1

INTRODUCTION TO NETWORKING

1.1 INTRODUCTION

We all use the Internet. However, few of us are aware of exactly what is happening once we click a mouse after typing a key word in a search engine. How is our question traveling to Google computers? What type of hardware

Network Infrastructure and Architecture: Designing High-Availability Networks,
By Krzysztof Iniewski, Carl McCrosky, and Daniel Minoli

is involved in this process? How are bits and bytes moved from point A to point B? How do they know which way to go? Who takes care of the process of data delivery?

In this chapter we provide a gentle start on our journey of discovering how data networks, which are responsible for the delivery of information across the Internet, work in today's world. The text serves as an introduction to broadband networking and deals with basic concepts in communications. We start by discussing the various transmission media used and show that optical fiber is a much better medium than copper wire or coaxial cable. This is followed by a description of the classes of various networks present today. We briefly mention the basic characteristics of access, local area, storage, and metropolitan and wide area networks. In the following section we go over possible network topologies. We show differences between point to point, hub, ring, and mesh architectures. Our discussion of network topologies is followed by a review of an open system interconnect (OSI) model. Following the OSI model we discuss three methods of multiplexing data: wavelength division, time domain, and statistical multiplexing. At the end of the chapter we describe various classes of networking equipment. We examine briefly the functionality of networking devices such as regenerator, hub, switch, or router.

1.2 TRANSMISSION MEDIA

Four types of media can be used to transmit information in the world of communications: copper wire, coaxial cable, optical fiber, and air (wireless communication).

We discuss all four options briefly to illustrate that optical fiber is by far the best option available when information is to be transmitted over long distances.

1.2.1 Copper Wire

In the old days, copper wire was the only means of transmitting information. Technically known as unshielded twisted pair (UTP), the connection consists of a number of pairs, typically two or four, of copper wires. The wire pairs are called "twisted" because they are physically twisted. This helps reduce both crosstalk between wires and noise susceptibility.

UTP cable does not have a shield, and therefore a high-frequency part of the signal can "leak out." Also, the twisting on the copper pair can be quite casual, designed as much to identify which wires belong to a pair as to handle transmission problems. Although not perfect, this cabling technology was for many years quite satisfactory for voice communication purposes. Consequently, there are millions of kilometers of copper wires in the public switched telephone network. However, not only did the copper wire itself have limitations, but things were done to the wiring to make it even more unsuitable for high-speed data transmission. These actions took many forms, although we only mention one here—load coils—as an example.

Load coils were frequently added to wiring loops longer than a few kilometers. The coils were essentially low-pass filters, which meant that they passed without a loss of the low frequencies that correspond to your voice, but blocked higher frequencies. High-frequency blocking is disastrous for data communications, as we actually rely on high frequencies to achieve the desired speed of data transmission.

In conclusion, what was once a good system for voice transmission has become a big problem for data transmission. In fact, out of all possible wires, UTP copper wire is probably the least desired, as it has the worst transmission properties. However, it is present virtually everywhere, as most households, at least in developed countries, have copper wire phone connections already in place.

For networking applications, the term UTP generally refers to 100-Ω category 5 cables, the characteristics of which are tightly specified in data networking standards. Techniques have been developed to provide means of having a phone conversation and high-speed data transfer using the same twisted-pair copper wire. These techniques, called DSL (digital subscriber loop) technology, come in many different flavors, asymmetrical DSL being the most popular for home Internet connectivity. However, even with state-of-the-art hardware, the highest bandwidth of the DSL connection might get to a level of 50 million bits per second (Mb/s). In practice, a more typical number would be around 1 Mb/s. As we will see in a moment, this is a relatively low number compared to that of optical fiber. The difference results from the poor transmission characteristics of unshielded copper wires compared with the excellent transmission characteristics of optical fibers.

The low bandwidth of UTP cabling is only one of its problems. The second problem is its signal attenuation, which is very high in this medium. For example, a DSL signal virtually disappears after transmission over a few kilometers. It needs to be recovered, amplified, and retransmitted by sophisticated electronic equipment. You can imagine that it would have taken many UTP wires and retransmission points to send a stream of 10 Gb/s from New York to San Francisco. So although UTP cable is useful at home, clearly we need something better to send significant amounts of data over long distances. Let us look next at coaxial cable.

1.2.2 Coaxial Cable

Coaxial cable consists of a single strand of copper running down the axis of the cable. This strand is separated from the outer shielding by an insulator made of a dielectric material. A conductive shield covers the cable. Usually, an outer insulating cover is applied to the overall cable. Because of the coaxial construction of the cable and the outer shielding, it is possible to send quite high frequencies. For example, in cable television systems in North America, 20 TV channels, each with 6 MHz of bandwidth, can be carried on a single coaxial cable.

Coaxial cable access was used originally for the purpose of broadcast video. As a result, the system is inherently well suited for high-speed data transfer. Techniques have been developed to provide broadcast video and high-speed Internet access on the same coaxial cable. We refer to these techniques collectively as *high-speed cable modem technology*. However, this term involves a bit of cheating. The term *coaxial cable* notes, in practice, a system that is really a hybrid of optical fibers and a coaxial cable with the optical fiber portion of the cable network being hidden from the end user. The reason for this "hybridization" of the coaxial cable network is that although coaxial cables have better signal transmission properties than UTP cables, they are inferior to optical fibers. As a result, your cable company, which provides you with TV programming, only needs to install optical fibers in its infrastructure in the last few kilometers adjacent to you when using coaxial cables.

Another issue with coaxial cable is the fact that cable access is not available universally. From this brief description, the conclusion seems to be that although coaxial cable is somewhat better than UTP wires, it is still not good enough for the purpose of transmitting large amounts of information over long distances. Let us look next at optical fiber.

