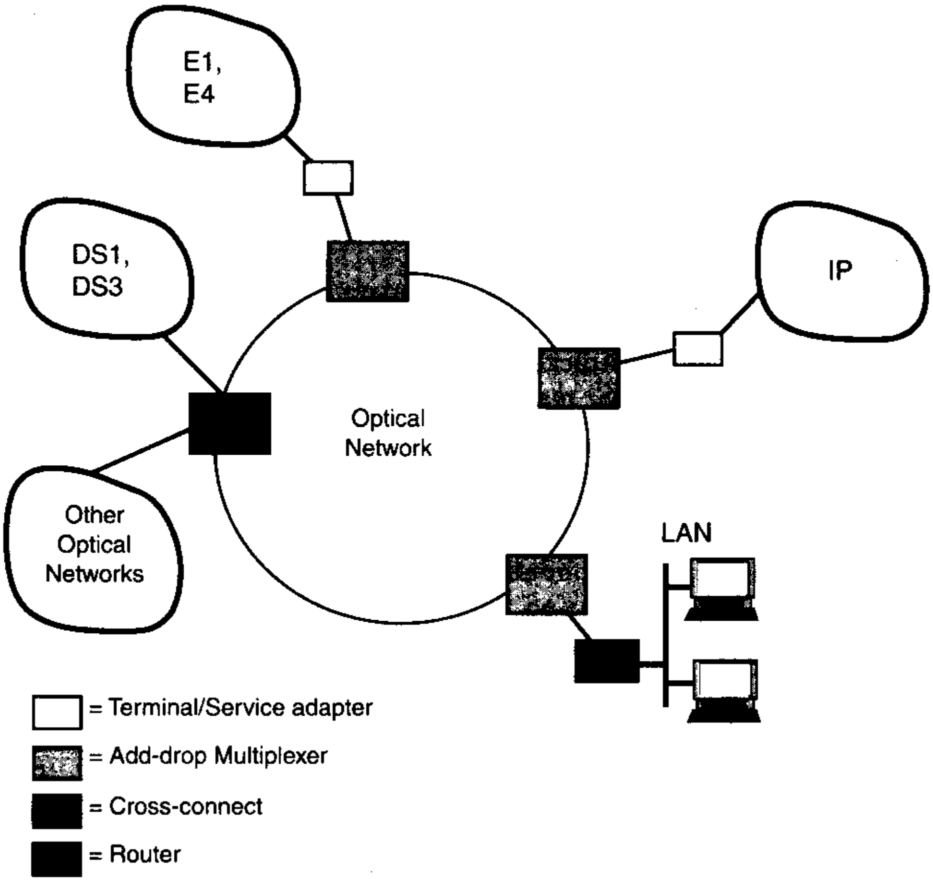

Figure 1–10 Key nodes in the optical network.

End-user devices operating on LANs and digital transport systems (such as DS1, E1, etc.) are attached to the network through a service adapter. This service adapter is also called an access node, a terminal, or a terminal multiplexer. This node is responsible for supporting the end-user interface by sending and receiving traffic from LANs, DS1, DS3, E1, ATM nodes, etc. It is really a concentrator at the sending site because it consolidates multiple user traffic into a *payload envelope* for transport onto the optical network. It performs a complementary, yet opposite, service at the receiving site.

The user signals, such as T1, E1, ATM cells, etc., are called *tributaries*. The tributaries are converted (mapped) into a standard format called the *synchronous transport signal (STS)*, which is the basic building block of the optical multiplexing hierarchy. The STS signal is an electrical signal. The notation STS-n means that the service adapter can multiplex the STS signal into higher integer multiples of the base rate, The STS signals are converted into optical signals by the terminal adapter and are then called OC (optical carrier) signals.

The terminal/service adapter can be implemented as the end-user interface machine, or as an add-drop multiplexer (ADM). The ADM implementation multiplexes various STS input streams onto optical fiber channels. OC-n streams are demultiplexed as well as multiplexed with the ADM.

The term *add-drop* means that the machine can add or drop payload onto one of the fiber links. Remaining traffic that is not dropped passes straight through the multiplexer without additional processing.

The *cross-connect (CS)* machine usually acts as a hub in the optical network. It can not only add and drop payload, but it can also operate with different carrier rates, such as DS1, OC-n, E1, etc. The cross-connect can make two-way cross-connections between the payload and can consolidate and separate different types of payloads. For example, the cross-connect can consolidate multiple low bit-rate tributaries into higher bit-rate tributaries, and vice versa. This operation is known as *grooming*.

Key Terms for the Cross-connect

The convention in this book is to use three terms to describe the optical cross-connect. There is a spate of terms to describe a cross-connect. I counted six terms in one paper alone. To make sure there is no ambiguity about the optical cross-connect in this book, the following terms are used:

- *Optical/Electrical cross-connect (OXC):* Receives optical signals, converts them to electrical signals, makes routing/switching and/or ADM decisions, then converts the electrical signals back to optical signals for transmission. These operations are also noted as O/E/O. This technique is also called an opaque operation.
- *Photonic cross-connect (PXC):* Performs the functions of the OXC, but performs all operations on optical signals. These operations are also noted as O/O/O, and are also called transparent operations.
- *Cross-connect (XC):* A more generic term, used when it is not necessary to distinguish between the OXC or the PXC.
- *Switch:* Some recent literature distinguishes between a cross-connect and a switch. This literature states that cross-connect is an outdated term! Well, the term switch has also been around for quite a while. Anyway, the book uses the terms cross-connect and switch synonymously.

Other terms and different definitions of optical nodes are used by various vendors, network operators, and standards groups. In some cases, they are the same as those just cited; in other cases, they are different. Where appropriate, I will distinguish and explain these other terms.

OTHER KEY TERMS

Other terms need to be defined and clarified in order for readers to understand the other chapters in this book. For the first few times I use these terms, I will refer you back to these definitions, or repeat them. Unfortunately, the industry is not consistent in the use of some of these terms; some varying interpretations are explained below.

- *Fiber link set:* This term refers to all the fibers (if there are more than one) connecting two adjacent XCs or other fiber nodes. The link set may consist of scores of individual optical fibers and hundreds of wavelengths.
- *Edge, ingress, egress nodes:* These terms refer to the placement of the XCs at the boundaries of the network. The term edge encompasses both ingress and egress. Ingress obviously means the XC sending traffic into the network, and egress is the node sending traffic out of the network.

- *Interior, transit, or core nodes:* These three terms refer to an XC that is located inside the optical network and communicates with other XCs for internal network operations or with the edge nodes for communications (perhaps) outside the network.
- *Optical switched path (OSP):* The optical path between two adjacent optical nodes. The OSP is one logical channel of a fiber link set.
- *Lightpath and trail:* The term lightpath defines an end-to-end optical path through one or more optical nodes or networks to the end users. This term is also used in some literature to identify the optical path between two adjacent nodes, so it must be interpreted in the context of its use. Also, some literature uses the terms lightpath and trail synonymously.
- *Label switched path (LSP):* The end-to-end MPLS path across one or more MPLS nodes (and perhaps optical as well) networks to the end users.

Another Look at the Optical Node

Figure 1–11 shows a more detailed view of the optical network node and its components [NORT99b]. This example shows the light

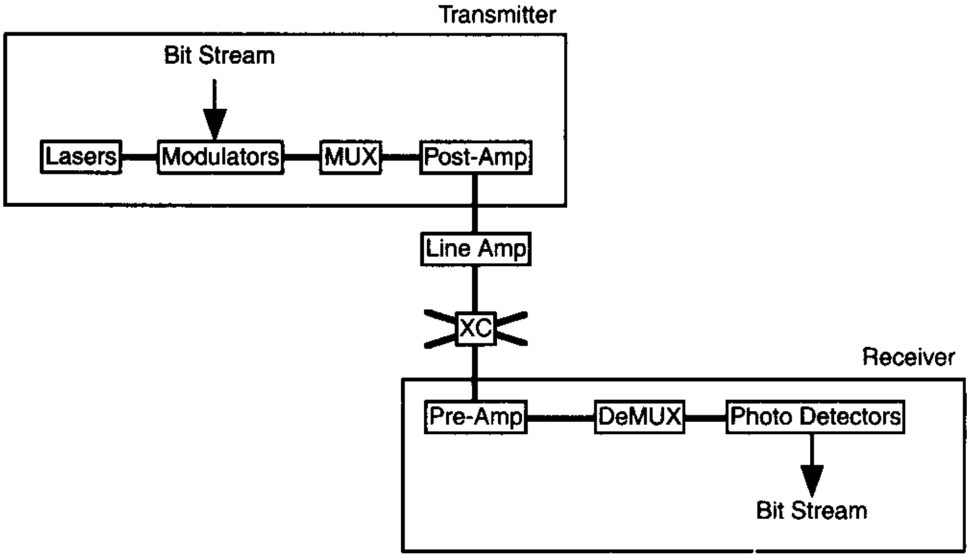

Figure 1–11 Optical components in more detail [NORT99b].

signal transmitted from the left side to the right side of the figure. The events for the operation are explained below, with references to the chapters in this book that provide more detailed explanations.

- First, laser devices generate light pulses tuned to specific and precise wavelengths, such as 1533 or 1557 nanometers (nm). Lasers are explained in Chapter 3.
- Next, the optical modulators accept the electrical signal (an incoming bit stream), and convert it to an optical signal. In addition to the conversion, the modulator uses the incoming bit stream to make decisions about turning the light stream on and off to represent the digital 1s and 0s of the incoming stream. Chapters 3 and 4 provide more information on this process.
- The multiplexer (MUX) combines different TDM slots or WDM wavelengths together. Chapters 5, 6, and 7 explain multiplexing in considerable detail.
- The signal is passed to an optical post-amplifier (Post-Amp). This amplifier boosts the strength of the power of the signal before it is sent onto the fiber. See Chapters 3 and 7 for more information on amplifiers.
- On the fiber, a dispersion compensation unit (not shown in Figure 1–11) corrects the dispersion of the signal as it travels through the fiber. As explained in more detail in Chapters 3 and 7, dispersion is the spreading of the light pulses as they travel down the fiber, which can cause interaction (and distortion) between adjacent pulses.
- As the signal travels down the fiber, it loses its strength. Therefore, the signal power is periodically boosted with an amplifier (Line Amp) to compensate for these losses, again as explained in Chapters 3 and 7.
- There may be an XC on the link to switch the signals to the correct destination. The manner in which the signals are relayed through a cross-connect is one of keen interest in the industry and is examined in Chapters 8, 9, 10, 12, and 14.
- At the final receiver, optical pre-amplifiers (Pre-Amp) boost the strength of the signal once again (Chapters 3 and 7).
- A demultiplexer separates the multiple wavelengths (Chapters 5, 6, and 7).
- Optical photodetectors convert the optical wavelengths into an electronic bit stream (Chapter 3).

EVOLUTION OF OPTICAL SYSTEMS

To set the stage for subsequent chapters, Figure 1–12 shows the evolution of optical systems since the late 1980s/early 1990s to the present time [GILE99].

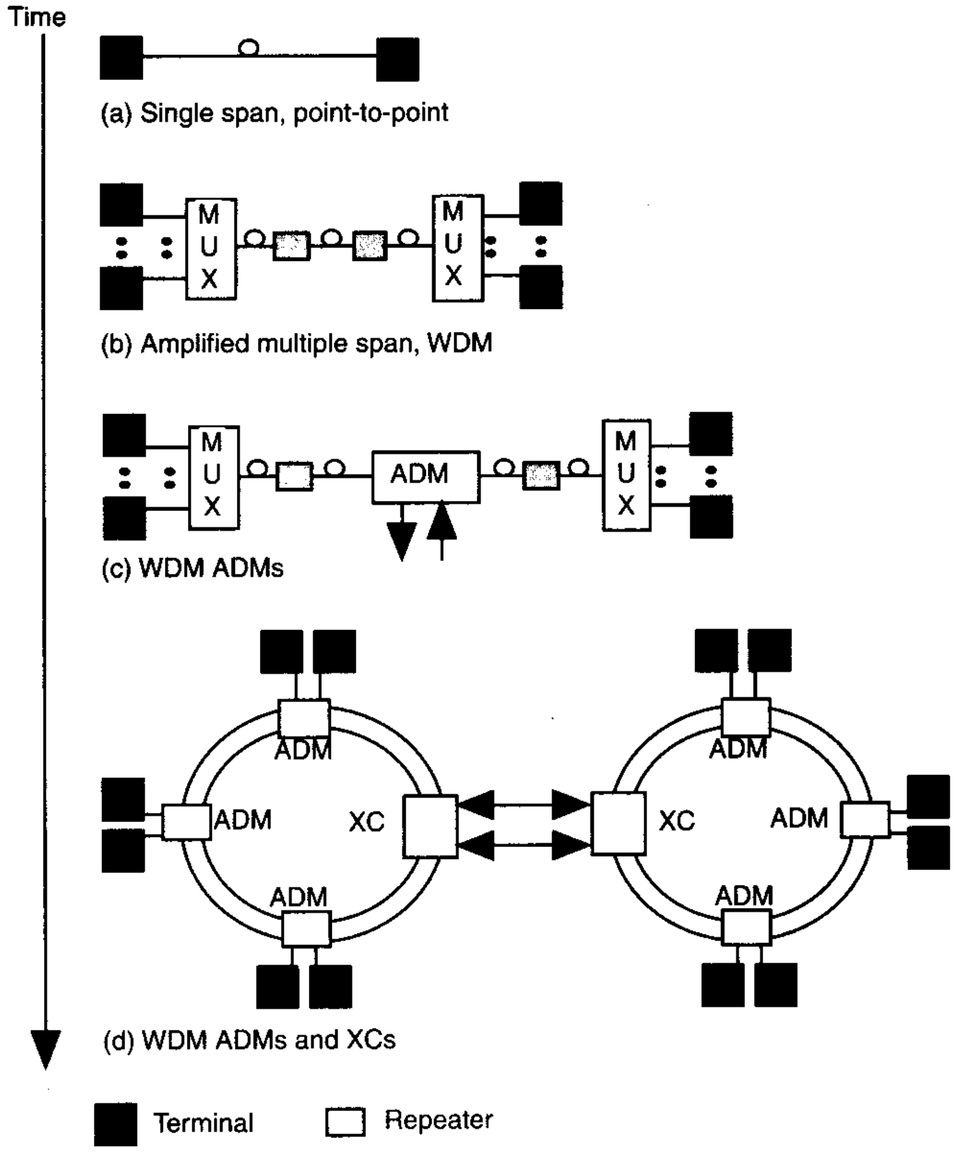

Figure 1–12 Evolution of optical systems [GILE99].

The early systems had a single channel point-to-point topology, as shown in Figure 1–12 (a). The short fiber spans did not use optical amplifier repeaters, so the span lengths rarely exceeded 40 km. The restriction of length was due partially to the limit of laser transmitter power of 1 mW. Optical systems at this time used a single fiber to increase the overall transmission capacity of the point-to-point link.

The discovery of erbium-doped fiber amplifiers (repeaters) was a major milestone in the ability to extend the fiber link. The optical amplifier allowed the optical transmission to extend to wide areas (see Figure 1–12 (b)).

Now that a link could span thousand of kilometers, it became feasible to deploy add-drop multiplexers (ADMs) to allow the connection of points along the link, shown in Figure 1–12 (c). (Of course, ADMs in optical networks have been around for over ten years, but wavelength ADMs (WADMs) are more recent.) As of this writing, the fixed-wavelength ADM (of selected channels) is the state-of-the-art implementation for WADMs. But WADMs that are dynamically reconfigurable are the next phase for WADMs.

As shown in Figure 1–12 (d), cross-connects (XCs) permit a more powerful grooming of traffic between optical networks (in this example, optical rings). This configuration has been deployed for over a decade in TDM optical systems; wavelength XCs, noted as PXCs in this book, are the next stage of the evolution.

KEY ATTRIBUTES OF OPTICAL FIBER

Finally, as a prelude to the remainder of the book, we conclude this chapter by reviewing the major attributes of optical fiber.

The advantages of fiber optics (compared to copper cable) include superior transmission quality and efficiency. Since the optical signal has none of the characteristics association with electrical signals, optical fiber does not suffer from common electromagnetic effects such as experiencing interference from other electrical components, such as power lines, electrical machines, and other optical links.

Because it does not emanate energy outside the fiber, the optical signal is more secure than copper and wireless media, which are easy to monitor and glean information from the residual energy emanating from these media.

Glass fiber is very small and of light weight, a significant attribute for network operators who must install communications links in buildings, ducts, and other areas that have very limited space for the communications links.

We learned earlier that fiber has a very wide bandwidth which allows for the transport of very large payloads, some in the terabits-per-second range.

Since the fiber is comprised of glass with a very small diameter, it is fragile and is somewhat difficult to connect and splice. Also, because glass is not a conductor of electrical current, it cannot carry power to the regenerators (which are used to strengthen signals on long spans). This situation is changing, as passive optical networks are deployed (a subject for Chapter 8).

SUMMARY

Optical fiber and optical-based networks have revolutionized the world of telecommunications, principally because of their extraordinary transmission capacity. They have replaced almost all the older media, such as copper, in the large backbone networks in the world, such as the telephone networks, and the Internet. Unfortunately, due to the dominance of copper wire in the telephone local loop plant, they have not been installed (to any significant) extent in residences. Nonetheless, optical fiber technology will continue to grow, and with the advent of WDM and powerless amplifiers, their presence will become commonplace in all high-speed networks.

2

The Telecommunications Infrastructure

This chapter introduces the basic components that make up a telecommunications network. The focus is on the components that are germane to the subject of optical networks. We begin with the part of the network which connects businesses and residences in a local/metropolitan area. This discussion is followed by an examination of the long-haul backbone network architecture. Next, the role of layered protocols in transport networks is explained. The chapter concludes with a discussion of the digital carrier systems, with examples of the digital multiplexing hierarchy.

The first part of this chapter is meant as a basic tutorial, so the more advanced reader can skip to the section titled "Considerations for Interworking Layer 1, Layer 2, and Layer 3 Networks." This section should be read by all.

THE LOCAL CONNECTIONS

Most businesses and residences are able to communicate with each other via the telephone and, increasingly, the CATV networks, as well as with mobile phones. See Figure 2–1. These communications connections are typically called local interfaces, local loops, subscriber loops, metropolitan networks, or simply the user-network interface (UNI).

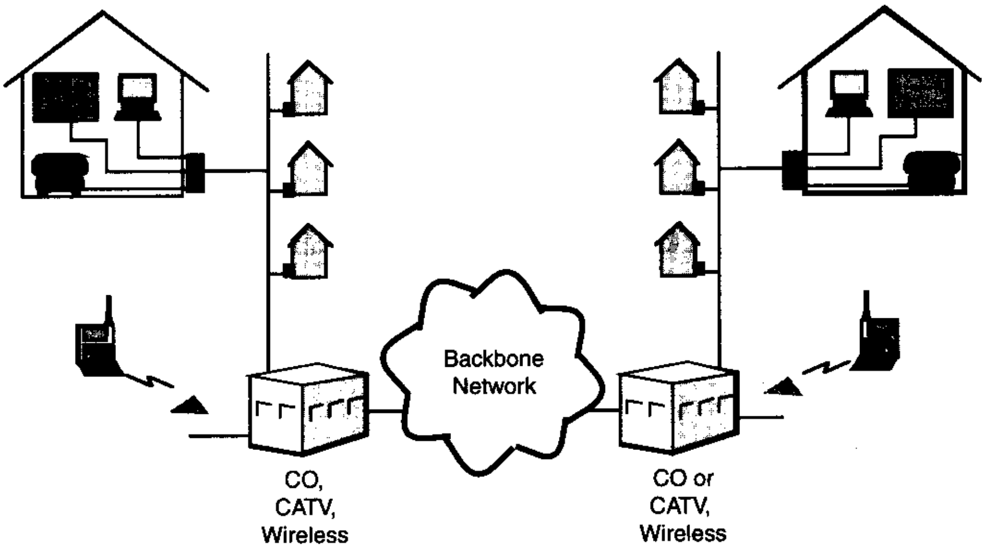

Figure 2–1 The local exchanges.

For the telephone company (telco) connection, the user gains access to the telephone system through the ubiquitous copper wires that run from the telco central office (CO) to the user's telephone. For the cable TV connection, the TV coaxial cable is the channel installed between the user and the cable provider. Of course, for the mobile phone user, access is accomplished through the air.