1.2.3 Optical Fiber

Fiber is the third transmission medium we discuss. We will say immediately that it is unquestionably the transmission medium of choice. Whereas transmission over copper utilizes frequencies in the megahertz range in the best of cases, transmission over fiber utilizes frequencies approximately 1 million times higher, in the terahertz range. Terahertz frequencies translate into terabytes per second (TB/s) of bandwidth; this is very large indeed. For example, the most congested Internet connection today still requires lower bandwidth than 1 TB/s. To say this another way, a single strand of optical fiber is at present sufficient to carry all Internet traffic between North America and Europe. Not bad.

Fiber-optic technology, offering virtually unlimited bandwidth potential, is the transmission medium of choice for long-haul networks. It is also widely considered to be the ultimate solution to the delivery of broadband access to the end user, called the *last mile*, the network space between the carrier's central office and the user location. The last-mile piece of the network is typically the area where the majority of bottlenecks occur that slow the delivery of data services. Right now, UTP wire and coaxial cable are still used in most cases for last-mile data delivery.

Beyond its enormous transmission bandwidth, optical fiber has another great advantage: its very low attenuation loss. Optical signals travel for kilometers without losing virtually any of their strength. Recall for comparison that electrical signals traveling over copper wires basically disappear after a few kilometers! If you think of the optical fiber as a piece of glass, you would be able to see objects several kilometers away when looking into it. This is tremendously impressive and useful. There is no need to regenerate optical signals for several kilometers of its transmission along the optical fiber.

Besides the available bandwidth, which is enormous compared to copper wire or coaxial cable, and the very low signal loss, which is again much lower, optical fiber has other advantages. For example, optical signal propagating in the fiber is insensitive to external noise. Chapter 2 is devoted to issues of optical transmission, as they lay the foundation for optical networking, so we defer a full discussion until then. For now, let us just summarize by saying that the data signal in an optical fiber has a very high bandwidth, a very low attenuation coefficient, and is very robust (it is quite difficult to disturb).

It appears as though we are done with the choices for suitable transmission media: Optical fiber wins hands down. But wait a moment, what about wireless? Everything is wireless these days: your CDMA or GSM cellular phone, your WiFi 802.11 connection in your laptop, and even your garage door opener. Could it be that wireless is more suitable than optical fiber? Not really, but for the sake of completeness we discuss wireless next.

1.2.4 Wireless Communication

Wireless communications is our final option for a transmission medium. In this case the medium is simply air. Wireless transmission can take several forms: microwave, low-Earth-orbit satellites, cellular, or wireless local area networks. In every case, however, a wireless system seems to obviate the need for a complex wired infrastructure. There are also other advantages. In the case of satellites, transmission can take place across any two points on the globe. With microwave communications there is no need to install cabling between microwave towers. Cellular phones, for example, afford significant mobility. Similarly, your laptop WiFi wireless card lets you surf the Net in any Starbucks with hotspot availability.

These advantages come at a price. Wireless communications is significantly less efficient than optical communication systems in terms of bandwidth available. The most advanced, third-generation cellular phone systems provide bandwidth of hundreds of kilobytes per second at best, whereas standard optical transmission runs routinely at rates of several gigabyte per second. Using advanced optical technology called wave-division multiplexing, which we discuss in detail in Chapter 3, transmissions of terabits per second in one optical fiber are possible.

At this point, it is difficult not to stop and wonder over the fact that there is a difference of a few orders of magnitude in terms of available signal bandwidth between wireless and optical networks. Looking at this comparison another way, it would take about 1 million times longer to transfer a large file using a wireless connection than to send it using optical fiber. Clearly, you would not want to send a 1-gigabyte file using air as the medium!

The reasons for such lower bandwidths in wireless systems are fairly obvious. Air is not a very suitable medium for transmission, as signals propagate in all directions. There are many signal distortions as signals bounce from numerous objects and interfere with hundreds of other signals present simultaneously. Most important, unlike in an optical fiber case, where the signal is confined to a very small physical space, a wireless signal loses it strength very quickly as its energy propagates in three-dimensional sphere. Finally, mobility of the user places an additional burden on the wireless communication system.

Therefore, the conclusion seems to be that although wireless systems are very useful in many applications, they are not suitable for sending large amounts of data over long distances. Optical fiber is a clear winner for this application, as we anticipated from the start. For this reason, for the remainder of the book, unless specified otherwise, we assume that data transmission takes place over optical fibers.

1.3 BASIC NETWORKING CONCEPTS

The global broadband network is a connection of thousands of different networks implemented in various countries and cities by companies using a variety of technologies and protocols, as shown schematically in Figure 1.1. For a number of historical reasons, the global broadband network is more complicated than it could have been if it were build today from scratch, but that is a nature of many technologies. We hope to shed some light on this complex jungle of networking hardware technologies by explaining the basic concepts. Armed with this knowledge, you should be able to navigate Internet infrastructure with better understanding, and you should be able to find additional details, if needed, in the appropriate reference sources.

1.3.1 LAN, SAN, MAN, and WAN

To start our discovery process in network complexities, we will divide the global network into five classes of networks that can be identified based on their geographical span and functionality:

- *Access networks*: networks that connect consumers and corporate users with Internet infrastructure. DSL and cable modems are examples of access technologies.
- *Local area networks* (LANs): networks that connect multiple users in a contained environment such as a corporate building or a campus. Typically, Ethernet is used as the protocol of choice.

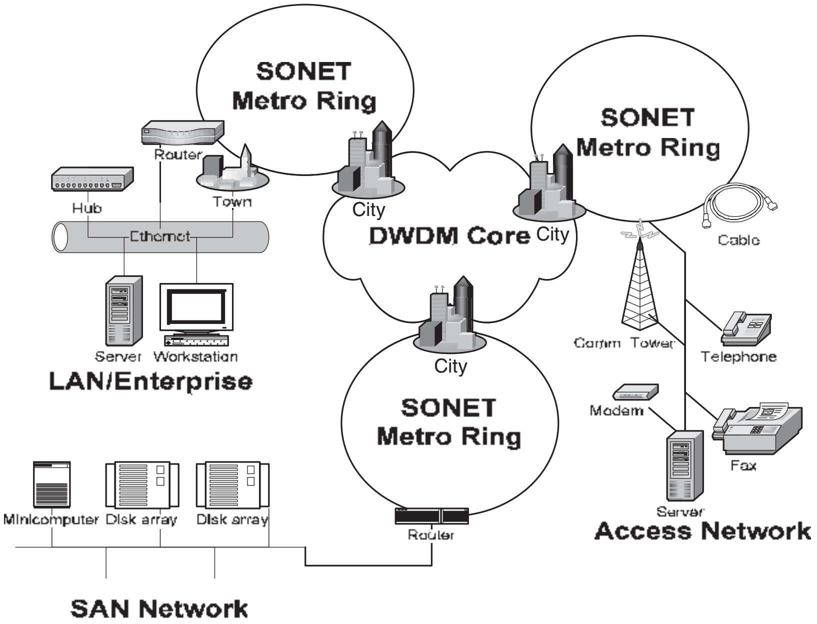

FIGURE 1.1 Broadband network elements consisting of a network core, synchronous optical network (SONET) rings, a local area network, storage, and access networks.

- *Storage area networks* (SANs): corporate data storage networks that connect backend storage disks via high-speed interfaces using primarily a fiber channel protocol.
- *Metropolitan area networks* (MANs): networks that connect data and voice traffic at the city level, typically using synchronous optical network (SONET) rings.
- *Wide area networks* (WANs): networks that connect multiple corporate locations or cities across long distances, also known as *core* or *long-haul networks*. WANs use optical fiber infrastructure exclusively.

Access Networks Access networks are used to gain access to broadband Internet infrastructure. At home we might use a conventional, old-fashioned telephone dial-up, a DSL, or a cable modem technology to make that connection. Some of us might use wireless via advanced cellular phones or by hooking up our laptop using a WiFi 802.11 wireless card. Only a lucky few have access to dedicated optical fiber connections with practically unlimited bandwidth. An *access network* is defined somewhat more precisely as the portion of a public switched network that connects central office equipment to individual subscribers. All these access connections, coming from multiple users, somehow have to be groomed and delivered to Internet service providers (ISPs), companies that provide Internet access service in terms of content. In general, access networks provide data, video, and voice connectivity to all required locations for both consumers and corporate customers using a number of protocols and technologies.

Local and Storage Area Networks University campus networks or networks in large companies are examples of LAN networks. Local area networks connect PCs, workstations, printers, and other devices inside a building or campus, typically using an Ethernet protocol. A LAN network starts with a single user in a corporate environment, but where does it end? Typically, a local area network is connected to the public network via a firewall. The firewall provides data and security protection for a business. Firewalls also provide a convenient demarcation point between LAN and WAN/MAN infrastructure. LAN networks typically use copper wires and coaxial connections, as these are more readily available than optical fiber, and the distances involved are short, hundreds of meters at most. As a result, the use of optical fiber in LAN networks is quite low, nonexistent in most cases.

Storage area networks (SANs) are specialized local area networks that deal exclusively with storage devices such as tape drives, disks, and storage servers (Figure 1.2). Storage networks that connect backend storage devices via high-speed interfaces use a protocol called a *fiber channel*. We do not discuss LAN or SAN networks much in this book, however, as our focus is on core data networks.

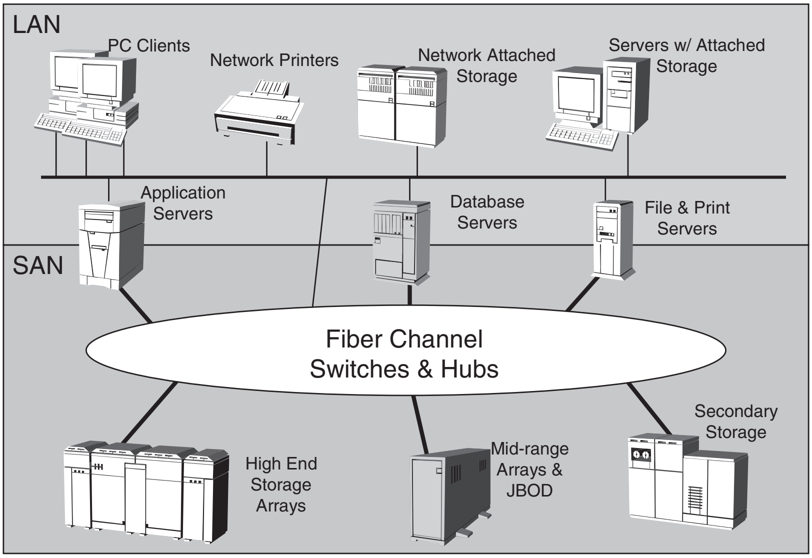

FIGURE 1.2 LAN and SAN networks.

Wide and Metropolitan Area Networks Wide area networks can be divided into metropolitan area networks (MANs) and long-haul networks. Long-haul networks provide transmission services over long distances, typically over hundreds of kilometers. Consequently, they use big "pipes" to carry the traffic and their main service is to deliver from point A to point B. Heavy trucks on a transcontinental highway system are a good example of a long-haul network. MANs encompass large metropolitan areas and therefore cover distances of about 80 to 120 km. The size of the data "pipe" is smaller than in a long-haul network, and services become more varied. Small distribution trucks in a city would be a good analogy for metropolitan area networks.