Increasingly, optical fibers are replacing some of the copper wire and coaxial cable installations, primarily to office buildings. Thus far, there is limited interest in extending fiber to residences, although a few new residential developments have deployed fiber to the home (FTTH). The main thrust is the deployment of fiber to the curb (FTTC), where the optical signals are converted to electrical signals for transmission back and forth on the copper wire or coaxial cables at the user's premises.

THE BACKBONE CONNECTIONS

For transport of the customer's traffic across a wide geographical area, the long-haul backbone network is employed, as shown in Figure 2–2. The common term in the industry for this network is a digital transport network or transport network, and it is the subject of this book.

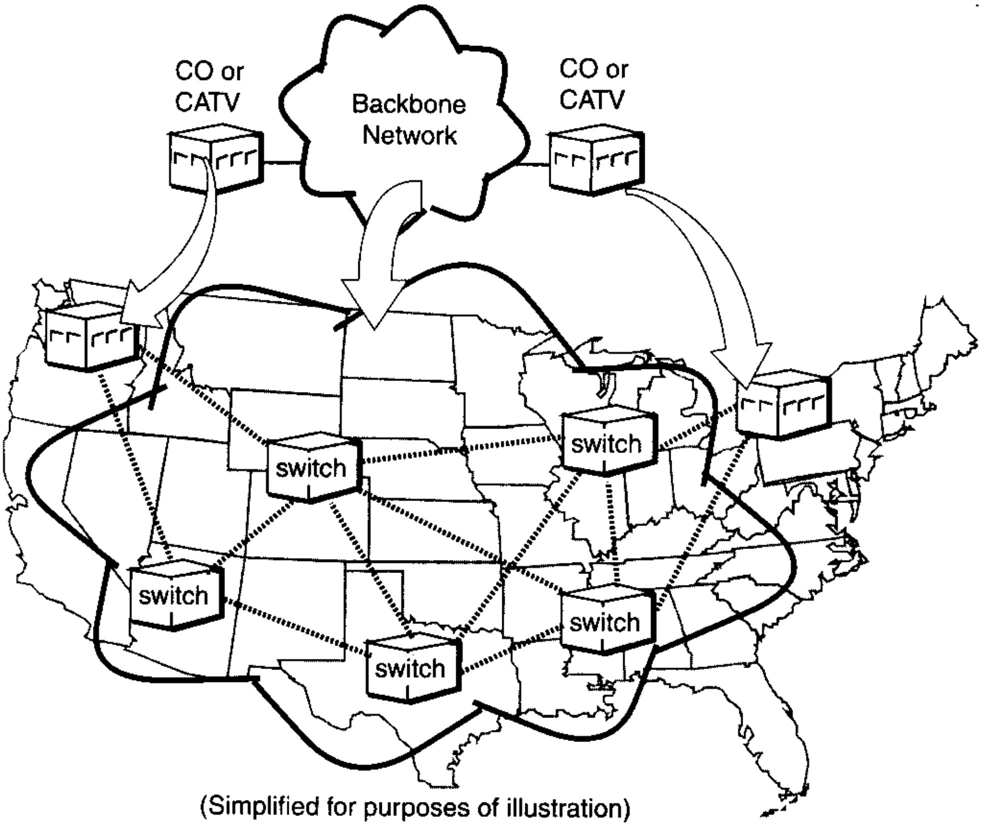

Figure 2–2 The backbone.

The term network in Figure 2–2 is a misnomer because the world's "transport network" consists of thousands of networks that are owned by network service providers such as the telephone companies, and the postal, telephone, and telegraph ministries (PTTs) in many parts of the world. Some Internet service providers (ISPs) have deployed their own networks as part of this backbone structure. Other companies build networks and rent them to other service providers; these companies are often called a carrier's carrier.

The blank network "cloud" at the top of Figure 2–2 has been filled in at the bottom of the figure, showing the positions of communications nodes and links. The communication among the nodes in the network is accomplished with procedures defined by a network-to-network interface (NNI).

Whatever the form taken, the important aspects of the backbone are the requirements for (a) supporting massive amounts of traffic (in the terabit-per-second rate), (b) assuring that the customers' traffic arrives safely at its destination, and (c) making sure the received traffic is of high quality (the voice image is pristine, the data message has no errors, the TV picture is clear). As explained in Chapter 1, the optical fiber technology is an indispensable tool to support these three critical requirements.

The dashed lines inside the backbone in Figure 2–2 are usually optical fiber, or some form of high-capacity radio medium, such as satellite or microwave. For the past ten years, the backbone links have been migrating to optical fiber.

The switches residing in the network are responsible for relaying the customer's traffic through the network to the final destination. With few exceptions, the switches are configured to examine a called party telephone number or a destination IP address to make their forwarding decision.

At this stage of the development of optical networks, these switches are O/E/O devices; as noted in Chapter 1, they convert the optical signal to electrical images so they can examine the information in the traffic to glean the telephone number or the IP address to make their routing decision. Then they convert the electrical signal back to optical for transmission to the outgoing optical link. As explained in more detail in subsequent chapters, the movement is toward photonic (PXC) switches (O/O/O devices) due to their capacity and speed.

THE DIGITAL MULTIPLEXING HIERARCHY

It was noted that the backbone network must support millions of customers and transmission rates in scores of terabits per second. In addition, this huge customer payload must be organized in a structured manner if it is to be managed properly. One key aspect of supporting, organizing, and managing this traffic is through digital multiplexing technologies.

The idea of digital multiplexing is to use multiplexing levels. As seen in Figure 2–3, at the lowest level, the user's traffic is not multiplexed. The most common example is a customer's voice traffic operating at 64 kbit/s (it is called a DS0 signal). In many networks (at a PBX in a building, or at the telephone central office), 24 DS0 signals are multiplexed together to make better use of, say, the copper wires going to the telephone office. These combined signals are called DS1 (and, in many circles, T1). To exploit the transmission capacity of other media, such as microwave, coaxial cable, and certainly optical fiber, 28 DS1s are multiplexed into what is called DS3 (or T3, in many circles).

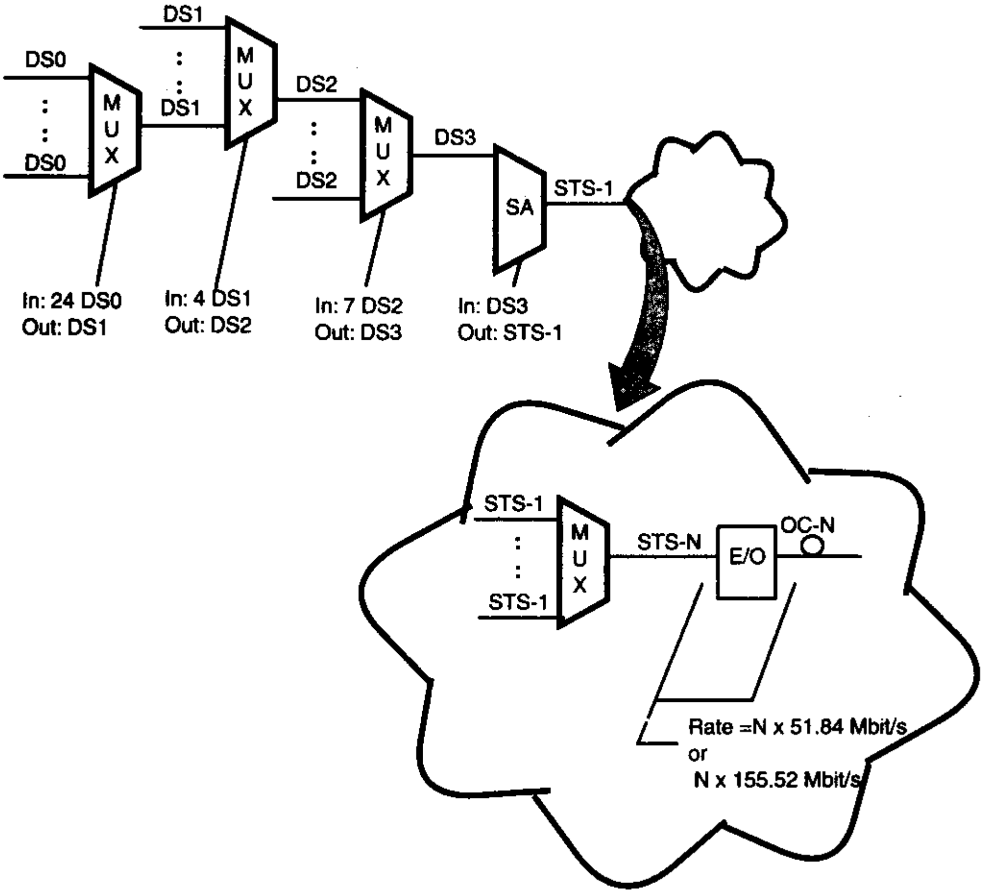

Figure 2–3 The North American multiplexing hierarchy.

Figure 2–3 shows another stage of the multiplexing called the DS2 signal. This operation is usually invisible to a customer; that is, a customer can purchase DS0, DS1, and DS3 payloads, but DS2 is usually not part of a commercial offering. It is placed in this figure because it plays a role in some of the SONET operations discussed in later chapters.

On the right side of Figure 2–3 is yet another stage of the operations. The DS3 signal is mapped into an STS-1 signal (synchronous transport signal number 1) by a service adapter (SA). The purpose of this operation is to add some overhead bytes to the DS3 and to align the DS3 properly with the network's other payloads.

The network cloud is expanded at the bottom part of Figure 2–3. Another multiplexer combines multiple STS-1s into another higher capacity

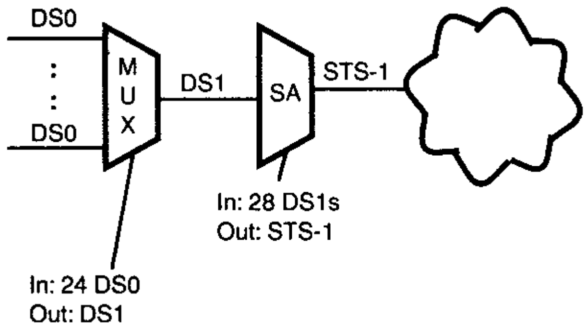

Figure 2–4 Another way to multiplex the users' signals.

payload, STS-N. The N can take different values (two values of N are shown in the figure), depending on specific configurations in the network, a topic explained in greater detail in Chapter 5.

Finally, an electrical-to-optical converter (E/O) converts the electrical signal to an optical signal, now called the OC-N signal. The small circle on the OC-N link is a common notation to identify optical (in contrast to electrical) links.

There are other ways to groom the traffic for transport across the optical network. As suggested in Figure 2–4, it is possible (and a common practice) to multiplex 28 DS1s into separate tributaries and carry them in the STS-1 signal.

THE DIGITAL SIGNALING HIERARCHIES

During the past 30 years, three different digital signaling hierarchies have evolved in various parts of the world. These hierarchies were developed in Europe, Japan, and North America, and are first generation digital transport systems. See Figure 2–5. Fortunately, all are based on the basic clocking rate of 125 µsec, and the basic 64 bit/s signal. Therefore, the basic architectures interwork reasonably well. However, their multiplexing hierarchies differ considerably, and the analog/digital conversion schemes are not the same.

Japan and North America base their multiplexing hierarchies on the DS1 rate of 1.544 Mbit/s. Europe uses an E1 2.048 Mbit/s multiplexing scheme. Thereafter, the three approaches multiplex these schemes in multiples of these rates (plus some overhead bits).

The Synchronous Digital Hierarchy (SDH) is the official international standard for second generation digital carrier networks. It is based

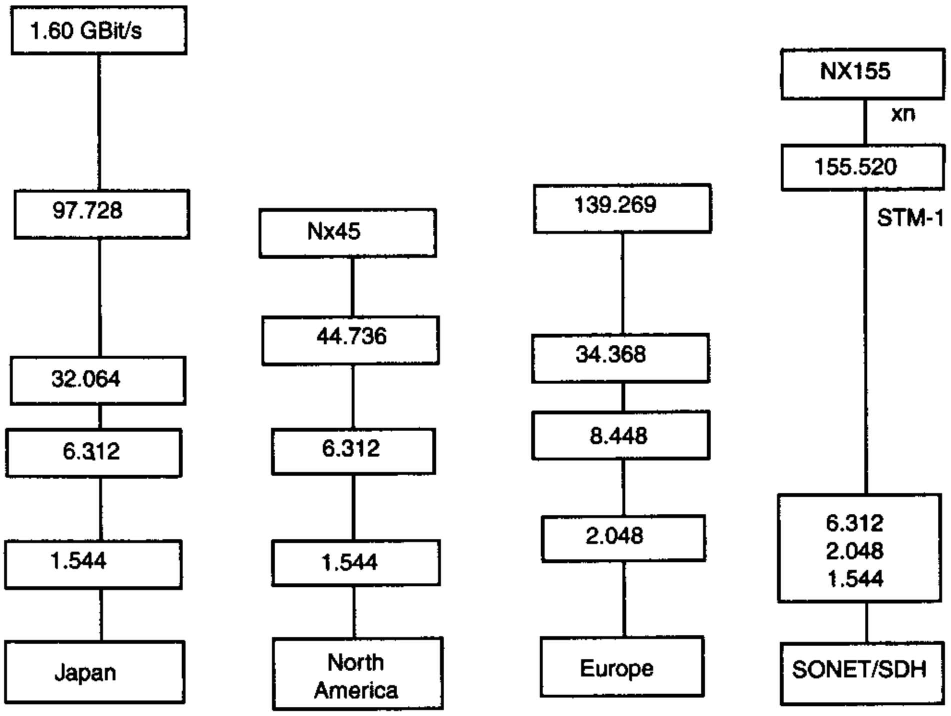

Note: Unless noted otherwise, speeds in Mbit/s

Figure 2–5 The digital multiplexing hierarchy.

on SONET, and specifies a different multiplexing hierarchy. The basic SDH rate is 155.52 Mbit/s. It then uses a N x 155.52 multiplexing scheme. A smaller rate than 155.52 Mbit/s is available for SONET. The smaller rate is 51.840 Mbit/s.

Shortly after the inception of T1 in North America in the early 1960s, the ITU-T published the E1 standards, which were implemented in Europe and most other parts of the world. Japan followed with a hierarchy similar to the North American specifications at the lower hierarchical levels, but not at the higher levels.

T1 is based on multiplexing 24 users onto one physical TDM circuit. T1 operates at 1,544,000 bit/s, which was (in the 1960s) about the highest rate that could be supported across twisted wire-pair for a distance of approximately one mile. Interestingly, the distance of one mile (actually, about 6000 feet) represented the spacing between manholes in large cities. They were so spaced to permit maintenance work such as splicing cables and the placing of amplifiers. This physical layout provided a convenient means to replace the analog amplifiers with digital repeaters.

T1 OR DS1; T3 OR DS3?

The term T1 was devised by AT&T (before divestiture in 1984) to describe a specific type of carrier technology. Strictly speaking, the term DS1 is the correct post-divestiture term. To keep matters simple, this

North American Multiplexing Hierarchy

Type	*Digital Bit Rate*	*Voice Circuits*	*T1*	*DS3*
DS1	1.544 Mbit/s	24	1	-
DS1C	3.152 Mbit/s	48	2	-
DS2	6.312 Mbit/s	96	4	-
DS3	44.736 Mbit/s	672	28	1
DS4	274.176 Mbit/s	4032	168	6

European Multiplexing Hierarchy

Type	*Digital Bit Rate*	*Voice Circuits*	*System Name*
E1	2.048 Mbit/s	30	M1
E2	8.448 Mbit/s	120	M2
E3	34.368 Mbit/s	480	M3
E4	139.264 Mbit/s	1920	M4
E5	565.148 Mbit/s	7680	M5

Japanese Multiplexing Hierarchy

Type	*Digital Bit Rate*	*Voice Circuits*	*System Name*
1	1.544 Mbit/s	24	F1
2	6.312 Mbit/s	96	F6M
3	34.064 Mbit/s	480	F32M
4	97.728 Mbit/s	1440	F100M
5	397.20 Mbit/s	5760	F400M
6	1588.80 Mbit/s	23040	F4.6G

Figure 2–6 The digital multiplexing schemes.

book uses the term T1 synonymously with the term DS1, and the term T3 synonymously with DS3. Figure 2–6 shows the more common digital multiplexing schemes used in Europe, North America, and Japan, and the terms associated with these schemes.

THE LAYERED PROTOCOL MODEL IN THE TRANSPORT NETWORK

Throughout this book, it is convenient to use the OSI layered protocol model to explain aspects of transport networks. If you are new to layered protocols and the OSI model, see [BLAC91]. Figure 2–7 depicts a view of this model and the relationships of the layers of transport networks, MPLS, and the principal Internet protocols for the model.

The first observation is that transport networks operate at the physical layer (layer 1) of the model. Originally, the OSI model defined the physical layer as one that had restricted, physical functions such as signal generation/reception, clocking, the definition of the media (fiber, copper), and so on. While this statement still holds, the 2G and 3G physical layers are quite powerful, and define many other operations, such as extensive diagnostics, backup/recovery, and bandwidth provisioning.

OSI	1G, 2G, 3G Transport Networks	MPLS	Internet
Application			Web, Email, etc.
Presentation			Not Used
Session			Not Used
Transport			TCP/UDP
Network			IP
Data Link		MPLS	MPLS
Physical	T1, E1, SDH, OTN, etc.	T1, E1, SDH, OTN, etc.	T1, E1, SDH, OTN, etc.

Figure 2–7 The layered model and other layered protocols.

For this book, the layer 1 remains the same for MPLS and IP operations. In addition, it is assumed the IP (and the upper layers above IP) will operate over MPLS, unless otherwise stated.

CONSIDERATIONS FOR INTERWORKING LAYER 1, LAYER 2, AND LAYER 3 NETWORKS

One of the fundamental concepts in the 3rd generation transport network is the interworking of the layer 1, layer 2, and layer 3 entities; in effect, the graceful interactions among what are called circuit-switched operations at layer 1 and packet/frame/datagram operations at layers 2 and 3.[1] This book covers this topic in many sections in subsequent chapters. For this discussion, it is noteworthy that 3G optical networks are essentially layer 1 networks, and they display some of the characteristics of 1G and 2G circuit-switched networks. The salient aspects of these systems for this book include the following points:

1. Circuit-switched nodes (such as an SS7 switch) may have thousands of physical links (ports).
2. These ports (unlike IP networks) do not have IP addresses. They are usually identified by (a) a switch-specific port number (interface) and (b) a channel ID on the port; for TDM, a slot number (say, DS0 # 6, or STS-1 # 2), and for WDM, a wavelength ID (say, wavelength # 4 on the fiber).
3. Neighbor nodes do not need to know about their neighbor's internal port number ID; they need to know the channel ID on the port in order to recognize each piece of traffic. The exception to this statement is when many links are bundled together as a fiber link set; in this situation, the neighbor nodes need to know the specific link in the set.
4. Many of the circuit-switches' features are configured manually, and many of the operations in the switch remain static throughout the sending and receiving of user traffic.

---

[1] The terms packet and frame are often used interchangeably. If they are distinguished from each other, frame refers to a unit of traffic associated with layer 2 (and in some transport networks, even layer 1), and packet refers to any unit of traffic associated with layers 3 and above. The term datagram is always associated with an IP packet at layer 3.

5. The switching technology on circuit-switches is based on a very fast hardware-oriented cross-connect fabric, wherein the input and output ports are very tightly synchronized to the receiving and sending of discrete, fixed-length TDM slots of traffic.

These important points are revisited in Chapter 10; see "Considerations for Interworking Layer 1 Lambdas and Layer 2 Labels."

SUMMARY

Most businesses and residences communicate with each other via the telephone, the CATV cable, and mobile phones. If the connections are made over a wide area, the long haul networks (the 1st, 2nd, or 3rd transport networks) come into play. Most of the long haul media is optical fiber, but the local connections from the long haul network to a customer are a mix of copper wire, coaxial cable, optical fiber, and wireless technologies.

Multiplexing is an indispensable tool for aggregating (and managing) customers' traffic into high bit rates for transport across the specific medium. Different multiplexing hierarchies have evolved over time, but the trend is to migrate to one multiplexing scheme—a subject for later chapters.