What is required to build economical wide area networks? First, a network should be flexible, which can be accomplished by providing a large number of interchanges to offload traffic at various points in the network. Second, means of going long distances need to be provided. With optical transmission this is not a problem under certain conditions. Finally, if many "cars" could use the same "highway," the networking system would be able to ensure large throughput. Fortunately, a technology that does precisely that wavelength-division multiplexing, has been developed and is discussed in detail in Chapter 3.

Optical Links Optical links are deployed in each segment (i.e., access, SAN, LAN, MAN, and WAN) of the global networking infrastructure. Networks such as long-haul have lots of optical fiber, whereas some, such as the LAN of a small company, have little or none. Very few networks will be purely optical, but more often than not, networks will involve both optical and electrical components.

1.3.2 Network Topologies

Many network topologies exist in optical data networks, from point to point, to hub, to ring, to fully meshed networks. Each network has advantages and disadvantages. Schematic representations of all of these topologies are shown in Figure 1.3.

Point-to-Point Topology The simplest network topology is point to point, shown schematically in Figure 1.3(a). In this configuration network management is straightforward, but the link has low reliability. If the fiber is cut between points A and B, there is no way to recover. Point-to-point links are sometimes used in long-haul networks, such as in cabling under oceans.

Hub Topology The hub network architecture accommodates unexpected growth and change more easily than do simple point-to-point networks. A hub concentrates traffic at a central site and allows easy reprovisioning of the

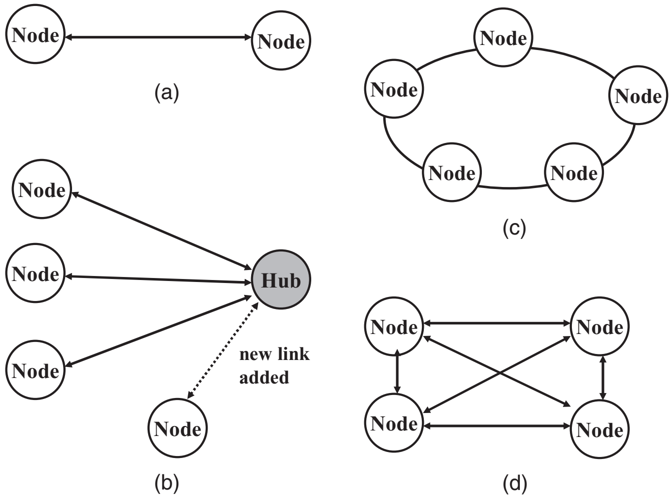

FIGURE 1.3 (a) Point-to-point, (b) hub (star), (c) ring, and (d) mesh network topologies.

circuits. A process of adding one more node in the hub configuration is simple and is shown in Figure 1.3(b). The hub configuration, however, still has the "single point of failure" problem. Ethernet hubs are frequently used in LANs.

Ring Topology Ring topology relies on a connectivity of the network nodes through a circular type of arrangement, as shown in Figure 1.3(c). The main advantage of the ring topology is its survivability: If a fiber cable is cut, the nodes have the intelligence to send the services affected via an alternative path. SONET/SDH uses ring topology predominantly.

Mesh Topology A meshed network consists of any number of sites connected together arbitrarily with at least one loop. A fully meshed topology is shown in Figure 1.3(d) for a simple network of four nodes. A typical characteristic of a meshed network is that nodes within the network can be reached from other nodes through at least two distinct routes. In practice, meshed networks are often constructed using large rings and numerous subrings. Meshed topology can be very expensive in terms of the hardware used. WDM technology frequently has a mesh infrastructure.

Network Topology Comparison All four topologies have been summarized in Table 1.1. It should be noted that other intermediate forms of network topologies exist. Star or bus topologies are very similar to the hub topology, tree topology is a combination of the hub and point-to-point configurations, meshed networks can be only partially meshed, and so on. These intermediate

TABLE 1.1 Network Topology Comparison

Topology	Benefits	Shortcomings	Example
Point-to-point	Very simple	Single point of failure	Long-haul links
Hub	Simple	Single point of failure	Ethernet LAN
Ring	Some redundancy	Scalability problems	Metropolitan SONET/SDH
Mesh	Full redundancy	Complex, hardware intensive	WDM core

forms retain, to a large extent, the benefits and shortcomings presented in Table 1.1.

1.3.3 Circuit vs. Packet Switching

Transmission in telecommunications networks is digital by nature, and the transmission medium of choice is fiber. But how are the ones and zeros to be arranged? At what speed are they to travel? What route should they take? Answers to questions such as these have taken many forms and have made for the most complicated aspect of telecommunications networks. In this section we introduce some basic networking concepts. We discuss circuit and packet switching and describe ways of multiplexing the data. To start, then, consider that points A and B in Figure 1.4 want to exchange information and that the network that will enable the desired connection is as shown. How can the connection get established? There are two ways in principle: circuit switching and packet switching.

Circuit Switching In circuit switching, a dedicated connection is established for the time of the data transfer between points A and B. This process seem to be straightforward (Figure 1.5), although in practice, finding an available

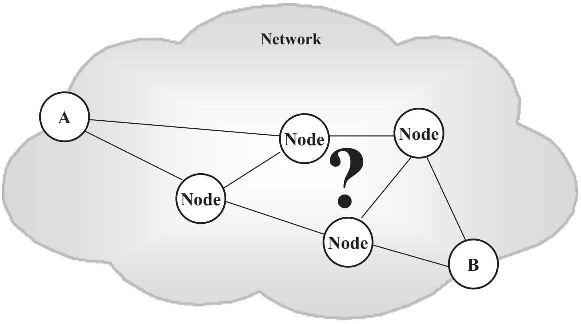

FIGURE 1.4 Network connection problem between points A and B.