3

Characteristics of Optical Fiber

This chapter is a tutorial on optical fiber. Hundreds of books and papers are devoted to this subject. The approach is to provide the newcomer with the requisite information needed to understand the role they play in 3G optical transport networks and in the emerging optical Internet. For those readers who want to delve into the engineering aspects of this technology, I provide several references as we proceed through the chapter.

THE BASICS

Before one can grasp the underlying concepts of fiber optics, several aspects of light should be considered. It is instructive to note that light is part of the electromagnetic spectrum. The light that is visible to the human eye is only a fraction of the entire spectrum range, and light frequencies are several orders of magnitude higher than the highest radio frequencies. Different colors or wavelengths that constitute light are nothing more than different frequencies that propagate at various speeds over a medium. The prism can be used to demonstrate this phenomenon. When white light (which is the composite of all visible colors) enters the prism, it refracts and bends differently because of the individual speeds of the various frequencies or colors. Because of this refraction process, the colors exit separately. So, in a sense, the prism is an optical demultiplexer.

Fiber optics is the technology of transmitting information over optical fiber in the form of light. The light energy consists of *photons*, which is the quantum of radiant energy. The electrical signal to be transmitted is converted at the source into a light signal, which is then modulated and sent to a light-emitting diode or a laser for transmission through the fiber. At the receiving end, the detector converts the modulated light signal back to its electrical equivalent.

The fiber used in communications is usually a fine strand of glass weighing on the order of one ounce per kilometer, and it is as thin as human hair (50–150 microns in diameter for multimode and 8 microns for single-mode fiber, discussed shortly). Plastic fiber also exists, and it is applied to short distance needs.

Some of the advantages of fiber optics noted in Chapter 1 (compared to copper cable) include superior transmission quality and efficiency, as well as the elimination of crosstalk, static, echo, and delay problems. Fiber also minimizes environmental effects such as weather, water, and freezing. Once again, optical fiber is very small and of light weight. It has a wide bandwidth which allows for the transport of very large payloads.

THE WAVELENGTH

Light is a form of energy radiation at very high frequencies. For example, the frequency of the color blue operates at 600 THz. In optical networks, the practice is not to use frequencies in defining an optical channel, but to use the term *wavelength*. The wavelength is defined by the length in nanometers (nm) of the optical signal. Figure 3–1 shows how a wavelength is measured. It is the distance (measured in meters) between successive iterations of the oscillating light signal. The frequency of the signal is the number of waves that are sent on the fiber (and measured at a fixed point) in one second. Frequency is designated by the term hertz (Hz).

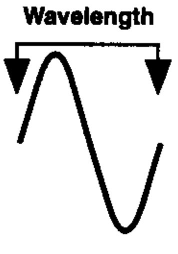

Figure 3–1 The wavelength.

Table 3–1 Common Terms

Multiplication Factor			Prefix	Symbol	Meaning
0.1	=	10^{-1}	deci	d	Tenth
0.01	=	10^{-2}	centi	c	Hundredth
0.001	=	10^{-3}	milli	m	Thousandth
0.000 001	=	10^{-6}	micro	μ	Millionth
0.000 000 001	=	10^{-9}	nano	n	Billionth
0.000 000 000 001	=	10^{-12}	pico	p	Trillionth
0.000 000 000 000 001	=	10^{-15}	femto	f	Quadrillionth
0.000 000 000 000 000 001	=	10^{-18}	atto	a	Quintillionth

The higher the frequency of the signal, the shorter the wavelength. This fact is demonstrated by:

Frequency (in Hz) = Speed of light in a vacuum (in meters) / Wavelength (in meters)

Thus, a frequency of 192.1 THz operates with a wavelength of 1560.606 nm (191.1 THz = 299,792,458 / .000001560606). As another example, a frequency of 194.7 THz operates with a wavelength of 1539.766 nm (194.7 THz = 299,792,458 / .000001539.766).

Table 3–1 should prove helpful as you read this book about wavelengths. As you can see, a nanometer wavelength is one billionth of a meter. A typical optical wavelength is 1552.52 nm.

THE BASIC COMPONENTS

Like other media technologies, a fiber optic system has three basic components (see Figure 3–2): the optic fiber or light guide, the transmitter or light source, and the receiver or light detector.

The Source of the Signal

For producing the light signal, the emission can be a spontaneous emission or a stimulated emission [DUTT98].[1] The spontaneous emission takes place when a electron (well, many electrons) is brought to a very

---

[1] For the reader who needs basic information on optical fiber beyond this chapter, Dutton's book is top-notch. It is, overall, the best-written book that I know of on the subject. It has very little overlap with the subject matter in this book, so it would serve as a good companion to the book you are now reading.

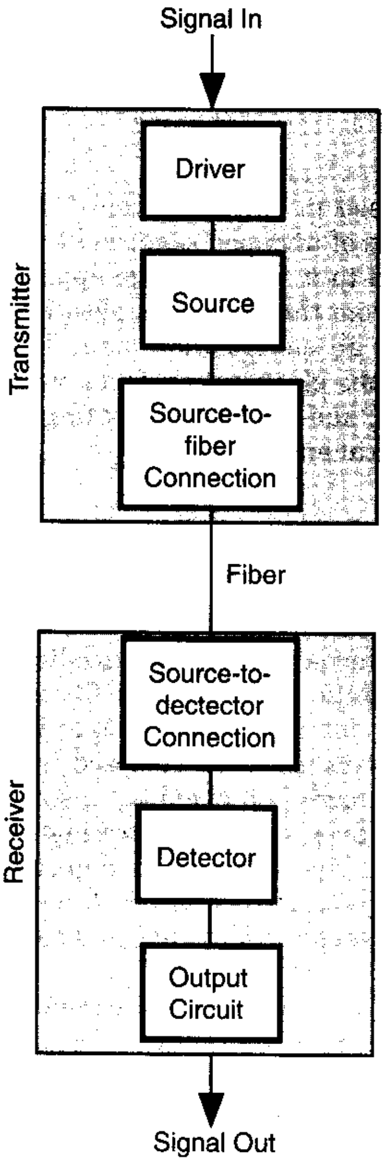

Figure 3–2 The basic fiber setup.

high energy level, and an unstable state. The electron will return spontaneously (in a few pico seconds) to a stable state, and will emit a photon in doing so. The optical wavelength is determined by the amount of energy the electron releases.

A laser operates with stimulated emission. For this process, the electron enters and stays in a high-energy state for a few microseconds. Then it changes its state spontaneously. During this state, it can be stimulated by a photon to emit its energy in the form of another photon. Thus, the laser produces the light signal which is sent into the fiber.

The Detector

Several kinds of optical detectors are used in optical systems. The photoconductor is a good example. It is a piece of undoped semiconductor material that has electrical contacts attached to it, with voltage applied across the contacts. When a photon arrives from the fiber, it is absorbed by the material, resulting in the creation of an electron/hole pair. The electron and the hole migrate toward one of the electrical contacts, with the electron attracted to the positive contact, and the hole migrating to the negative contact, thus creating the electrical energy.

STRUCTURE OF THE FIBER

Figure 3–3 (a) shows the parts that make up an optical fiber cable. The optical fiber works on the principle of total internal reflection. Once light begins to reflect down the fiber, it will continue to do so. The fiber is constructed of two layers of glass or plastic, one layer surrounding the other, as shown in Figure 3–3 (a). These layers are then enclosed in a protective jacket. The jacket surrounding the core and cladding is some type of polymer protective coating. The inner layer, the *core*, has a higher refractive index (n_1) than the outer layer (n_2), the *cladding,* as depicted in Figure 3–3 (b). Light injected into the core and striking the core-to-cladding interface at greater than the critical angle will be reflected back into the core.

The overall size of the optical fiber is quite small. As one example, in a single mode fiber (explained shortly), the core is about 10 micrometers in diameter (0.000010m). The cladding is around 125 mmeters (0.000125 mm).

Angles

Figure 3–4 shows several aspects of the fiber and the optical signal. Two aspects of the signal's propagation are of interest here: the angles of incidence (identified as θ_1) and reflection (identified as θ_{1r}). The light signal propagates down the fiber going through a series of reflections off the cladding, back into the core, then to the cladding, then back into the core, and so on.

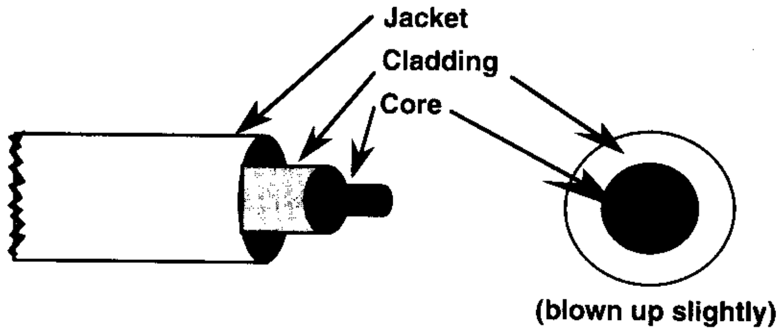

(a) Structure of Optical Fiber

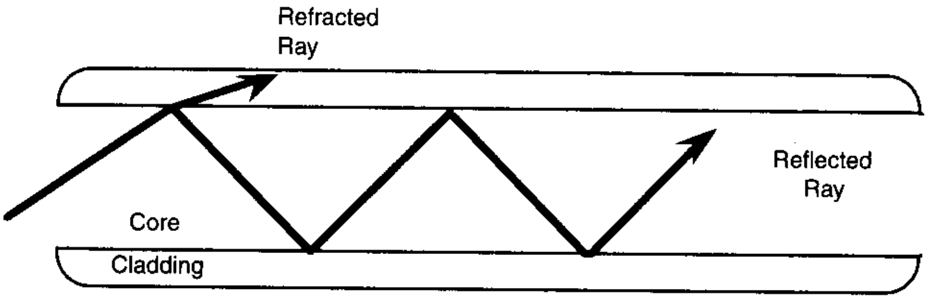

(b) Signal Propagation

Figure 3–3 Structure of the fiber.

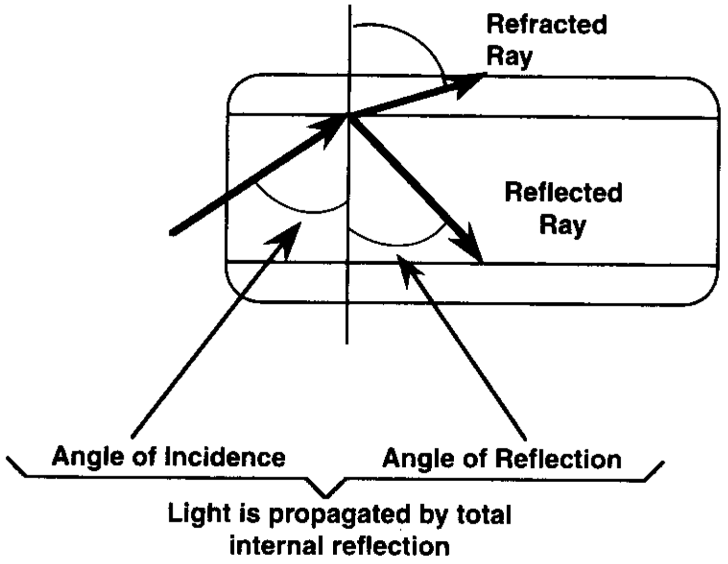

Figure 3–4 Total internal reflection.

Since the angles of incidence (θ_1) and reflection (θ_{1r}) are equal (a law of geometrical optics), the light will continue to be reflected down the core. This concept is not unlike a very long and very narrow billiards table. A ball shot at an angle into the cushion will bounce off at the same angle and continue bouncing from cushion to cushion down the length of the table, always at the same angle. Light striking the interface at less than the critical angle will pass into the cladding and will be absorbed and then dissipated by the jacket. (However, a significant amount of the energy of the light can propagate down the cladding.)

As the angle of incidence increases, so too does the angle of reflection. For larger values of θ_1, there is no refracted ray; thus, all the energy from the incident ray is reflected, and is called total internal reflection. The angle of refraction is the angle of the refracted signal (θ_2). To summarize this verbal description [RAMA98]:

$$\theta_{1r} = \theta_1$$
$$n_1 \sin\theta_1 = n_2 \sin\theta_2$$

FIBER TYPES

This part of the chapter examines the optical fiber technology in more detail, and I highly recommend [REFI99], if you want more details on this subject. Several of the examples in this section are sourced from this reference. I have already given my recommendation for [DUTT98] (see Footnote 1). Two other fine technical references on this subject are [ARGA92] and [RAMA98]. Any examples that are sourced from these references will be cited as appropriate.

Let's examine the concept of *mode*. See Figure 3–5. In its simplest form, a mode is a path that a light signal takes though a fiber. A multimode fiber is one in which the guided light ray takes different paths through the cable. A single-mode fiber is one in which the guided light ray takes one path through the cable.

The simplest fiber type is the multimode fiber. This fiber has a core diameter ranging from 125–400+ μm (.005 to over .016 of an inch), which allows many modes, or rays of light propagation. The most common multimode fiber deployed in the late 1970s and early 1980s had a core diameter of 50 μm. This fiber was used on interoffice and long-distance trunks. Today, multimode fiber is used almost exclusively for private premises systems in which distances are usually less than 1 km.

The larger core diameter in multimode fiber permits more modes. Since light reflects at a different angle for each mode, some rays will

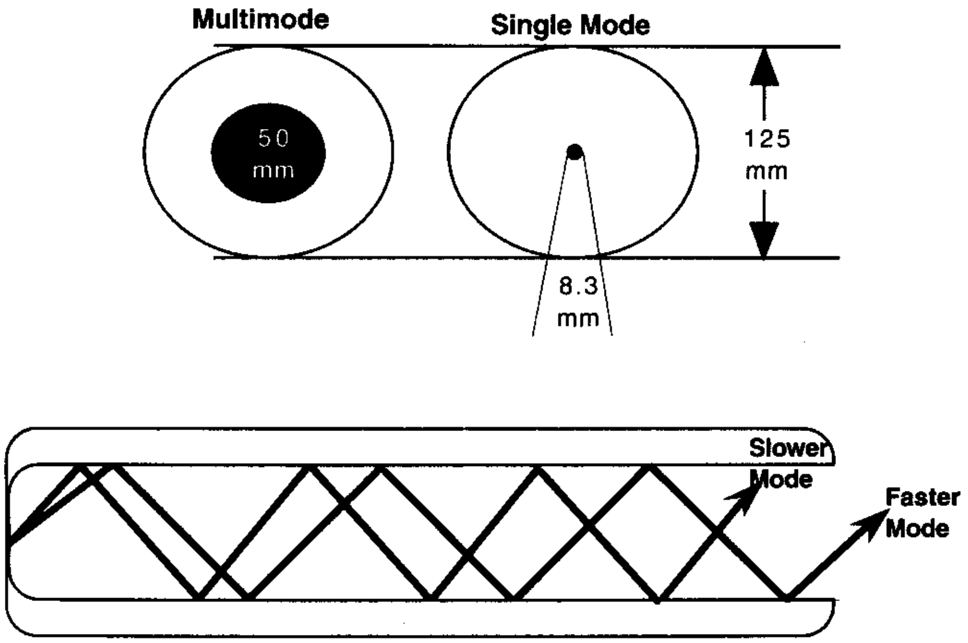

Multimode Fiber

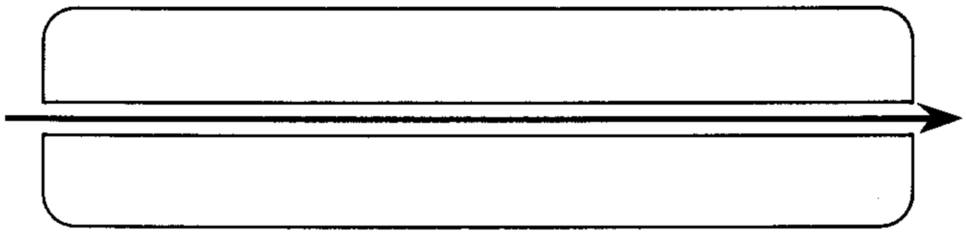

Single-mode Fiber

Figure 3–5 Basic fiber types.

follow longer paths than others; that is, each mode travels down the path at slightly different velocities.

The ray that goes straight down the core without any reflecting will arrive at the other end sooner than other rays. Other rays arrive later, and the more times a ray is reflected, the later it arrives. Thus, light entering the fiber at the same time may arrive at the other end at slightly different times. When this phenomenon occurs, we say that the light has "spread out." This spreading of light is called *modal dispersion*. The effect of modal dispersion is that the signal (the digital pulse) becomes smeared as it travels down the fiber.

A way to eliminate modal dispersion is to reduce the core's diameter until the fiber will propagate only one mode efficiently. In effect, the energy of the light signal travels in the form of one mode. This approach is used in the single-mode fiber, which has a core diameter of only 2–8 μm, with the most common diameter of about 8.3 μm. These fibers are by far the most efficient. Their advantages are minimum dispersion, high efficiency, and high operating speeds.

Single-mode fiber eliminates modal dispersion, but another potential problem may occur. It is called *chromatic dispersion*. The idea of chromatic dispersion is the same as modal dispersion. Even in single-mode fiber, different frequency components (wavelengths) travel at different speeds in the fiber. We will have more to say about chromatic dispersion shortly.

KEY PERFORMANCE PROPERTIES OF FIBER

Optical fiber's key performance properties discussed in this section are (a) attenuation, (b) chromatic dispersion, and (c) polarization mode dispersion. We start with a look at the decay of the signal's strength, called attenuation.

ATTENUATION

Figure 3–6 shows the spectral attenuation performance of a typical silica-based optical fiber. The attenuation decreases with increasing wavelength, except at wavelengths above 1.6 μm, due to bending-induced loss and silica absorption. Also, notice the three peaks in the figure. They are attenuation absorption peaks associated with the hydroxyl ion (OH-). The figure also shows five wavelength windows (labeled 1st through 5th), which represent implementations of several optical fiber wavelengths.

If low power levels are introduced into an optical fiber, the signal propagation behaves as if the medium were linear. This means that the power of the signal determines attenuation and the refractive index. However, high power levels have the effect of creating nonlinear effects; they are dependent upon the power level itself. Thus, at higher power levels, optical fiber exhibits nonlinearities. These nonlinearities, known

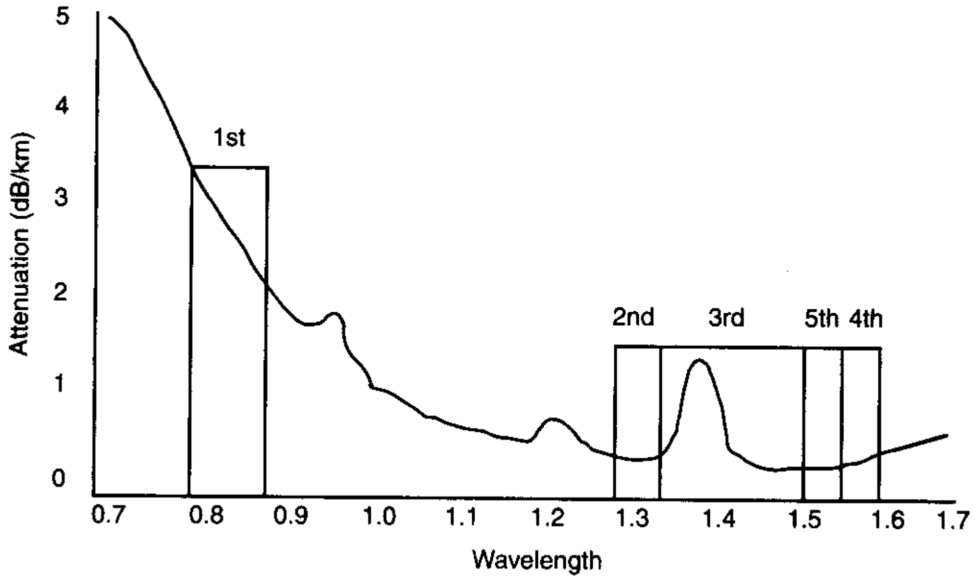

Figure 3–6 Attenuation performance of silica-based optical fiber [REFI99].

as *thresholds* when they reach the point of manifesting themselves, place limitations on the transmission capacity of optical fiber systems.

Early systems operated in the first window using multimode fiber at about 0.85 μm, and later in the second window at about 1.3 μm. Initially, single-mode fiber also operated in the second window with attenuation at about 0.35 dB/km or less. In the late 1990s, lasers supporting longer wavelengths (1.55 μm) became available operating in the third window with attenuation at around 0.2 dB/km.