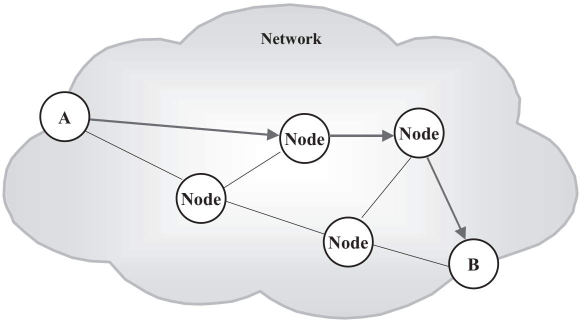

FIGURE 1.5 Circuit-switching principle. A permanent connection is established between points A and B.

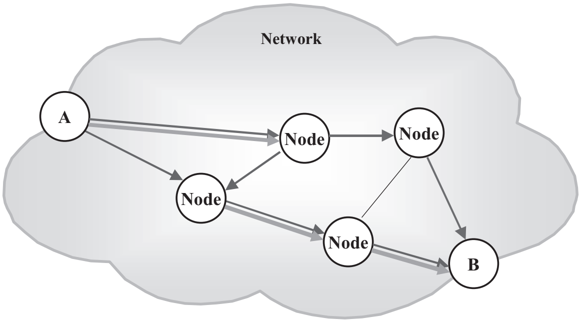

FIGURE 1.6 Packet-switching principle. Each packet travels using a different route from paint A to paint B, as indicated by arrows.

connection in a congested network is often difficult. Telephone services use circuit-switching connections. Once the connection is established, you can talk as long as you want. During your talk the network circuits (i.e., in practice, the wires and switch ports) are open, while other hardware pieces remain dedicated to your phone conversation. This means that nobody else can use them at the same time. As a result, the circuit-switching system is simple and reliable, but not efficient. Even if you do not say anything for minutes, the connection is reserved for you.

Packet Switching In packet switching, each piece of data, or *packet*, can travel from A to B using a different path. The process is illustrated in Figure 1.6 for three packets taking different routes. Packet switching is an example of connectionless technology. Internet protocol (IP) is a routing protocol used to find suitable routing paths in Internet networks. IP uses a packet-switching principle.

Switching Technology Comparison What are the benefits and shortcomings of packet-switching concept compared to circuit switching? Packet switching is more complex, as each piece of data is switched differently. It is also less reliable, as the packets reach destination point B at different times. In fact, because a network has to ensure that each packet reaches its destination in a reasonable time, looping around in large networks has to be avoided. Packet switching has two huge advantages, though: resource sharing and flexibility. Network resources are used only when needed, and as soon as a resource becomes available, it can be used to serve another connection. In this sense, resources are shared between various connections. In addition, this model is flexible, as data packets can be sent using the routes that are available at that particular moment in time.

The battle between circuit switching and packet switching has been one of the most interesting technology battles to watch. On one side you have public switched telephone network and traditional telephone service. On the other side, you have IP and related voice over IP (VoIP) technologies. In a general sense, circuit switching is a centralized model. To establish a global path between A and B, some authority, in the form of a central management system, has to decide where that dedicated path is established. Packet switching, on the other hand, is a decentralized system. Local paths between network elements can be established locally only by viewing large networks in close proximity. In this way, there is less need for centralized management, and large portions of system intelligence can be distributed through the network.

1.3.4 Wavelength vs. Time vs. Statistical Multiplexing

There are several way of multiplexing communications signals. Consider your cable television system. To get multiple TV channels, each channel is broadcast on a different frequency. All signals are "mixed" or, to use a more technical term, *multiplexed*, using frequency-division multiplexing (FDM). FDM is rarely used in optical networks but has a close relative called wavelength-division multiplexing (WDM).

Wavelength-Division Multiplexing WDM works on a principle similar to that of FDM. It is a technique of "combining" or multiplexing multiple wavelengths in an optical fiber. The basic concept is illustrated in Figure 1.7. WDM takes optical signals, each carrying information at a certain bit rate, gives them a specific wavelength, and then sends them down the same fiber. Each input optical signal has the illusion of possessing its own fiber. WDM and FDM are conceptually similar, as there is a direct fundamental relationship between frequency and wavelength. On the other hand, FDM equipment for multiplex electrical signals and WDM equipment for multiplex optical signals are very different.

In this chapter we have started to use a highway analogy for communication along optical fiber. In this context, WDM gets more cars to travel. It

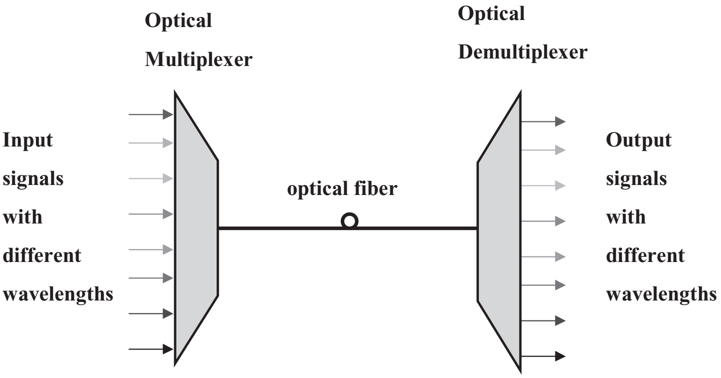

FIGURE 1.7 Principle of wavelength-division multiplexing.

does this not by increasing their speed but by making them travel in parallel in their own dedicated lanes. Traffic in each lane can travel at different speeds, as each lane is independent. The wavelengths used for WDM are chosen in a certain range of frequencies, and details of this selection are discussed in Chapter 3.

Synchronous Time-Division Multiplexing Another method of multiplexing signals is called time-division multiplexing (TDM). The basic concept is illustrated in Figure 1.8. In this example, three data streams (A, B, and C) are multiplexed in a time domain in the following way: Each stream is assigned a fixed time slot in a multiplexed output stream. In time-synchronized TDM, it does not matter whether the stream has any data to send. If it does not, the

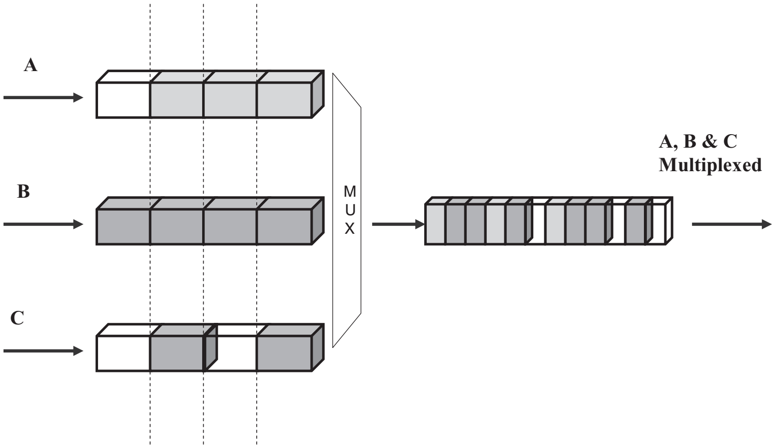

FIGURE 1.8 Principle of synchronous time-division multiplexing.

time slot is simply wasted. As a result, the multiplexed system retains all empty slots of the input data stream, so the TDM system is not efficient in this regard.

Note that for the system shown in Figure 1.8, the bit period of the output stream is three times smaller than the bit period of the input stream. To put this differently, the output bandwidth is three times greater than the bandwidth of an individual input stream. Therefore, in the example given, the networking clock for the output network needs to be three times as fast as the networking clock for the input network. As an example of networking application, TDM has been used in the multiplexing of voice signals. In the early 1960s, Bell Labs engineers created a voice multiplexing system that digitized a voice sample into a 64-kB/s data stream. They organized this data stream into 24-element frames with special conventions for determining where all the bit slots were positioned. The frame was 193 bits long and created an equivalent data rate of 1.544 Mb/s. This rate is referred to as *T1*. European public telephone networks modified the Bell Lab approach and created *E1*, a multiplexing system for 30 voice channels running at 2.048 Mb/s.

Based on these T1 and E1 concepts, the entire hierarchy of T- and E-type signals was created. For example, T3 was created to have a bandwidth three times greater than that of T1. Figure 1.8 can serve as a conceptual representation of multiplexing three T1 signals into a T3 data stream. Multiples of T1 were used to create the SONET multiplexing hierarchy. For example, 84 T1 signals create STS-1, which serves as a building block for SONET and has the rate of 155 Mb/s (84 times 1.544 Mb/s). Similarly, E1 was used to create the synchronous digital hierarchy (SDH) system in Europe. We talk about SONET and SDH in much more detail in Chapter 4.

Asynchronous Time-Domain Multiplexing A different technique for multiplexing relies on an asynchronous principle and, as it turns out, is well suited for packet switching. Asynchronous multiplexing takes packets from each data stream and allocates them in order in the output queue. The process shown conceptually in Figure 1.9 looks deceivingly simple. In

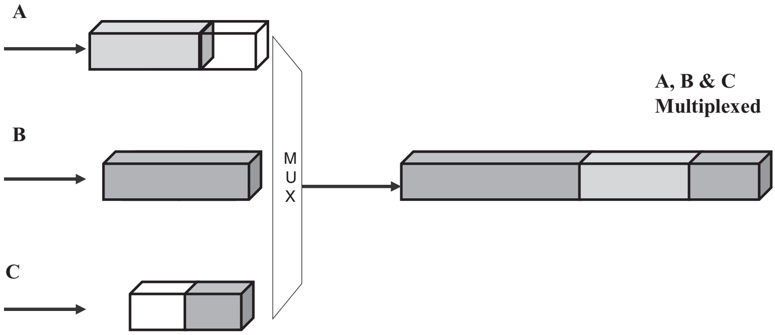

FIGURE 1.9 Principle of statistical multiplexing.

practice, one has to realize that packets arrive at different times and have different lengths, so some form of packet buffering is required. In addition, a recipe for packet ordering is required. Overall, the process of packet switching can be a very complex, and we discuss it later in the book. The big advantage of asynchronous multiplexing is that empty slots in input data streams are effectively eliminated. Only "real" data packets are included in the output queue. As a result, this technique has the potential for a statistical throughput gain.

1.4 OPEN SYSTEM INTERCONNECTION MODEL

1.4.1 Basic Concept

The open system interconnect (OSI) model is somewhat abstract, but is nevertheless a very useful tool in understanding networking concepts. To describe it clearly, we will first use a loose analogy: Imagine that two ships on the Atlantic Ocean are passing each other. Assume that there is a Chinese cook on the first ship and a French cook on the second ship, and they want to exchange recipes. There are a few problems they have to overcome. First, the French cook does not speak Mandarin or Cantonese. In addition, the Chinese cook does not speak French. Fortunately, English translators are available on both ships, so both cooks get their recipes translated into English.