The third window is where DWDM operates today, and commercial systems will soon support transmissions in the fourth window. The fifth window may come into use in the future.

AMPLIFIER SPONTANEOUS EMISSION (ASE)

Due to the attenuation of the optical signal, optical amplifiers must be placed on the link to boost the strength of the signal. This situation results in the following effects [CHIU01]. First, amplifiers do not perform perfectly, and during their amplification operation, they emit power not associated with the signal. This effect is called amplifier spontaneous emission (ASE).

ASE degrades the optical signal-to-noise ratio (OSNR). An acceptable optical SNR level (SNRmin) that depends on the bit rate, transmitter-receiver technology, and margins allocated for the impairments, needs to be maintained at the receiver. Vendors often provide some general engineering rule in terms of maximum length of the segment and number of spans. For example, current transmission systems are often limited to up to 6 spans each 80km long.

Assume that the average optical power launched at the transmitter is P. The lightpath from the transmitter to the receiver goes through M optical amplifiers, with each introducing some noise power. Unity gain can be used at all amplifier sites to maintain constant signal power at the input of each span to minimize noise power and nonlinearity. A constraint on the maximum number of spans can be obtained with well-known and published equations. If you want to delve into these details, refer to [KAMI97]. Some examples of implementations of fiber spans are discussed in Chapter 7.

CHROMATIC DISPERSION

Chromatic dispersion was introduced earlier. Recall that it is the effect of the different wavelengths traveling at different speeds in the fiber. In effect, the spectral components of the pulse travel at different speeds.

The first component of chromatic dispersion is called *material dispersion*. It arises because the refractive index of silica is frequency-dependent; that is, the dispersive characteristics of the dopants and the silica mean that the frequency components of the signal travel at different velocities.

The second component is called *waveguide dispersion*. It deals with the fact that the refractive index (the refractive indices of the cladding and the core) changes with wavelength. Waveguide dispersion is explained as follows [RAMA98]:

- The energy of the light propagates in the core and in the cladding.
- The effective index of a mode lies between the refractive indices of the core and the cladding.
- The value of the index between the two depends on the proportion of power that is in the core and the cladding. If most of the power

is in the core, the effective index is closer to the core refractive index; otherwise, it is closer to the cladding reflective index.

- The power distribution of a mode between the core and the cladding is a function of the wavelength.
- If the wavelength changes, the power distribution changes, resulting in the effective index or propagation constant of the mode to change.

Figure 3–7 shows the dispersion material component and two types of waveguide components: dispersion-unshifted fiber (USF) and dispersion-shifted fiber (DSF). The DSF component includes a variation of DSF, called non-zero-dispersion fiber (NZDF). These fibers are explained next, with references back to this figure.

Chromatic dispersion is measured in units of ps/nm-km. The term ps refers to the time spread of the pulse; nm is the spectral width of the pulse, and km is the length of the fiber cable.

A method to combat chromatic dispersion is called *dispersion-shifted* fiber (appropriate for single-mode fiber). It is designed to exhibit zero dispersion in the 1.55 μm wavelength.

Figure 3–8 compares the chromatic dispersion characteristics of three different single-mode fibers:

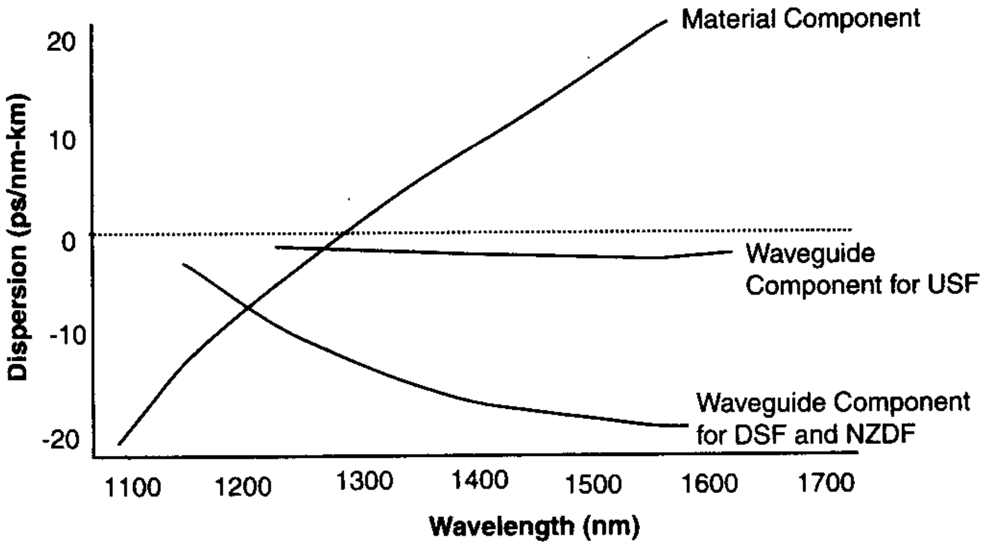

Figure 3–7 Material and waveguide dispersion components [REFI99].

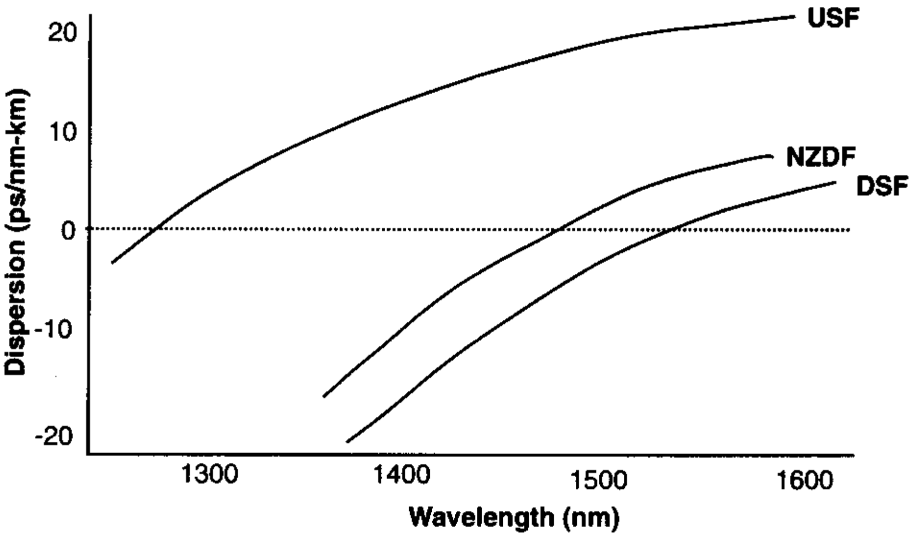

Figure 3–8 Chromatic dispersion versus wavelength [REFI99].

- *Dispersion-unshifted fiber (USF):* Zero-dispersion wavelength resides near 1.31 μm.
- *Nonzero-dispersion fiber (NZDF):* Zero-dispersion wavelength resides near 1.53 μm (may range to 1.565 μm). The initials for nonzero-dispersion fiber are also noted as NDF.
- *Dispersion-shifted fiber (DSF):* Zero-dispersion wavelength resides near 1.55 μm.

This discussion leads us to a more detailed explanation of the types of optical fibers in use. The International Telecommunication Union (ITU) has published recommended standards for four types of single-mode fibers.

Dispersion-unshifted fiber (USF) is an older technology that was introduced in 1983 (ITU G.652). It exhibits zero chromatic dispersion at 1310 nm. USF is very widely-used in the industry, and is also referred to as standard or conventional fiber. It is found as CATV operating in the second- and third-wavelength windows. It exhibits high chromatic dispersion at 1550 nm, and does not operate well at rates beyond 2.5 Gbit/s.

Dispersion-shifted fiber (DSF) was introduced in 1985 (ITU G.653). It exhibits zero chromatic dispersion at 1550 nm. Its attenuation performance (third window) is attractive. Due to its attractive attenuation and chromatic dispersion characteristics, it was considered a good candidate for 1550 nm systems. This initial view has been changed in view of the nonlinear effects of optical signals (discussed later). Notwithstanding, the technology is found in submarine systems and long-distance telephone networks. However, DSF does not perform well with multiple wavelengths in the third wavelength window, and it is being replaced by nonzero-dispersion fiber (NZDF), which will be discussed shortly.

1550 nm loss-minimized fiber is a special form of DSF and is published as ITU Recommendation G.654. It has very low loss in the 1550 nm signal. However, USF is expensive and difficult to manufacture. Its principal use has been in nonrepeated submarine systems.

Nonzero-dispersion fiber (NZDF or NDF) was introduced in the early 1990s. It is designed for DWDM systems, and it has been standardized by the ITU, as well as the Telecommunications Industry Association (TIA). It operates over a portion of the third wavelength window, with the chromatic dispersion small enough to support individual channel rates of 10 Gbit/s over distances of over 250 km. The chromatic dispersion ranges between 1 and 6 ps/nm-km in the 1550 nm wavelength window. Today, NZDF is deployed extensively in submarine as well as long-haul terrestrial networks.

Polarization-mode Dispersion (PMD)

The third property (actually, impairment characteristic) of optical fiber is called *polarization-mode dispersion* (PMD; see Figure 3–9). This problem occurs because the fiber is not consistent along its length. Due to bending and twisting, as well as temperature changes, the fiber core is not exactly circular. The result is that the modes in the fiber exchange power with each other in a random fashion down the fiber length, which results in different group velocities; the signal breaks up. In effect, the light travels faster on one polarization plane than another.

The distribution of the energy of the signal over the state of polarizations (SOPs)[1] changes slowly with time [RAMA98]; that is, the instantaneous values conform to a Maxwellian function whose mean value increases with the square root of the fiber's length:

---

[1] The SOP refers to the distribution of light energy in two polarization modes. This idea is found in any electromagnetic signal. Light energy (the pulse), traveling down the fiber that is uniform along its length, is divided into two polarization modes, one on the x-axis and one on the y-axis. In ideal conditions, the two modes have the same propagation constant.

$$(\Delta\tau) = D_{PMD}\sqrt{L}$$

where $\Delta\tau$ is the time-averaged differential time delay, L is the length of the length, and D_{PMD} is the fiber PMD parameter, measured in ps/$\sqrt{\text{km}}$.

For a uniform fiber, the time spread is constant. In reality, the light energy of the pulse, propagating at different velocities, leads to pulse spreading, or PMD. The two modes are on the x-axis (slow mode) and y-axis (fast mode), as shown in this figure. The right side of the figure shows the PMD experienced with non-consistent fiber.

The various standards groups are now publishing parameters on PMD. The International Electrotechnical Commission (IEC) is working on a standard under IEC SC86A Working Group 1.

It should be noted that PMD might (depending on who is talking) be a problem in systems that use older optical fiber. For example, problems can occur using older fiber with transmissions rates in the 10 Gbit/s range, and some in the industry think that running high rates (such as 40 Gbit/s) is not operationally feasible. As illustrationed, [CHIU01] defines the following PMD performance parameters. For older fibers with a typical PMD parameter of 0.5 picoseconds per square root of km, based on the constraint, the maximum length of the fiber segment should not exceed 400 km and 25 km for bit rates of 10 Gbit/s and 40 Gbit/s, respectively. For newer fibers with a PMD parameter of 0.1 picosecond per square root of km, the maximum length of the transparent segment (without PMD compensation) is limited to 10000 km and 625 km for bit rates of 10 Gbit/s and 40 Gbit/s, respectively. In general, the PMD re-

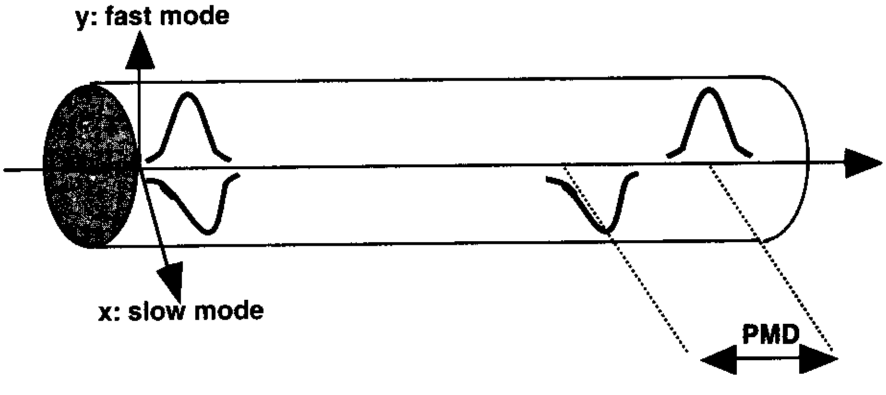

Figure 3–9 PMD [REFI99].

quirement is not an issue for most types of fibers at 10 Gbit/s or lower bit rate.

Components are available to reduce the effect of PMD. They are called PMD compensators, and are now available for rates as high as 40 Gbit/s. But they cost money, and some companies think the problem is overblown. As of this writing, the issue of PMD and how serious a problem it is has not been resolved. Many of the solutions are still in the vendors' labs. So, stay tuned, and stay alert to the issue.

LASERS

Lasers were introduced briefly at the beginning of this chapter. The term laser is a shortened form for "light amplification by stimulated emission of radiation." Lasers are used in high-capacity optical networks to send the optical signals to the fiber. They are very small semiconductor devices designed to transmit very specific and precise wavelengths. These transmitters operate on the principal of an "excited state," which means that the electrons in certain parts of the semiconductor material have more energy than other electrons. When the electron loses some of its energy and falls to a ground state, the energy released is in the form of photon, light energy. Appling electrical current to the laser produces many electrons in an excited state. Of course, as these many electrons decay to a ground state, they give off light energy in the form of many photons.

SUMMARY

While the optical fiber medium is clearly superior to copper and wireless media in regards to bandwidth and signal quality, optical signals still suffer from transmission impairments. They are broadly classified as attenuation, chromatic dispersion, and polarization-mode dispersion. The effects of these problems depend on the type of fiber used, as well as other factors such as the wavelength(s) transmitted on the fiber, the spacing (and nature) of the amplifiers, and the condition of the fiber cable.

4

Timing and Synchronization

This chapter explains the synchronization and clocking functions used in optical networks. Asynchronous and synchronous networks are examined. Clock variations and controlled and uncontrolled clocking slips are analyzed, and compared to each other. After these subjects are covered, clock distribution systems are examined. The chapter concludes with a discussion of synchronization messages.

TIMING AND SYNCHRONIZATION IN DIGITAL NETWORKS

With the advent of digital networks and the transmission and reception of binary pulses (1s and 0s), it became important to devise some method for detecting these signals accurately at the receiver. Figure 4–1 illustrates the problem.

The ideal system is one in which the binary pulses arrive at the receiver in a very precise and concise manner. This means that the receiver knows the exact time that the signal (a binary 1 or 0) manifests itself at the receiver interface. This synchronization between the transmitter and the receiver is achieved because each machine knows about the other's "clock," that is, at which frequency the sending machine sends its traffic to the receiving machine. Fortunately, it is a relatively easy task to determine this timing, because the receiver can derive (extract) the clock from the incoming bit stream by examining when the pulses arrive at the receiver. For example, in Figure 4–1(a), the signal and the clock are perfectly aligned when the signal reference mark occurs with the zero crossing in the physical clock wave form. However, errors can occur if the clock is not aligned with the signal, as seen in Figure 4–1(b). This problem is usually called phase variation, and it may translate into an incorrect interpretation of the binary 1s and 0s in the transmission stream.

Therefore, it is not enough that signals be aligned in the same frequency domain (the same rate of ticking of the clocks). The signals must also be aligned in the phase domain (the same instant that the clocks emanate the tick).

In older systems that operated at a relatively low bit rate, the clocking did not have to be very accurate because the signal on the line did not change very often. As the digital networks became faster and more bits were transmitted per second, the time the bits were on the channel decreased significantly. This meant that if there was a slight inaccuracy in the timing of the receiver's sampling clock, it might not detect a bit or, more often, it might not detect several bits in succession. This situation leads to a problem called slips. Slipping is the loss of timing and the resultant loss of the detection of bits.

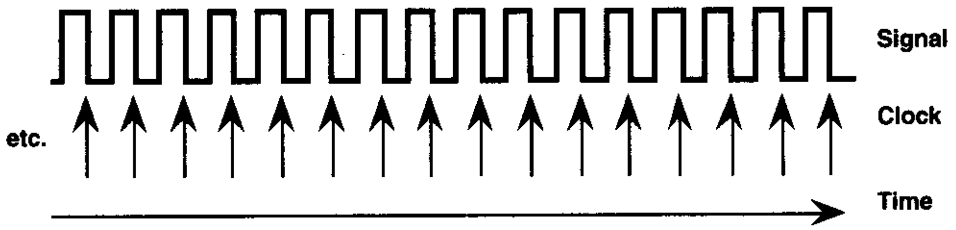

(a) Clocking is aligned

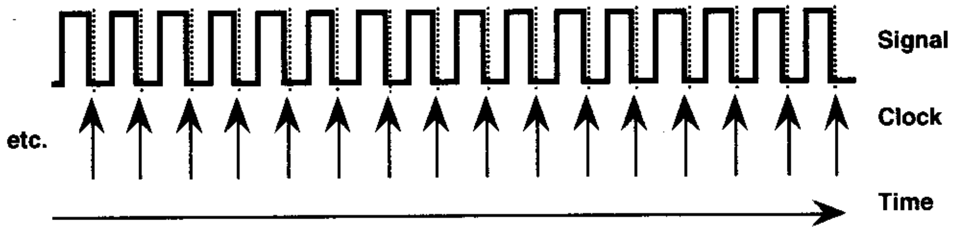

(b) Clocking is not aligned

Figure 4–1 Clocking and phase variation.

EFFECT OF A TIMING ERROR

Without the use of clocking mechanisms, an error will be passed to customers, perhaps through multiple nodes in the network. This creates a ripple effect throughout the network, but its effect depends on the type of traffic being transmitted. The typical impact of clocking errors on different types of user traffic is:

Type of Traffic	Result of Error
Digital data	Reduced throughput
Encrypted data	Resend key
Fax	Missing lines/distorted page
Video	Picture freeze, or dropouts
Voice	Audible click
Voice band data	Carrier drop

THE CLOCKING SIGNAL

In its simplest terms, a clocking signal carries information about time. This information is represented by the signal crossing a reference mark. In Figure 4–1, for example, when the signal crosses a certain voltage level, this crossing allows the receiver to determine when the clocking signals occur.

Different terms are used to describe the clocking signal. The one used in most of the literature is the unit interval (UI), which corresponds to one cycle of the clocking signal. Another term is the phase, which is measured in radians (1 UI = 2 radians). UI is equal to the reciprocal of the data rate. As examples, one UI for a DS1 1.544 Mbit/s rate is 648 nsec and one UI for a DS3 44.736 Mbit/s rate is about 22 nsec.

TYPES OF TIMING IN NETWORKS

The systems that existed in the 1950s and early 1960s were not synchronized to any common clocking source because they consisted of analog circuits and did not need a precise timing setup. However, as digital networks were deployed, and especially with the advent of the T1 technology, timing became a greater concern.

These early digital networks were not synchronized to a common frequency, and thus they were called asynchronous networks. Each machine in the network ran its own "free-running clock," and the clocks between two machines could vary by many unit intervals.

Today, optical networks are called synchronous networks because the timing is tightly controlled. In reality, they are actually plesiochronous networks. The prefix plesio means "nearly." Each portion of the network, as shown in Figure 4–2, is synchronized to a highly accurate primary reference source (PRS) clock. Because of the superior level of performance and the fact that this technique is fairly inexpensive, PRS clocks are a cost-effective way to improve network performance. In Figure 4–2, a portion of this network is referenced with primary reference source x (PRS x) and another portion is synchronized with PRS y. Thus, the term plesiochronous distinguishes this type of network from a truly synchronous network, which has only one PRS. The distinction between synchronous and plesiochronous is not used much today, so we will use the term synchronous throughout the book.

Ideally, a single synchronous network uses one PRS, which is also known as the master clock. As shown in Figure 4–3, the components derive their clocking from this master clock. Timing is derived first from the master clock, and then from a slave (in this example, a toll office).

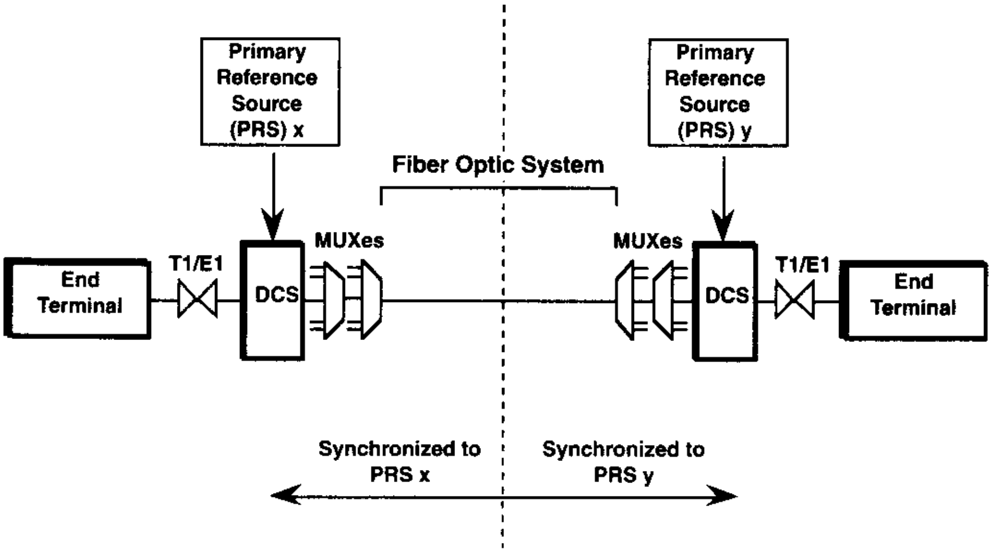

Figure 4–2 Synchronous networks.

Then, the timing is passed to digital switches, digital cross-connects, end offices, etc. Therefore, timing is "cascaded down" to other equipment, such as channel banks and multiplexers.

Figure 4–3 also shows the employment of different types of clocks called stratum n clocks. Each stratum n clock is required to perform within a certain degree of accuracy. The stratum 1 clock must meet the most stringent timing requirements, whereas the stratum 4 clock needs to meet only the least stringent requirements.

The Synchronous Clock Hierarchy

Table 4–1 summarizes the synchronous network clock hierarchy and shows long-term accuracy for each stratum level, as well as typical locations of the clocking operations. Long-term accuracy for the stratum 1

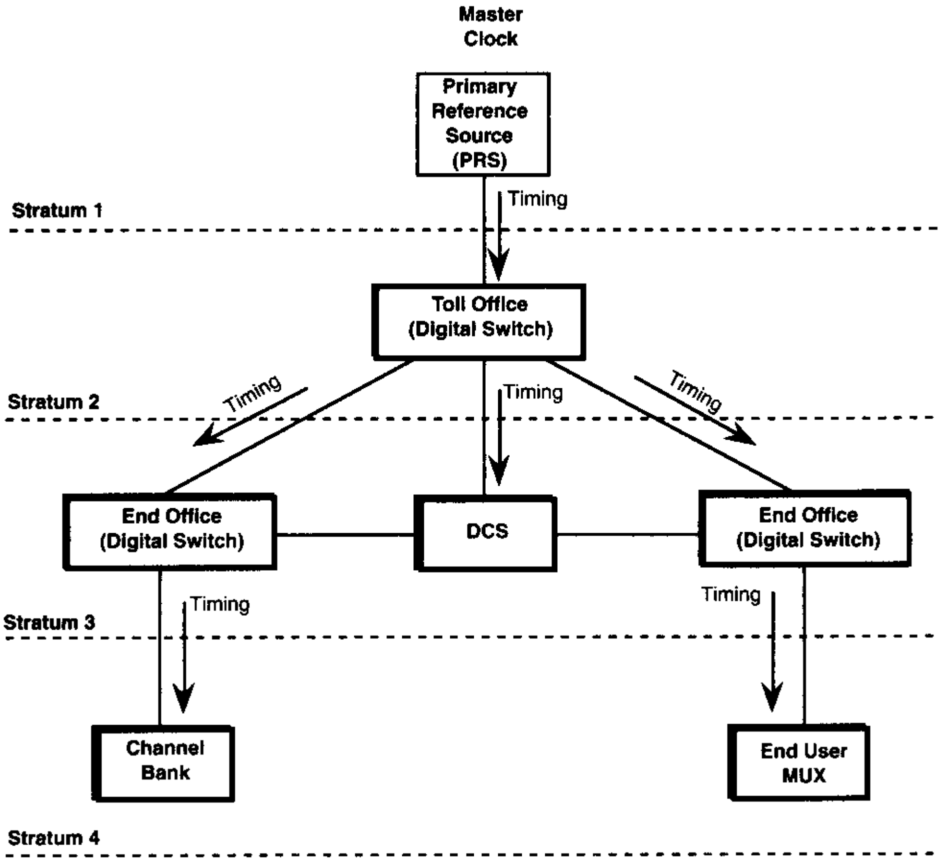

Figure 4–3 Synchronous clocking hierarchy.

Table 4–1 Clock Hierarchy for Synchronous Networks

Clock Stratum	Typical Location(s)	Free-Run Accuracy (Minimum)
1	Primary Reference Source (PRS)	$\pm 1.0 \times 10^{-11}$*
2	Class 4 office	$\pm 1.6 \times 10^{-8}$
3 and 3E[1]	Class 5 office, DCS	$\pm 4.6 \times 10^{-6}$
4 and 4E[1]	Channel bank, end-user Mux	$\pm 32 \times 10^{-6}$

* = Also annotated as .00001 ppm (parts per million)

[1] Stratum 3E and 4E clocks are not part of ANSI standards (ANSI T1.101). They are used by (the former) Bellcore (GR-1244-CORE) and stipulate more stringent requirements with regard to wander and holdover (a loss of a previously connected external reference). These enhanced clocks are compatible with the ANSI T1.101 clocks.

clocks is ± $1.0 * 10^{-11}$. The next level of accuracy is the stratum 2 clock, which is usually located in class 4 toll offices. The long-term accuracy for these clocks is ± $1.6 * 10^{-8}$. Next in the order of accuracy are the stratum 3 clocks, typically located in the class 5 end office or a digital cross-connect (DCS). The long-term accuracy of these clocks is ± $4.6 * 10^{-6}$. The last level of the synchronous network clock hierarchy is the stratum 4 clock. These clocks are usually located in channel banks or multiplexers at the end-user site. Their accuracy is ± $32 * 10^{-6}$

TIMING VARIATIONS

While synchronous networks exhibit very accurate timing, some variation will exist between the network elements within a network as well as network elements between networks. This variation is generally known as phase variation.

Phase variation is usually divided into *jitter* and *wander*. Jitter is defined as a short-term variation in the phase of a digital signal, which includes all variations above 10 Hz. In effect, jitter is the short term for variation of the digital signal's optimal position in time. Causes of jitter include common noise, the bit stuffing processes in multiplexers, or faulty repeaters.

In contrast, wander is the long-term variation in the phase of a signal and includes all phase variations below 10 Hz. Wander may also include the effects of frequency departure, which is a constant frequency difference between network elements. Wander is almost inevitable in any network due to the slight variations in clock frequency differences, transmission delay on the path, or bit stuffing operations.

Jitter and wander are dealt with in many digital networks through the use of buffers. These buffers exist at each interface in any machine where the signal is processed (multiplexed, switched, etc.). Buffers act as windows to receive and transmit traffic. Additionally, for digital systems, they can be used to accommodate to frequency departure or phase variations. Buffers are carefully designed to handle the most common variations.

Frequency Accuracy

Most systems in North America describe clocking accuracy as the degree that a clock's frequency deviates from its ideal value. It is defined as:

$$FF_{os} = (f - f_d)/f_d$$

where FF_{os} = fractional frequency offset, f = actual frequency output of a clock, and f_d = ideal or desired frequency.

In addition to the requirements cited above, Bellcore establishes the following requirements for holdover stability and the pull-in/hold-in range. The legend in Box 3–1 explains the column entries in Table 4–2.

Table 4–2 Other Timing Requirements

Stratum	Holdover Stability	Pull-in/Hold-in Range
1	N/A	N/A
2	$\pm 1 * 10^{-10}$/day	$\pm 1.6 * 10^{-8}$
3E	$\pm 1 \times 10^{-10}$/day*	$\pm 4.6 * 10^{-6}$
3	<255 slips ($\pm 3.7 \times 10-^{7}$)*	$\pm 4.6 * 10^{-6}$
4E	NA	$\pm 32 * 10^{-6}$
4	NA	$\pm 32 * 10^{-6}$

* For initial 24 hours

Box 4–1 Legend for Table 4–2

Holdover:	Operating condition of a clock once it has lost a previously connected clocking reference. During this time, the machine must exhibit the stability cited in this column.
Pull-in range:	The largest band of input signal frequency for which the clock will acquire lock. Assures that synchronization can be achieved with a clock of equal stratum level that may be operating at the limits of its permissible frequency offset, while the clock-under-test is operating at the opposite frequency offset limit.
Hold-in range:	The largest band of input signal frequency for which the clock will maintain lock. Specified so that a clock of a given stratum level will be able to maintain lock with a reference from a clock of the same stratum level as the upstream clock varies in frequency.

METHODS OF CLOCK EXCHANGE

Clearly, it is in the best interest of all concerned to use a common clocking source for all machines in the network. Some systems use this approach and some systems do not. This part of the chapter expands on our earlier discussions about clocking and describes several methods of achieving synchronization between machines.

Five methods of clock exchange can be employed in a network. They are as follows: (a) free-running, (b) line-timed, (c) loop-timed, (d) external, and (e) through-timed. Some systems use a combination of these methods. The reader should check vendors offerings carefully, because each vendor probably uses these operations in slightly different ways. In addition, the design of each network element may place limitations on how some of the clocking distributions operations are implemented.

Free-Running

The free-running method has each machine generate its own timing from a (highly stable) crystal oscillator. In most systems, this oscillator has a long-term accuracy better than ±4.6 ppm. Figure 4–4 shows a free-running/free-running configuration. We use the term free-running twice to denote that both machines on the line are running with an oscillator.

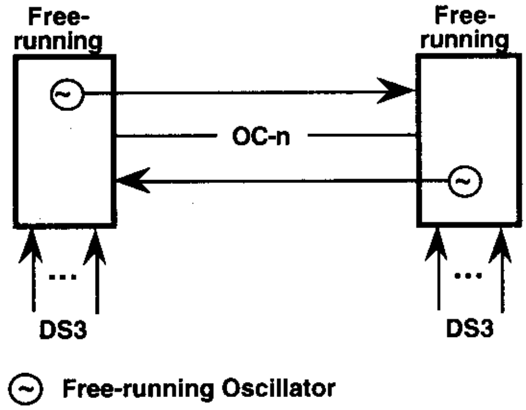

Figure 4–4 Free-running/free-running configuration.

With this approach, no external clocking source is used, which has its advantages and disadvantages. The advantage is obviating the expense of connecting to an external timing source. The disadvantage is the requirement for buffers to compensate for the delta between clock differences (i.e., an incoming DS3 signal and the outgoing signal to the other machine). Nonetheless, this approach is quite effective for point-to-point linear configurations with asynchronous (in this example, DS3) interfaces.

Line-Timed

The line-timed mode derives clocking from the signal on the incoming line. The clock extraction from this signal is fed into the local timing generator module, which, in turn provides the timing to the outgoing signals. This configuration is shown in Figure 4–5. The line-timed mode is simple, but it does not perform very well in configurations where several machines are connected linearly to each other in a path. Clocking inaccuracies tend to accumulate at each node and can lead to distorted signals.

Loop-Timed

Loop-timed mode is also called gateway or master/slave mode. It is used in systems where different timing generator modules are employed or where machines tied to different stratum levels must interact with each other.[2] The frequency sent from the master unit (in Figure 4–6 (a),

---

[2] Bellcore defines loop-timing as a timing mode for nodes that have only one synchronous interface. Therefore, it is a special case of line-timing. We consider this definition helpful but too restrictive.

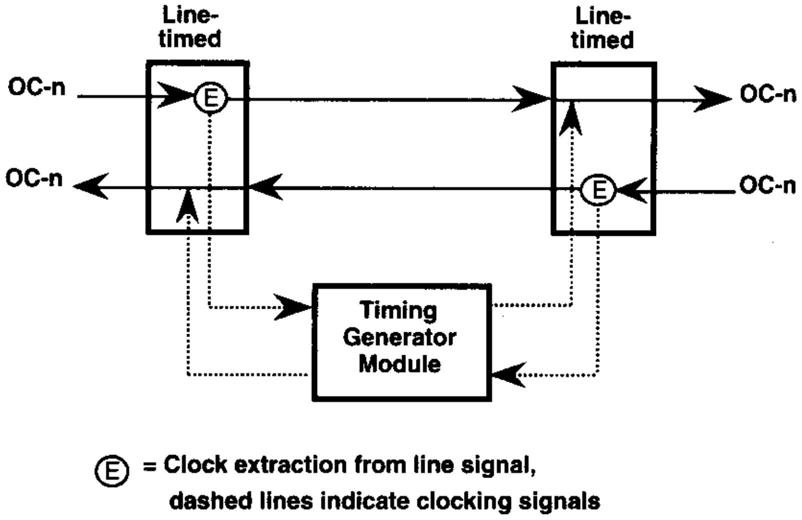

Figure 4–5 Line-timed configuration.

the free-running unit) is used to derive the clock at the slave unit (in Figure 4–6(a), the loop-timed unit). The slave unit loops the receive clock back across the line as the transmit clock. The loop-timed mode may also employ an external clock, as shown in Figure 4–6(b). The external clock is typically a stratum 3 level clock or better. This configuration in this figure is also called the phase locked/loop-timed mode. Another configuration is shown in Figure 4–6(c). The slave unit can also furnish clocking utilizing the Building Integrated Timing Supply (BITS, described later in this chapter).

External

We introduced the external clocking mode in the previous section. As the name implies, the machines time their transmitted signals from internal oscillators that are locked to an external clocking source. This configuration requires local office clocks at each end terminal, so it is used in many interoffice applications. The clock sources illustrated in Figure 4–7 must be stratum 3 or better, and they may emanate from more than one primary reference source (that is, they may be plesiochronous systems). Figure 4–7 (b) shows that this configuration can provide synchronization outputs to the office BITS clock (or to other central office equipment).

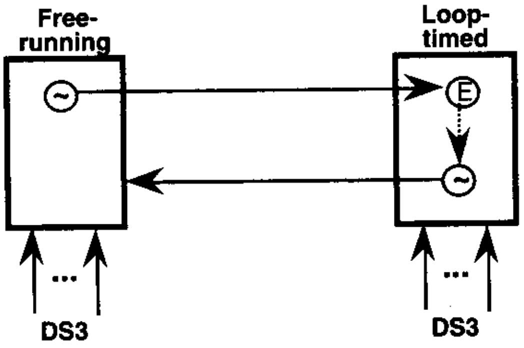

(a) Loop-timed at one end/free-running at the other end

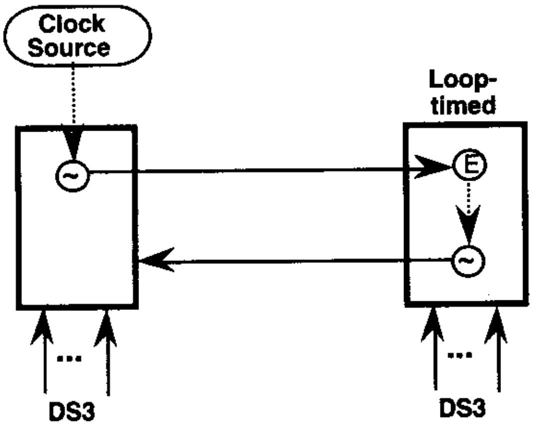

(b) Loop-timed at one end/external at the other end

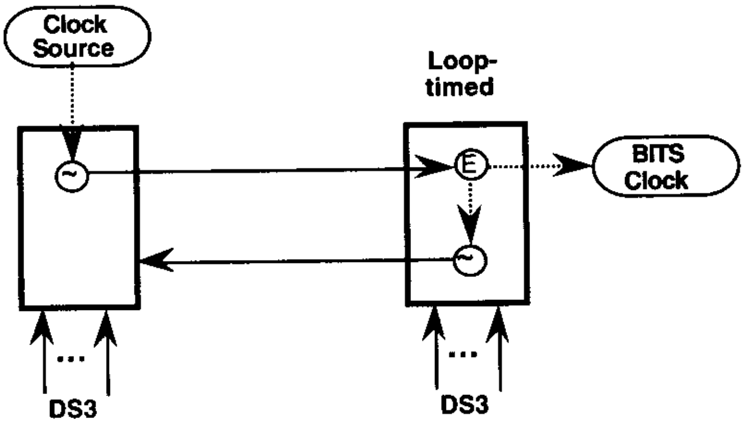

(c) Loop-timed at one end/external at the other end, and timing to BITS.

Figure 4–6 Loop-timing.

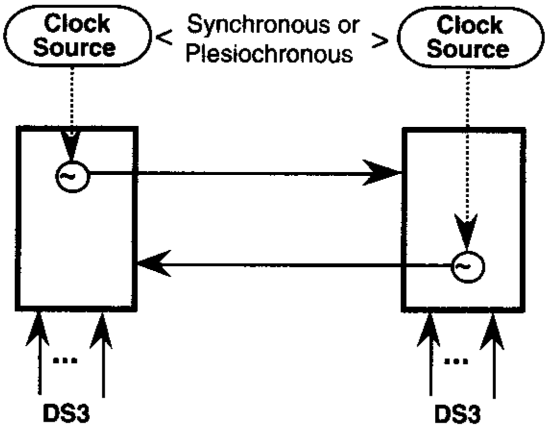

(a) External at both ends

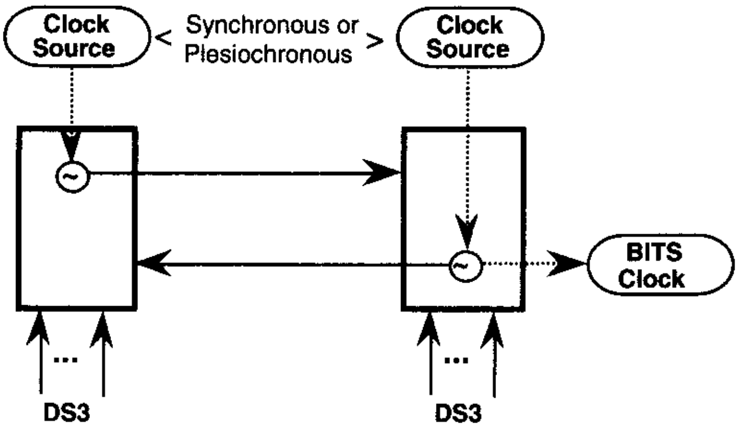

(b) External at both ends, and timing to BITS

Figure 4–7 External timing.

Through-Timed

The last example of timing distribution is called through or through-timed mode. We use a ring topology for this example. Two configurations are shown in Figure 4–8. In Figure 4–8(a), a free-running clock distributes timing signals to the other nodes on the ring. These nodes on the ring are through-timed. This term means that the network element derives its transmitted timing in the "east" direction from a received line signal in the "west" direction, and the transmit timing in the west

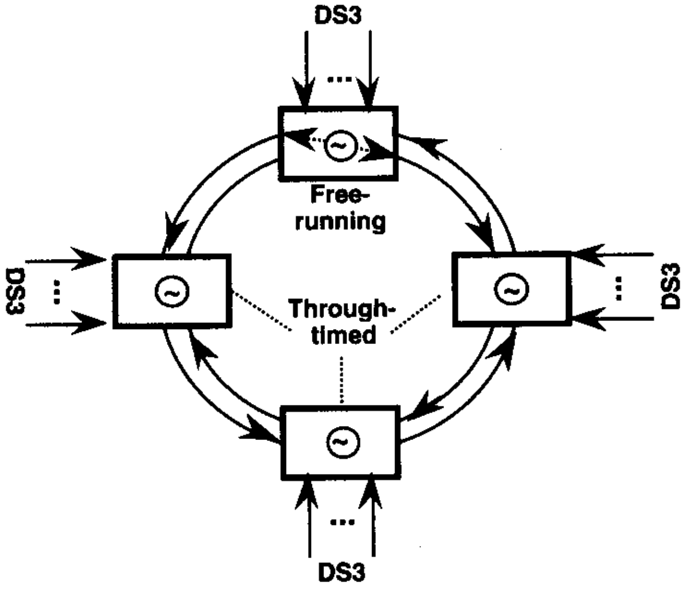

(a) Free-running clock and through-timed

Figure 4–8 (a) Through-timed.