Although both sets of recipes are translated into a common language, there is still a problem. Due to the distance involved, neither party can talk to the other. Fortunately, the ships are close enough that people can see each other. Both have specialists in semaphore on board who can communicate using flags. The translation process goes like this. First, the Chinese recipe is translated from English into semaphore and signaled by the flag so the other party can read the message. On the other side the semaphore signal is received and decoded into English. Subsequently, the English text is translated into French so that the French cook can read it and cook a Chinese dinner for his crew. The process can be repeated as many times as needed as long as visual flag connection is maintained between the two parties (and obviously, this process works in reverse as well).

As we can see from this simple example, the communications problem can be viewed hierarchically. The first level, layer 1, is the physical layer at which the communication has occurred. In our example this is the semaphore signaling layer using flags. At this level the signals are physical in nature. The flag is up or down, and it would be very difficult for an untrained observer to understand what the message is saying. In digital communication, the corresponding physical signal is a stream of zeros and ones.

The second layer, layer 2, is the English language layer. Anyone who speaks English can participate in communications at this level. But people who use different language, or in a digital communication sense, different protocols, cannot participate effectively unless English text is translated into their languages. Obviously, people can speak two or more languages, but that is another story.

Finally, the third layer, layer 3, is the cooking layer. Only cooks or people skilled in cooking can participate in communications at this level. But people who do not cook will not know what a phrase such as "make French sauce" means and will therefore be excluded from communication.

1.4.2 OSI Model and Data Encapsulation

The OSI model presented in Figure 1.10 is built on the layering principles outlined in our example. *Layer 1* is a *physical* or *transport layer*. This layer is concerned with electrical and optical signals, their amplitude, jitter, frequency of operation, and so on. An example of the transport layer is the SONET protocol, discussed in detail in Chapter 4.

Layer 2 is a *data link layer*. The data link is concerned with moving data across the physical links in the network. Layer 2 is responsible for switching data, typically using some form of data address to determine where to direct the message. The data link layer ensures that the connection has been set up, formats data into chunks called *frames*, and handles acknowledgments from a receiver that the data have arrived successfully.

Layer 3 is a *routing layer*. The routing layer is concerned with finding the next routing point in the network to which a message should be forwarded toward its end destination. It decides which way to send the data based on the

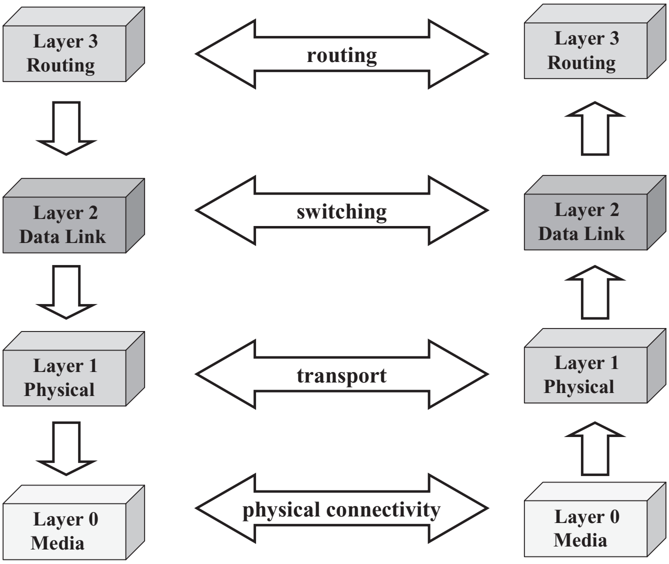

FIGURE 1.10 OSI model for hardware layers.

device's understanding of the state of the network to which the routing device is connected. An example of the routing protocol is IP, which uses a number of storage mechanisms to create and maintain tables of available routes and their conditions. The routing table information, coupled with distance and cost algorithms, is used to determine the best route for a given packet. IP routing is a complex technology that we discuss in later chapters.

There is one additional layer indicated in Figure 1.10 which has not been described so far. It is *layer 0*, a *media layer*. This additional layer represents media and technologies that operate at the media level. As discussed earlier, the media might be UTP copper wire, coaxial cable, or optical fiber, with optical fiber being the medium of choice in this book. Wave-division multiplexing is an example of a layer 0 technology, as it effectively "multiplies" the fiber capacity, as explained in Chapter 3.

One word about the naming convention is in order before we continue. Depending on where we are in the OSI hierarchy, we might use different names for the communication signal. At layer 0 we will probably talk about *optical pulses* or *electrical currents*. At layer 1 we will probably use the term *frame*: for example, *SONET frame*. At layer 2 we might also use the term *frame*, or if the frame is of constant length, we will call it a *cell*, as in asynchronous transfer made (ATM) *cell*. Finally, at layer 3 we usually refer to a data signal as a *packet*.

It is important to realize that frames, cells, and packets have generally similar structures, as shown in Figure 1.11. They all consist of the real data they are carrying and some additional overhead information. The real data are the

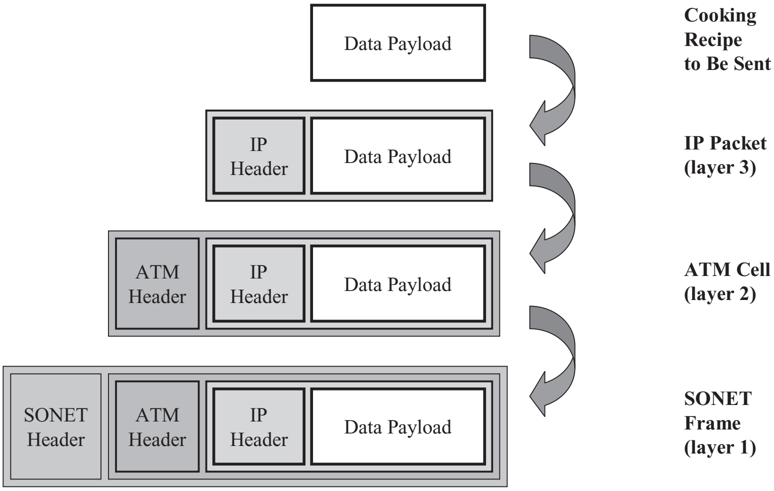

FIGURE 1.11 Process of data encapsulation.

messages to be transferred across the network, referred to as a *data payload*. Most of the data are in the data payload section, although the amount of data overhead will vary depending on the particular protocol implementation. The overhead information can either be appended before the payload in the form of a header or after the payload in the form of a trailer, and is frequently appended at both ends. For simplicity of illustration we assume in this chapter that only the header is attached.