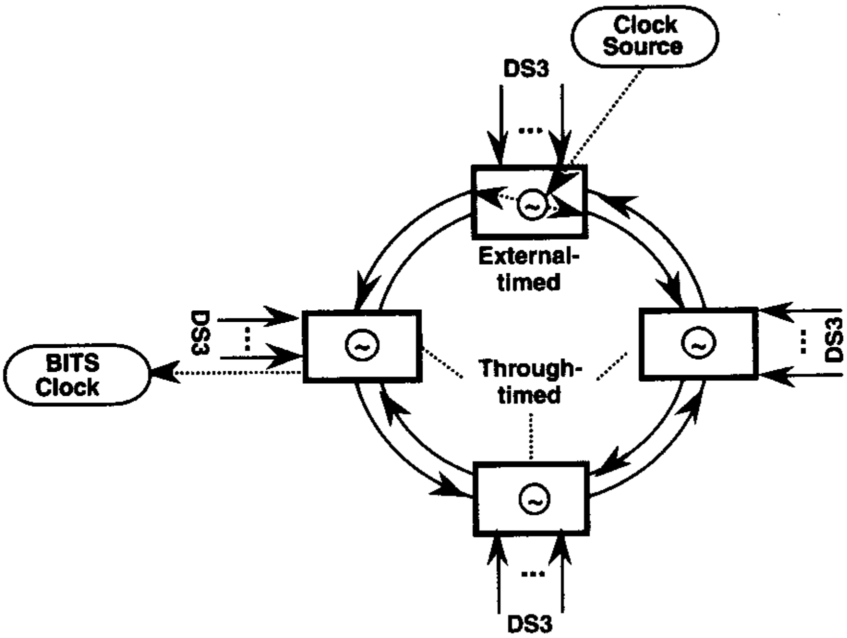

(b) External clock and through-timed

Figure 4–8 (b) Through-timed, continued.

direction from the received line signal in the east direction. Through-timing is typically employed when the network node is passing the signal transparently through the node and does not want to change the timing. Through-timing is used in signal regeneration and echo cancellers.

Timing can be provided with an external clock as well, as depicted in Figure 4–8(b). This figure shows also that timing on the ring can be distributed to BITS.

DISTRIBUTION OF TIMING USING SONET AND DS1

Figure 4–9 shows a configuration for using both DS1 and SONET to distribute timing information. DS1 has been used over the years to distribute timing information throughout the network. As seen in this figure, the information is sent between the master (source) and slave clocks. These DS1 signals may also carry traffic. As networks have been upgraded to lightwave systems, SONET has become the preferred facility to transmit these signals among COs, interoffice networks, and access networks. The incoming OCn signal provides timing information, which is traceable back to a highly accurate reference clock, exhibiting low jitter and wander. The DS1 timing signal can also be used to feed the local BITS clock, or, if BITS is not available, the DS1 signal can be provided directly to other equipment.

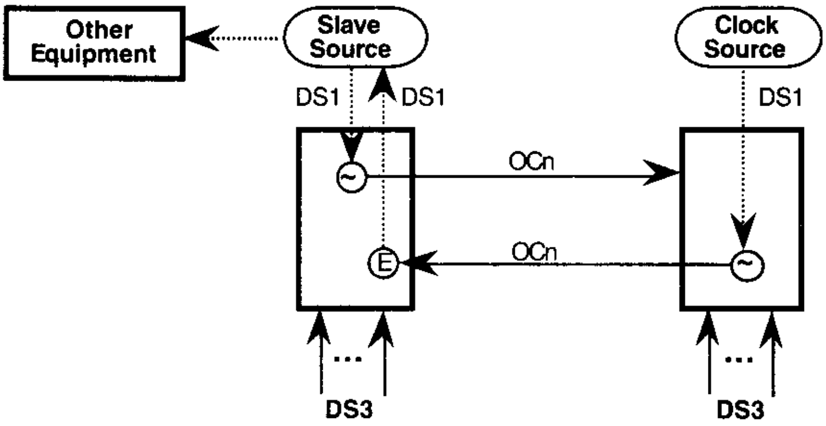

Figure 4–9 Using DS1 and SONET to distribute timing information.

TIMING DOWNSTREAM DEVICES

We have learned that synchronization can be provided with an external synchronization interface (ESI) and "cascaded" downstream to other nodes. An earlier discussion showed how this occurs in a ring configuration. Figure 4–10 shows several linear configuration options for downstream timing. These options are not all-inclusive.

Figure 4–10 depicts where timing is inserted in the line and the number of interfaces (as (n)) that are clocked by building integrated supply timing (BITS). In Figure 4–10(a), a terminal provides primary and backup BITS to 10 external interfaces, with the timing cascading down through add-drop multiplexers (ADMs) to an end terminal. Figure 4–10(b) is provided to emphasize that timing signals can be branched

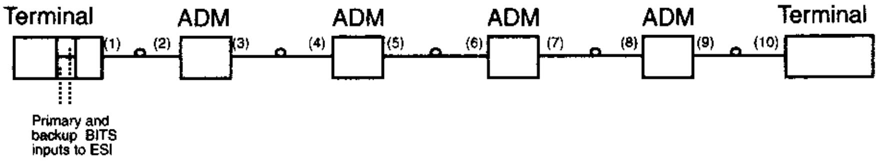

(a) Consecutive Interfaces

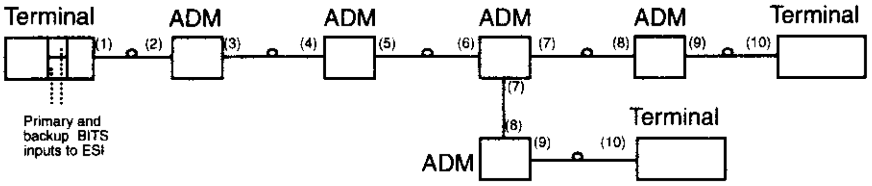

(b) Branching Topology

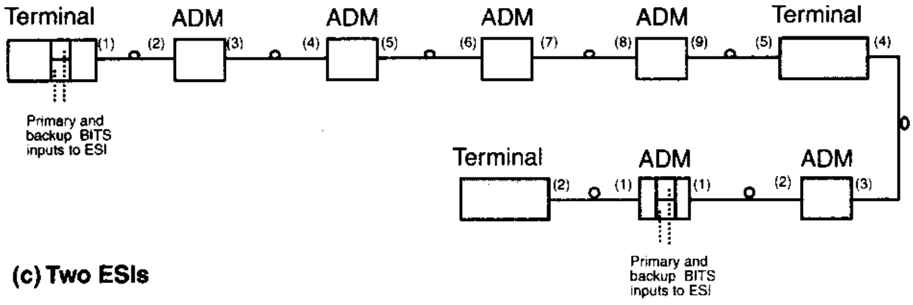

(c) Two ESIs

Figure 4–10 External synchronization interfaces (ESI).

across more than one link and to more than one end terminal. Figure 4–10(c) shows how two ESIs are terminated in one terminal. This configuration might occur if a terminal is connected to two different network providers. BITS does not define which ESI is the authoritative clock source—that issue must be resolved by the network providers. Some providers are reluctant to have their machines clocked by other providers. Others will not allow the customer to provide the clock.

THE BUILDING INTEGRATED TIMING SUPPLY (BITS)

The building integrated timing supply (BITS) is now being employed throughout the United States to provide synchronization for digital networks. BITS has two timing references: a primary source called reference A and a secondary source called reference B. Regardless of the source reference, timing must be traceable to a stratum 3 clock or better.

BITS can provide timing for a wide range of equipment, such as channel banks, cross-connects, SONET terminals, digital loop carriers (DLCs), signaling system #7 (SS7) components, asynchronous transfer mode (ATM) machines, frame relay nodes, and switched multimegabit data service (SMDS) devices.

BITS provides a composite clock for equipment with extractable timing of 64/8 kHz. It also provides DS1 timing as well as other timing, if needed. Figure 4–11 shows the 64/8 kHz composite clock. The term 64/8

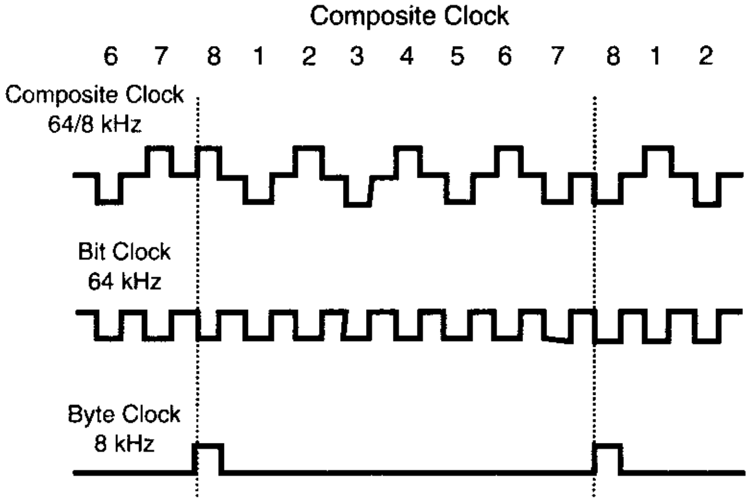

Figure 4–11 BITS clocks.

refers to a 64 kHz bit clock and a 8 kHz byte clock which can be derived from the composite clock stream. The 64 kHz bipolar clock is derived with every 8th bit bipolar violation (BPV). Additionally, the 8 kHz byte clock is derived by counting the BPVs in the byte clock.

Therefore, two types of phase criteria must be satisfied: (a) bit phase and (b) byte phase. We learned earlier that phase is any stage in a series of changes, and two signals are said to be "in phase" when the two signals are in the exact same stage; that is, at the same percent of amplitude (rising or falling) at the same time.

Figure 4–12 shows how a composite clock actually clocks data across a channel unit in an office. The system clocks data in at one end and out at the other end. Therefore, if a 1 data bit is clocked out as a 0, then the results are, of course in error. This error occurs when a transmitter's clock is out-of-phase with the receiver's clock. This situation can occur when two channel banks are referenced to two different composite clocks. While they both may be operating with the same frequency of 64 kHz, their signals are lacking phase sync.

While achieving correct bit phase synchronization is quite important, when used alone, bit phase synchronization does not guarantee accurate sync operations at the machine. Figure 4–13 shows that byte clock phase synchronization must also be provided by the network. The top part of the figure shows that the byte sent from the sender is in exact phase with the receiver's clock. Therefore, the sender and the receiver have both the same frequency and the same phase relationships. The lower part of the figure shows that the sender's and receiver's frequen-

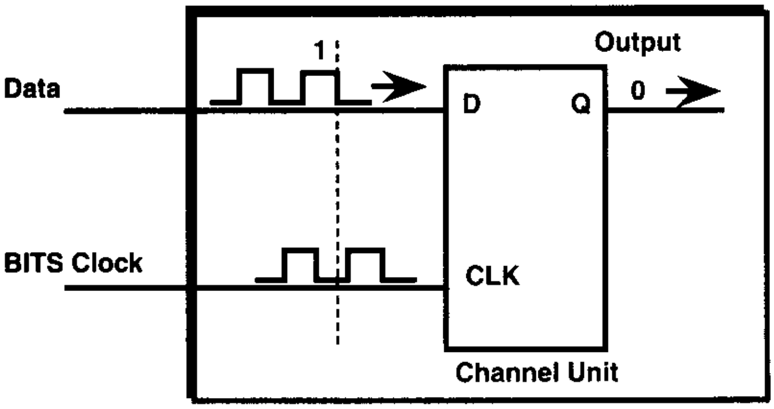

Out-of-phase Condition

Figure 4–12 Bit sync.

Both frequencies are the same and in phase:

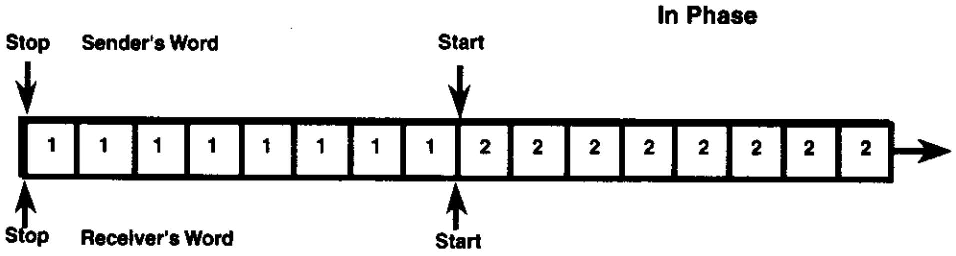

Both frequencies are the same but out-of-phase:

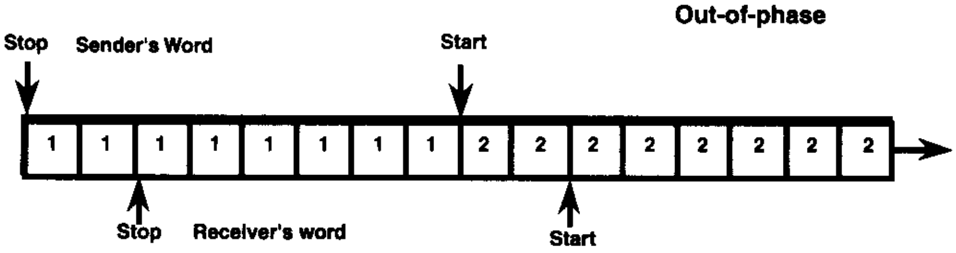

Figure 4–13 Byte sync.

cies are the same but their signals are out of alignment. Simply stated, the receiver is using a different set of 8 bits as a byte from the sender.

As we learned earlier, the byte BPV of the composite clock must be exactly the same with regard to amplitude polarity and the time it is measured. Figure 4–14 shows an out-of-phase signal from the perspective of transmitter and receiver.

Be aware that phase sync failure will not cause office alarms; therefore, no protection switching operation is invoked nor is an alternate data path chosen. Additionally, carriers normally do not monitor loss of a composite clock, but they will receive information via customer complaints.

In contrast to a system that does not employ BITS, timing problems are filtered out at a tandem office and Bellcore standards allow no more than four slips per day per circuit.

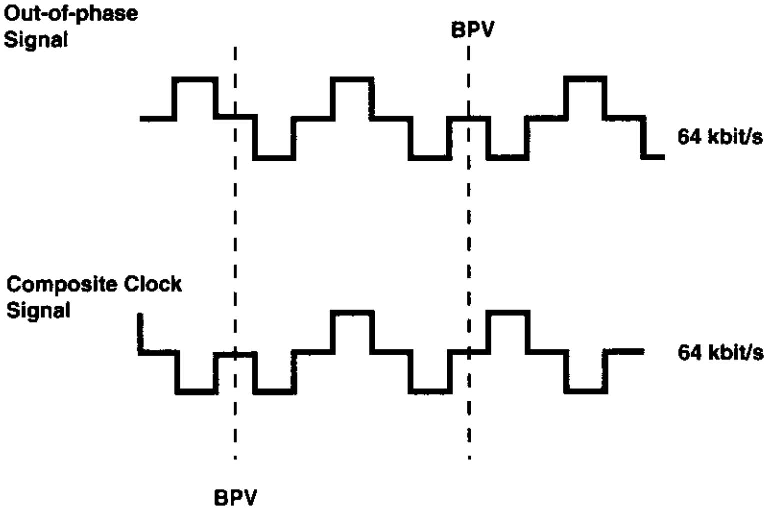

Figure 4–14 Composite clock and out-of-phase signal.

SYNCHRONIZATION STATUS MESSAGES (SSMS) AND TIMING LOOPS

One of the overhead bytes of the SONET header (overhead bytes are discussed in Chapter 5) contains the synchronization status message (SSM). SSM indicates the status and quality level of the SONET signal. It allows the network provider to define which clocking source (and its accuracy) is being used in the network. The SSM operation is often referred to as S1-byte synchronization messaging.

The SSM operation is quite helpful in a situation where a node (say, node A) loses its synchronization clock from a primary source [NORTa99]. In Figure 4–15(a), node A is providing clocking information to node B from BITS 1. Further, node B is also receiving clocking from yet another node, say, node C. However, the message sent from B to A is coded as "Don't use." This prevents A and B from going into a timing loop.

In Figure 4–15(b), node A loses its BITS 1 timing, or the timing is degraded (say to stratum 3). Node A informs node B of this situation with a message stating that is running a statum 3 (holdover mode). Node B realizes that its clocking from node C is better than that from node A. Therefore, it sends a message to node A informing A that C has

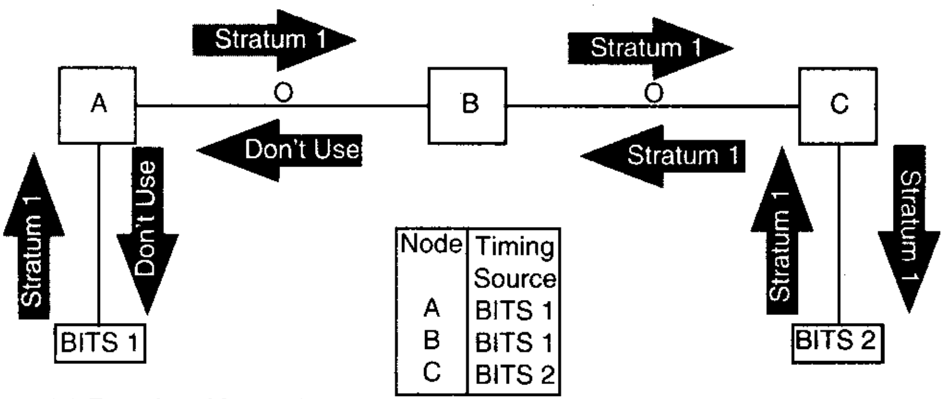

Node	Timing Source
A	BITS 1
B	BITS 1
C	BITS 2

(a) Fault-free Messaging

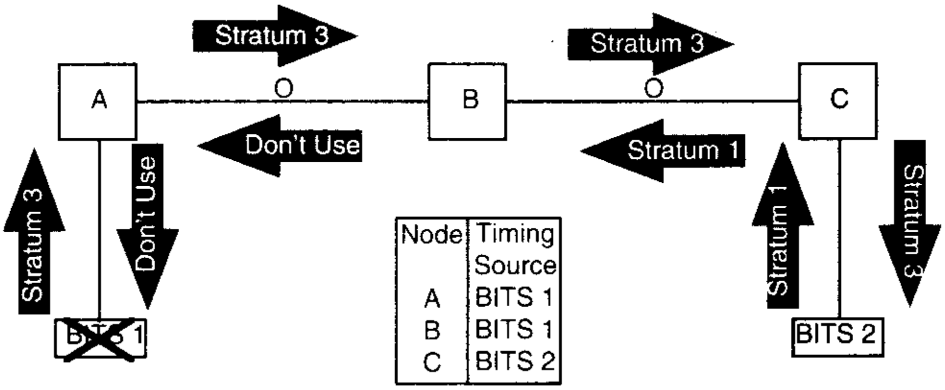

Node	Timing Source
A	BITS 1
B	BITS 1
C	BITS 2

(b) Timing Problem

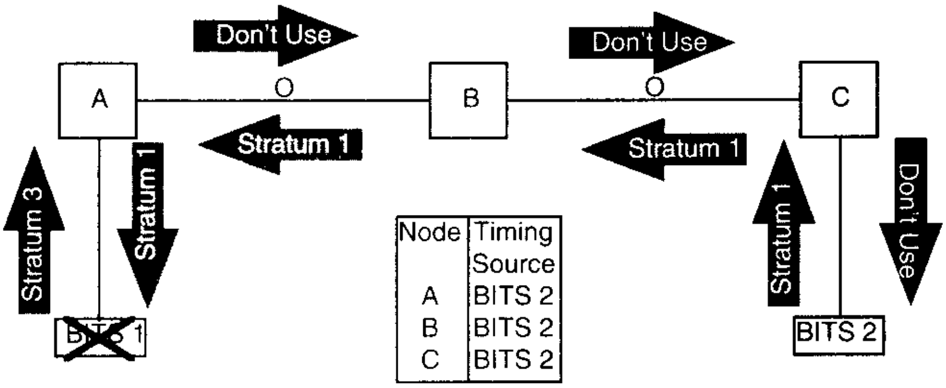

Node	Timing Source
A	BITS 2
B	BITS 2
C	BITS 2

(c) Reconfiguration

Figure 4–15 Synchronization messaging examples.

access to statum 1 clocking. This message is relayed from node A to the BITS source.

Now, all components are clocked from BITS 2, as depicted in Figure 4–15 (c). All nodes are receiving their timing from BITS 2, through node C.

SUMMARY

Synchronization and clocking functions are vital operations in digital networks. Older carrier systems were asynchronous in nature, but today, all implementations in backbone networks use highly accurate and precise synchronous clocks. With these newer systems, clock variations and slips are tightly controlled (and are rare in occurrence relative to—especially—1G systems). Third-generation optical transport networks operating in the terabit range require very precise and accurate synchronization services.

5

SONET and SDH

This chapter surveys the second generation digital transport network technology, as implemented in SONET and SDH. The topic is introduced by explaining the origins of SONET and SDH, and how the two standards came about. The chapter then concentrates on the principle features of SONET and SDH, citing their main functions and how they perform these functions.

Following are some points to keep in mind during the reading of this chapter. SONET and SDH use the terms frame and envelope to describe a discrete unit of traffic on a communications link. The terms are used interchangeably. The term synchronous payload envelope (SPE) denotes that part of the envelope that contains the user traffic; the remainder of the envelope contains overhead bytes used to help manage the SPE.

HOW SONET AND SDH CAME INTO BEING

SONET and SDH were developed to replace the aging T1 and E1 technologies that were first deployed in the early 1960s to provide a high-speed (1.544 Mbit/s and 2.048 Mbit/s, respectively) digital carrier system for voice traffic.

Extensive research had been underway for more than a decade on many of the features that are found in SONET/SDH. One notable achievement began in 1984. It focused on the efforts of several standards groups and vendors to develop optical transmission standards for what is known as the *mid-span meet* (also known as *transverse compatibility*). The goal was to publish a specification that would allow different vendors' equipment to operate with each other at the fiber level.

In addition, due to the breakup of the Bell System in 1984, there were no standards developed beyond the T3 technology. Prior to the divestiture, all equipment was built by AT&T's manufacturing arm, Western Electric (WECO), which ensured that there would be no compatibility problems of any components in the telephone network.

After the breakup, there was little incentive for the other carriers (such as MCI and Sprint) to purchase AT&T-based equipment. Indeed, there was an incentive *not* to purchase AT&T equipment, since AT&T, MCI, and Sprint had become competitors with each other for long-distance services. This situation led to the rapid growth of alternate equipment vendors (such as Nortel Networks), who were developing advanced digital switching technologies.

The 1984 divestiture paved the way for alternate long-distance carriers through the equal access ruling. The alternate carriers were given equal access to the local exchange carrier (LEC) infrastructure and connections to AT&T for end-to-end long-distance service. The LEC could connect to MCI, Sprint, and others through their switching facilities, at an interface in the LEC or long-distance carrier offices called the point of presence (POP).

During this time, higher capacity schemes beyond T3 became proprietary, creating serious compatibility problems for network operators who purchased equipment from different manufacturers. In addition, the early 1980s witnessed the proliferation of incompatible and competing optical fiber specifications.

The various standards groups began the work on SONET after MCI sent a request to them to establish standards for the mid-span meet. The SONET specifications were developed in the early 1980s and Bellcore submitted its proposals to the American National Standards Institute (ANSI) T1X1 Committee in early 1985, based on a 50.688 Mbit/s transfer rate. The initial SONET work did not arouse much interest until the Metrobus activity became recognized.

Later, using the innovative features of Metrobus, the SONET designers made modifications to the original SONET proposal, principally in the size of the frame and the manner in which T1 signals were mapped into the SONET frame.

From 1984 to 1986, various alternatives were considered by the ANSI T1 Committee, who settled on what became known as the synchronous transport signal number one (STS-1) rate as a base standard. Finally, in 1987, the ANSI T1X1 committee published a draft document on SONET.

PARTICIPATION BY ITU-T

During this time, the international standards body now known as the International Telecommunication Union-Telecommunication Standardization Sector (ITU-T) had rejected the STS-1 rate as a base rate in favor of a base rate of 155.520 Mbit/s. For a while, it appeared that the North American and European approaches might not converge, but the SONET frame syntax and structure were altered one more time to a rate of 51.84 Mbit/s, which permitted this rate to be multiplexed (concatenated) by an integer of three to the European preference of 155.52 Mbit/s. This work has resulted in almost complete compatibility between the North American and European approaches. The ITU-T Recommendations are now considered the "official" standards and are collectively called the Synchronous Digital Hierarchy (SDH).

Once the major aspects of the standards were in place, vendors and manufacturers began to develop SONET and SDH equipment and software. These efforts came to fruition in the early 1990s and, as of this writing, SONET and SDH have been deployed throughout the world.

REASONS FOR SUCCESS OF SONET/SDH

The SONET and SDH provide a number of attractive features when compared with the first generation transport networks. SONET/SDH is an optical-based carrier (transport) network utilizing synchronous operations between the network components/nodes, such as multiplexers, terminals, and switches. SONET/SDH's high speeds (some systems operate at the gigabit rate) rely on high-capacity fiber. Much of the T1 and E1 technology was geared toward the use of copper (twisted pairs) media, which operate at more modest transmission rates.

As just stated, the SONET/SDH network nodes are synchronized with each other through very accurate clocking operations, which insures that traffic is not "damaged," or lost due to clocking inaccuracies. T1/E1 clocking systems are very accurate, and they now use the same clocking mechanisms as SONET/SDH, but they were not originally so designed.

SONET/SDH is quite robust and provides high availability with self-healing topologies. In the event that a link is lost due to node or fiber failure, SONET/SDH can recover by diverting the traffic to back up facilities. Most T1/E1 systems can be configured for backup, but "robustness" is not an inherent part of a T1/E1 architecture.

SONET and its ITU-T counterpart SDH are international standards. As such, they pave the way for heterogeneous, multivendor systems to operate without conversions between them (with some exceptions).[1]

Unlike T1/E1, SONET/SDH give the network node a direct access to low-rate multiplexed signals, without the need to demultiplex the signals back to the original form. In other words, the payloads residing inside a SONET/SDH signal are directly available to a SONET/SDH node.

SONET/SDH provide extensive Operations, Administration, and Maintenance (OAM) services to the network user and administrator. Indeed, about 4% of the bandwidth in a SONET/SDH network is reserved for OAM. A T1 system allows only one bit per 193 bits for OAM&P, and E1 provides only one byte per frame. With this comparison in mind, it is easy to conclude that SONET and SDH have the capability for more extensive and powerful network management operations than T1 or E1.

THE SONET MULTIPLEXING HIERARCHY

We have emphasized several times that one of the most important functions of a transport network is the multiplexing of user traffic on low-capacity links into much larger payloads on higher capacity links. We also explained that a multiplexing hierarchy is essential to keep the many payloads (a) organized and (b) identified on the media. This section explains how SONET accomplishes this critical operation.

The synchronous transport signal-level 1 (STS-1) forms the basis for the optical carrier-level 1 signal. OC-1 is the foundation for the complete synchronous optical signal multiplexing hierarchy. The higher level signals are derived by the multiplexing of the lower level signals. The high-level signals are designated as STS-N (or electrical signals) and OC-N (for optical signals), where N is an integer number.

As illustrated in Table 5–1, OC transmission systems are multiplexed by the N values of 1, 3, 9, 12, 18, 24, 36, 48, to 192. Presently, signal levels OC-3, OC-12, OC-48, and OC-192 are most widely supported multiples of OC-1. Table 5–1 also shows the number of DS1 and DS3 signals that are carried in the OC envelopes.

---

[1] First generation digital carrier systems (such as T1 and E1 in Europe) are not standardized on a worldwide basis, and different systems exist in various parts of the world.

Table 5–1 SONET Transmission and Relationship to Asynchronous Payloads

Electrical	Optical Hierarchy	Transmission Line Rate (Mbit/s)	DS-3 Equiv.	DS1 Equiv.	DS0 Equiv.
STS-1	OC-1	51.840	1	28	672
STS-3	OC-3	155.520	3	84	2,016
	OC-9	466.560	9	252	6,048
	OC-12	622.080	12	336	8,064
	OC-18	933.120	18	504	12,096
	OC-24	1,244.160	24	672	16,128
	OC-36	1,866.240	36	1,008	24,192
	OC-48	2,488.320	48	1,344	32,256
	OC-96	4,976.640	96	2,688	64,512
	OC-192	9,953.280	192	5,376	129,024

SONET AND SDH MULTIPLEXING STRUCTURE

Table 5–2 illustrates the similarities and differences of the SONET and SDH multiplexing structure. We use this table to compare the basic terminology used in the SONET and SDH multiplexing operations. The multiplexing operations of SONET and SDH are quite similar. Part of the challenge to understanding their relationships is to understand that the two hierarchies use different terms to describe very similar functions.

Table 5–2 Comparison of SONET/SDH Rates and Services

SONET	SDH	Bit Rate-Mbit/s	Mux. Rate
VT1.5		1.728	4 x 1.728 = 6.912
	VT2	2.304	3 x 2.304 = 6.912
VT3		3.456	2 x 3.456 = 6.912
VT6		6.912	1 x 6.912 = 6.912
STS-1	STM-0	51.84	7 x 6.912 or 1 x 51.84 = 51.84
STS-3	STM-1	155.52	3 x 51.84 = 155.52
STS-12	STM-4	622.08	4 x 155.52 = 622.08
STS-48	STM-16	2488.32	4 x 622.08 = 2488.32
STS-192	STM-64	9953.28	4 x 2488.32 = 9953.28

SONET uses the term virtual tributary (VT) to describe a specific user payload that it transports across a communications link. SDH uses the term virtual container (VC) to describe the user payload. For example, a VT might consist of a DS1 payload, and a VC might consist of an E1 payload. In some SDH documents, the term VT is also used.

Other key terms for the SDH signals are as follows: container (C-n), virtual-container (VC-n), tributary unit (TU-n), TU group (TUG-n), administrative unit (AU-n), and AU Group (AUG). The "n" following the various designations represents an integer number. The designations can have different numeric values following them. So "n" is substituted with a numerical value depending on what part of the SDH mapping hierarchy is being discussed.

SDH payloads are called synchronous transport modules (STMs), and are: STM-1, STM-4, STM-16, and STM-64. This notion also applies to both the container (C) designations and virtual container (VC) designations, as they can be represented with double digits. However, with the STM signals, the decimal number following the letters STM represents its multiplexing level.

In SONET, STS-3 is comprised of 3 DS-3s plus the SONET overheads, STS-12 has the bandwidth capabilities of 12 DS-3s, STS-48 has 48 DS-3s worth of bandwidth, and finally, there are 192 DS-3s equivalency of bandwidth in an STS-192 signal. Note that each jump in level is equal to four times, just as it is in the SDH STM hierarchy.

THE SONET/SDH FRAME STRUCTURE

An effective way to keep the many payloads organized on the fiber is to make certain that each bit and each byte is concisely and precisely structured and identified. This task is accomplished by defining how the user payloads and supporting overhead bytes are placed onto the fiber at the sending node and how they are removed at a receiving node. In this section, we learn how the organization is accomplished.

The basic transmission unit for SONET and SDH is the envelope (frame), as illustrated in Figure 5–1. It is comprised of 8 bit bytes (octets) that are transmitted serially on the optical fiber. For ease of documentation, the payload is depicted as a two-dimensional map. The map is comprised of n rows and m columns. Each entry in this map represents the individual octets of a synchronous payload *envelope*. (The "F" stands for flag, and is explained later.)

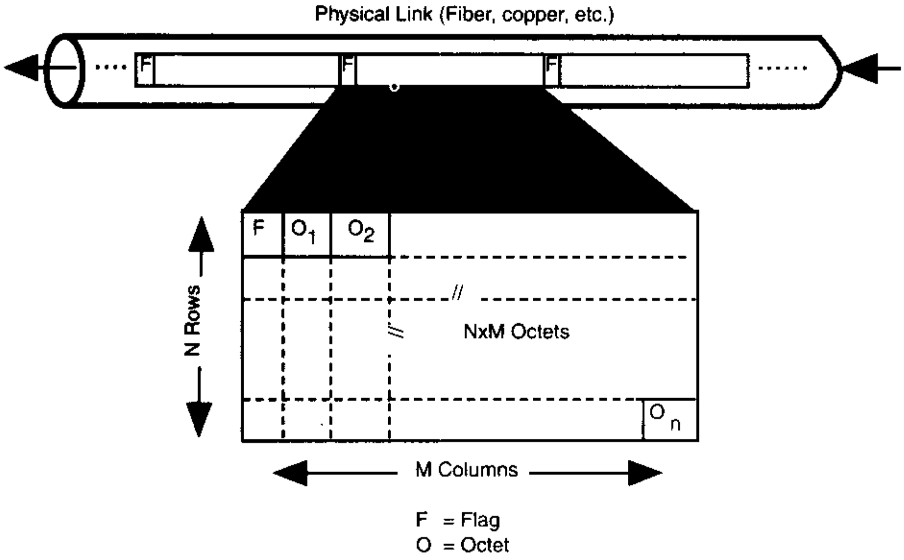

Figure 5–1 The SONET/SDH envelope (frame).

The octets are transmitted in sequential order, beginning in the top left-hand corner through the first row, and then through the second row, until the last octet is transmitted, which is the last row and last column.

The envelopes are sent contiguously and without interruption, and the payload is inserted into the envelope under stringent timing rules. In addition, a user payload may be inserted into more than one envelope, which means that the payload does not need to be inserted at the exact beginning of the part of the envelope that is reserved for this traffic. It can be placed in any part of this area, and a pointer is created to indicate where it begins.

This approach allows the network to operate synchronously while accepting asynchronous traffic. That is, traffic arriving at a node (say, a cross-connect) does not need to be synchronized with the clock at this node. By simply placing the user payload into the payload envelope as the user traffic arrives, the pointer can be set to indicate where the payload is. Thereafter, the payload is encapsulated into the SONET/SDH and sent on its way, using the ongoing network clocks to keep things synchronized.

Rationale for the 51.840 Mbit/s Envelope

In the second generation transport technology, all multiplexing rates start at the SONET rate of 51.840 Mbit/s or the SDH rate of 155.520 Mbit/s. As noted, the SONET basic rate was the foundation for all higher level rates in SONET and SDH. Therefore, it will be helpful to understand how this rate is created and the reasons for using the rate in the first place.

A key consideration in the design of SONET and SDH was to be able to support the T1 and E1 125 μsecond clocking increment. This timing unit is fundamental to the North American, Japanese, and European digital voice transport system because the universal 64 kbit/s voice channel is derived as:

- A voice signal is sampled at 8,000 times per second.
- Therefore, each sample duration is 125 μsec (1 sec / 8000 = .000125).
- 8 bits per samples × 8,000 samples per second = 64,000 bit/s.

The basic transmission unit for SONET is the STS-1 frame. See Figure 5–2. The frame consists of 90 columns and 9 rows of 8-bit bytes

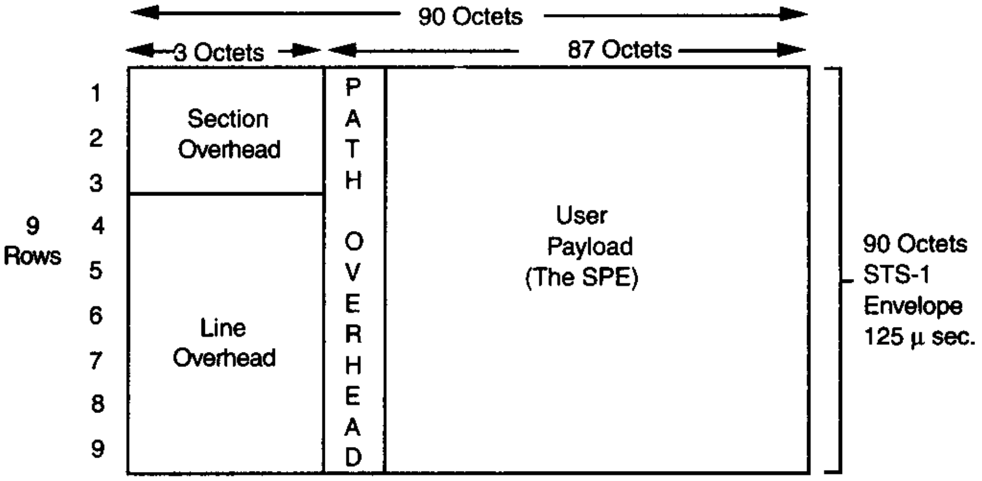

90 Octets
x9 Rows
x8 Bits per Octet
6,480 Bits
x8000 125 μ sec Slots per Second
51,840,000 or 51.840 Mbit/s

Figure 5–2 The basic SONET frame structure.

(octets). Therefore, the frame carries 810 bytes or 6480 bits. SONET transmits at 8000 frames/second. Therefore, the frame length is 125 μsec. This approach translates into a transfer rate of 51.840 Mbit/s (6,480 x 8000 = 51,840,000). Obviously, this approach interworks well with the 64 kbit/s digital voice signal. And don't forget, SONET and SDH were designed to support digital voice traffic.

Overhead and User Areas in the Envelope

Notice that the first three columns of the frame contain transport overhead, which is divided into 27 bytes with 9 bytes allocated for section overhead and 18 bytes allocated for line overhead. The other 87 columns comprise the STS-1 envelope capacity (although the first column of the envelope capacity is reserved for STS path overhead). Several fields are used for signaling control, administrative alarms, etc. Several other fields are used to identify the type of equipment being used as well as the types of payloads that reside in the envelope. Other bytes are used for pointers to locate the position of the user payload in the frame. This overhead is explained shortly.

The 87 columns are also called the synchronous payload envelope (SPE). The actual user payload consists of 86 columns or 774 bytes. Therefore, the user payload operates at 49.536 Mbit/s (774 x 8000 = 49,536,000). Obviously, the user payload can support VTs up to the DS3 rate (44.736 Mbit/s).

The STS-1 frame is transmitted row by row from left to right. Each byte is transmitted with the most significant bit first.

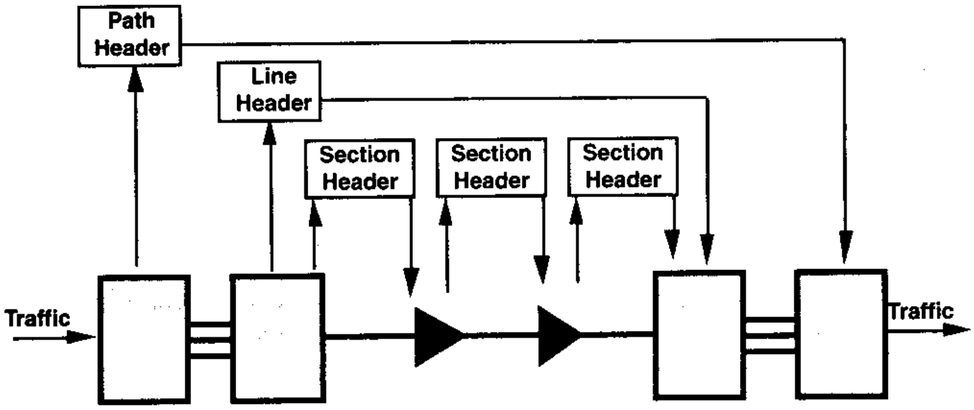

Figure 5-3 SONET functional components.

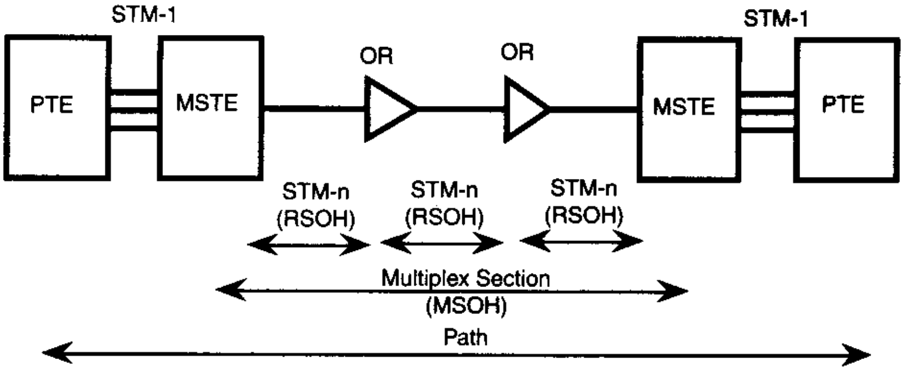

Figure 5–4 SONET functional components.

SONET AND SDH FUNCTIONAL COMPONENTS

Figure 5–3 shows the major functional components on a SONET link: the section, line, and path operations. The concept is to partition activities and responsibilities among the three major components. This modular approach facilitates the ability to add and remove components, because from a protocol standpoint (logical operations), the three components are somewhat independent of each other, and their operations are described shortly.

Figure 5–4 shows the major SDH components. Once again, the similarities between SONET and SDH are evident, but SDH further distinguishes differences in the path overheads. The first one is the higher order path overhead (HO-POH) and the second is the lower order path overhead (LO-POH). There is some difference in how these overhead octets are used. In the SONET arena, the overheads are referred to as section, line, and path overheads. There is no distinction among types of Path overheads.

SONET AND SDH PROBLEM DETECTION

SONET and SDH define many rules and actions for the detection of problems and failures on a link or in a node. Additionally, they provide information on how the detection of these problems is signaled to various parts of the span. Figure 5–5 shows how the detection of certain events can trigger alarm signals. There are six failure conditions:

- Loss of signal (LOS)
- Loss of frame (LOF)

- Loss of pointer (LOP)
- Alarm indication signal (AIS)
- Far end receive failure (FERF)
- Remote alarm indication (RAI)

The detection of LOS, LOF, or LOP at the section terminating equipment (STE) or line terminating equipment (LTE) causes the generation of alarms on that device's output port to the downstream network element. For example, if an STE detects a LOS or LOF, its output port generates an AIS of all 1s to the downstream LTE, which in turn generates an AIS with H1, H2 to the path terminating equipment (PTE); this in turn generates an AIS with V1, V2 to the VT PTE. The VT PTE might then generate a DSn AIS on the tributary. These events also invoke upstream signals. The LTE sends a foreign receive failure using a K2 to its

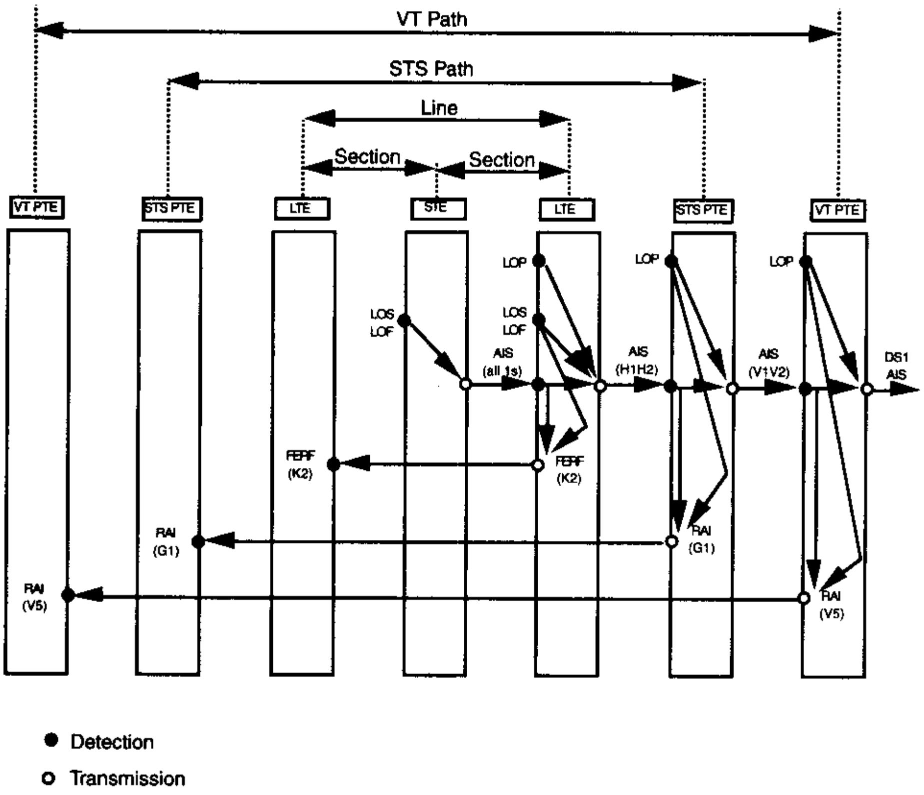

Figure 5–5 Examples of alarm signals.

associated LTE upstream. Likewise, the VT PTE sends an RAI from V5 to its peer VT PTE.

LOCATING AND ADJUSTING PAYLOAD WITH POINTERS

SONET and SDH use a pointer concept to deal timing (frequency and phase) variations in a network. The purpose of pointers is to allow the payload to "float" within the payload area of the envelope. Figure 5–6 shows this idea. The pointer is an offset value that shows the relative position of the first byte of the payload.

During the transmission across the network, if any variations occur in the timing, the pointer needs only to be increased or decreased to compensate for the situation.

Several options are available for how the payload is mapped into the frame. The option just discussed is called the floating mode, for obvious reasons.

Another option is called the locked mode. With this approach, the pointers are not used and the payload is fixed within the frame. It cannot float. This approach is much simpler, but it requires that timing be maintained throughout the network. Since all signals have a common orientation, the processing of the traffic is efficiently performed.

VIRTUAL TRIBUTARIES IN MORE DETAIL

Virtual tributaries (VTs) are used to support sub-STS-1 levels, which are simply synchronous signals used to support low-speed signals. To support different mixes of VTs, the STS-1 SPE can be divided into seven

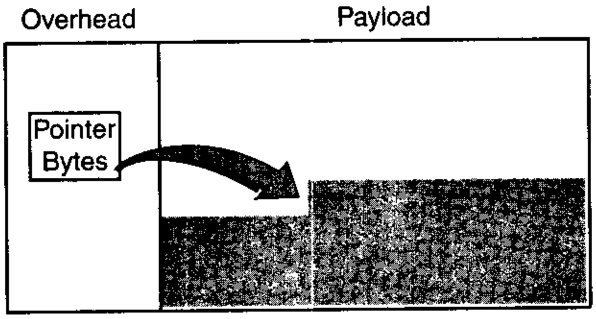

Figure 5–6 Payload pointers.

groups. Each group occupies 12 columns and 9 rows of the SPE and may actually contain 4, 3, 2, or 1 VTs. For example, a VT group may contain one VT 6, two VT 3s, three VT 2s, or four VT 1.5s. Each VT group must contain one size of VTs, but different VT groups can be mixed in one STS-1 SPE.

The four sizes of the VT are as follows: VT 1.5 = 1.728 Mbit/s, VT 2 = 2.304 Mbit/s, VT 3 = 3.456 Mbit/s, VT 6 = 6.912 Mbit/s.

Figure 5–7 shows a VT 1.5 group. The 1.5 bytes occupy three columns and nine rows. The actual user traffic consists of 24 bytes in

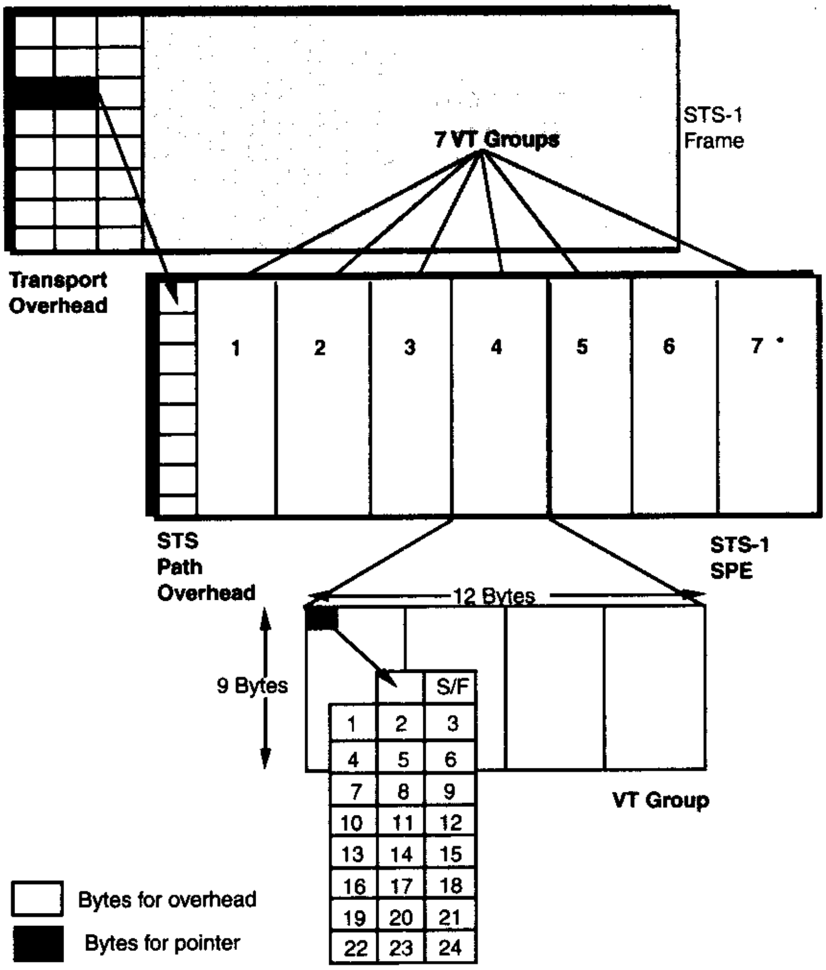

Figure 5–7 Payload management with VTs.

accordance with a T1 24 slot frame. The remaining three bytes are used for SONET control. A VT 1.5 group supports four VT 1.5 transmissions to occupy the full 12 columns of the VT structure.

At the bottom of Figure 5–7 reside the 24 user samples of a, say, voice signal. Each sample is 8 bits, and, since the SPE is sent 8,000 frames a second, one sample of the voice image is carried in each SPE. Once again, supporting the 64,000 bit/s voice signal is fundamental to the operation of a digital transport network.

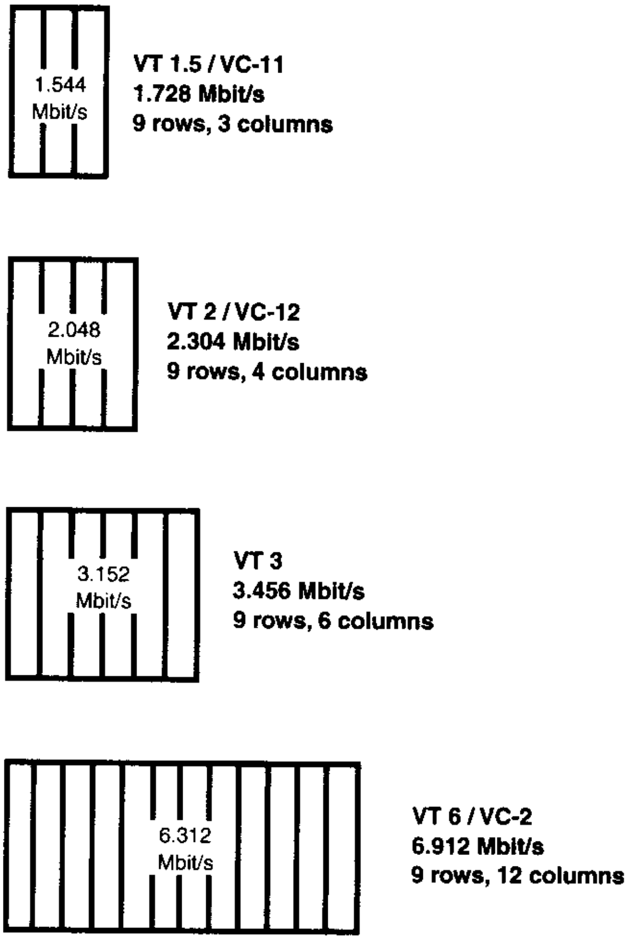

Figure 5–8 VTs and VCs.

VIRTUAL TRIBUTARIES AND VIRTUAL CONTAINERS

The various administrations and standards groups from Japan, North America, and Europe worked closely together to accommodate the three different regional signaling standards. The initial SONET standards published in the United States in 1984 were reviewed by Japan and the European PTTs to see if these requirements would meet their needs. During this time, the ANSI T1 committee had become involved with Bellcore in the development of SONET.

The original SONET standard made no provision for the European rate of 140 Mbit/s and, for a time, the various administrations and standards groups tried to accommodate all multiplexing rates of all three regions. Discussions continued through the European Telecommunications Standard Institute (ETSI) and agreement was reached on a subset of the multiplexing schemes of the three regions. ETSI stressed the importance of the intermediate rates of 8 and 34 Mbit/s, in contrast to the ease of doing international networking. Reason prevailed and compromises were reached. The importance of international internetworking came to the fore and multiplexing schemes based on 1.5, 2.48, 6.312 were accommodated. Figure 5–8 shows the structure of the virtual tributaries (VT) and virtual containers (VC). VT1.5 is called VC1–11 in Europe and it accommodates the T1 rate. VT2 is called VC-12 in Europe and it accommodates the Europe CEPT 1 rate of 2.048 Mbit/s. VT3 is not employed in Europe. It is used in North America to optimize multiplexing DS1c transport signals. VT6 is called VC-2 in Europe and it accommodates the 6.312 Mbit/s rate from all three regions.

THE OVERHEAD BYTES

Figure 5–9 shows the structure and names of the SONET OAM headers and the bytes within the headers. The reader may recall that each major component is responsible for creating a header at the transmitting network element and processing the header at the receiving network element. Let's start with a general look at the section overhead bytes:

- *A1 and A2 (framing):* Flags used by a receiving machine to synchronize onto the SONET signal.
- *J0 (trace):* Used for STS-1 identification; a unique number that is assigned to each STS-1 of an STS-n signal.

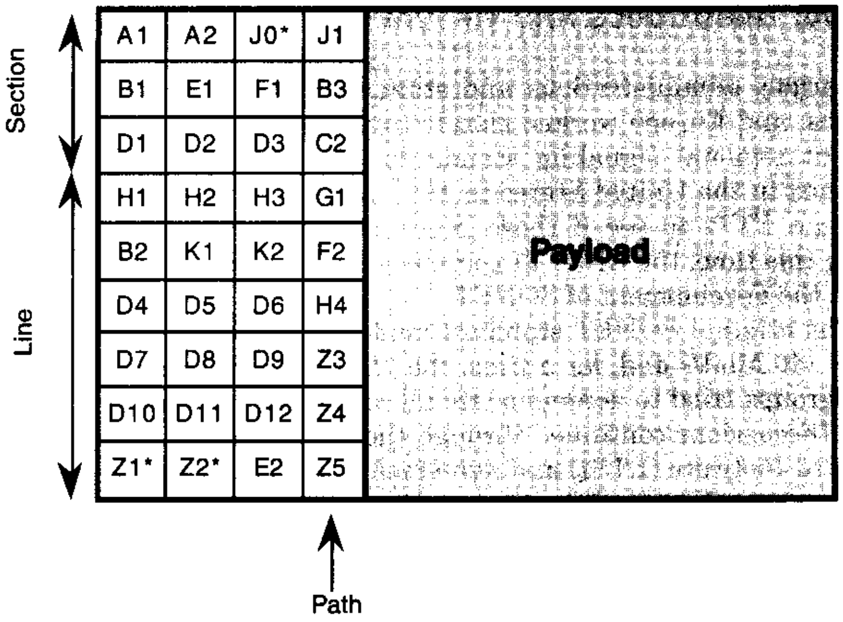

A1	A2	J0*	J1
B1	E1	F1	B3
D1	D2	D3	C2
H1	H2	H3	G1
B2	K1	K2	F2
D4	D5	D6	H4
D7	D8	D9	Z3
D10	D11	D12	Z4
Z1*	Z2*	E2	Z5

Notes: for J0*, Z1*, and Z*, see text for additional byte designations

Figure 5–9 Abbreviated names for the OAM bytes.

- *B1 (bit interleaved parity):* A parity check on the previously sent STS-1 frame with the answer (the parity) in the current frame. The BIP-8 byte checks for transmission errors over a section. If excessive errors are occurring, alarms and diagnostics (as shown in Figure 5–5) may be generated.
- *E1 byte (orderwire):* A 64 kbit/s voice channel which can be used for maintenance communications among terminals, hubs, and regenerators.
- *F1 (user):* Set aside for the network provider to use in any manner deemed appropriate.
- *D1, D2, and D3 (data communications channel (DCC)):* Used for data communications channels and are part of 192 kbit/s operations that are used for signaling control, administrative alarms, and other OAM&P. In later chapters, discussions will explain how the DCC can be used for MPLS, and other more advanced operations.

The line overhead bytes occupy the bottom 6 octets of the first three columns in the SONET frame. Line overhead is processed by all equipment except for the regenerators. Here is a general view of their functions:

- *H1 and H2 (pointers):* Indicate the offset in bytes between the pointer and the first byte of the SPE. This pointer allows the SPE to be located anywhere within the envelope, as long as capacity is available.
- *H3 (pointer action):* Used to frequency justify the SPE; that is, to allow for possible slight timing differences that may exist between nodes.
- *B2:* Use is identical to the BIP byte found in the section header, except that this byte pertains only to the line header.
- *K1 and K2 automatic protection switching (APS):* Used for detecting problems with the line terminating equipment, and for alarms and signaling failures, as well as network recovery.
- *D4-D12:* (data communications channel): Used for line communication and are part of a 576 kbit/s message which is used for maintenance control, monitoring, alarms, etc. Originally, the Common Management Information Protocol (CMIP) was defined for use in the bytes, and some vendors still use CMIP. Others use the Transaction Language 1 (TL 1) in these bytes, or in the F1 byte. Still others use the Simple Network Management Protocol (SNMP).
- *Z1 and Z2 bytes:* Originally these bytes were reserved for future growth, and are now partially defined. The Z1 byte is also designated as the S1 byte. It is used to convey synchronization information about the SPE, and it allows the node to make decisions about potential clocking sources. The Z2 byte is also designated as byte M0 or byte M1. Its use is to convey information about error conditions back to the source of the SPE.
- *E2 (orderwire):* Explained earlier.

The path overhead remains with the payload until the payload is demultiplexed at the far-end node (the STS-1 terminating equipment). This processing is usually at the customer premises equipment (CPE). The functions of the path overhead bytes are as follows:

- *J1(path trace):* Used to repetitively transmit a 64 byte fixed length string in order for the recipient path terminating equipment to verify a connection to the sending device.
- *B3:* Function is the same as that of the line and section BIP fields, except that it performs a parity check calculated on all bits in the path overhead.
- *C2 (path signal label):* Used to indicate the construction of the STS payload envelope. The path signal label can be used to inform the

network that different types of systems are being used, such as ATM, FDDI, etc.

- *G1 (path status):* Carries maintenance and diagnostic signals such as an indication for block errors, and several identification functions.
- *F2:* Used by the network provider.
- For certain support functions, the growth bytes of SONET Z3 are to be used for DQDB mapping. Z4 for SONET and SDH is still a growth byte. Z5 for both SONET and SDH is for error monitoring.

SONET AND SDH CONCATENATION

It is possible to concatenate (for example) three STS-1s into one STS-3c frame. This operation is common when payload such as ATM or IP is transported, and a channelized structure is not needed. Beyond STS-3c, SONET and SDH support a variety of concatenation operations, and all are identified with the letter "c" appended to the back of the multiplexing designation. Chapter 6 shows provides more information on concatenation and [BLAC01] has the details if you need them.

SUMMARY

SONET and SDH represent the second generation transport network, and in many networks, especially long-haul networks, the T1/E1 first generation systems no longer exist. However, in order to provide a graceful evolution to SONET and SDH, the 2G networks are backward compatible and able to support 1G operations. Thus, the well-known 125 μsec timing increment is carried over to SONET and SDH. Likewise, the TDM concepts that are fundamental to T1 and E1 are retained in SONET and SDH as well.

The 2nd generation systems are quite different from their predecessors in that they provide (a) much more capacity to users, (b) have powerful OAM capabilities with the three headers, and (c) operate with extensive backup and protection arrangements. We explained points (a) and (b) in this chapter. Point (c) is covered in Chapter 8.