As illustrated in Figure 1.11, when the data travel in the OSI hierarchy, its overhead information changes. Assume that the aforementioned cooking recipe is the data payload to be sent across the network. At the IP routing layer the packet consists of the data payload (the recipe) and IP header. When sent to a data link layer, a frame appends its own header (in the example shown, it is the ATM header). Finally, when sent down to layer 1, the transport protocol, in this case SONET, appends its own header. The entire process is reversed on the opposite end. First, the transport header is stripped off, and the remaining information is sent to the data link layer. At layer 2 the switching header is removed and the remaining information is sent to the routing layer. Finally, at the routing layer the layer 3 header is removed and only the data payload remains intact.

1.4.3 Network Overlay Hierarchy

As a result of the OSI layer structure, the networks are frequently overlaid one on top of the other. A typical overlay situation is shown in Figure 1.12. IP networks are overlaid on top of layer 2 networks such as ATM and Ethernet

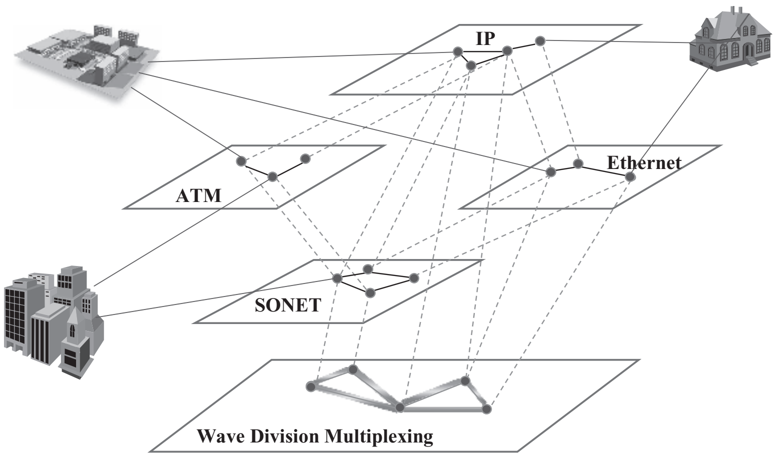

FIGURE 1.12 Overlay network hierarchy.

networks. These in turn can be overlaid on top of layer 1 networks such as SONET networks. Finally, all of the networks above can be overlaid on top layer 0 WDM networks. The degree of overlay can vary and in numerous cases, some layers are omitted. As you can imagine, this fact creates a large number of possible scenarios. One popular example of this networking overlay is the following combination: IP on top of ATM on top of SONET on top of WDM.

1.5 NETWORKING EQUIPMENT

Networking hardware equipment consists of various networking "boxes" that are used to build broadband data networks. To connect your computing resources and enjoy the benefits of network computing, you need these "boxes" filled with processors, switches, and cabling interfaces. One might argue here that the term *boxes* is somewhat of an understatement, as the box can be quite powerful, as illustrated in Figure 1.13, which uses a core router as an example. We might also want to mention that a core router can easily cost over $1 million. After this example, hopefully, networking boxes have gained some respect in your eyes. Corresponding to their new, more appreciated status for networking boxes, we instead use the terms *network elements* or *networking equipment* for the remainder of the book.

We can divide networking elements according to the OSI layer number at which they operate, as illustrated in Figure 1.14. For example, IP routers

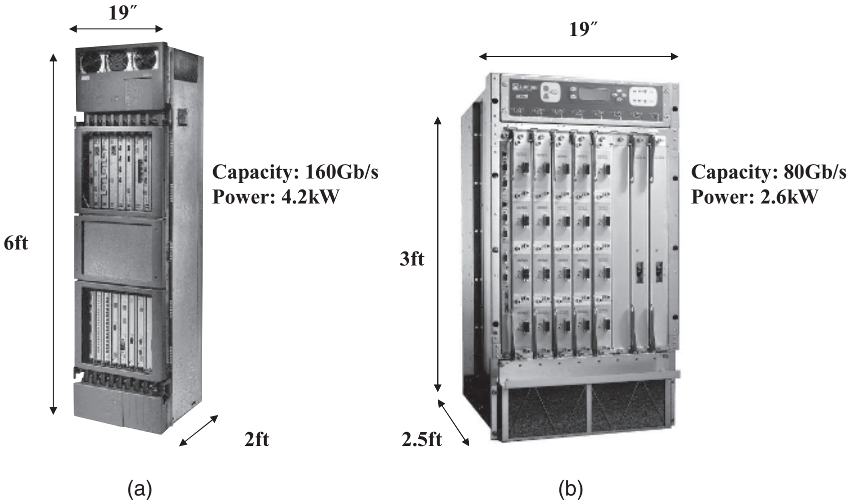

FIGURE 1.13 Examples of core routers: (a) Cisco GSR 12416; (b) Juniper M160. (Courtesy of Cisco and Juniper Corporations.)

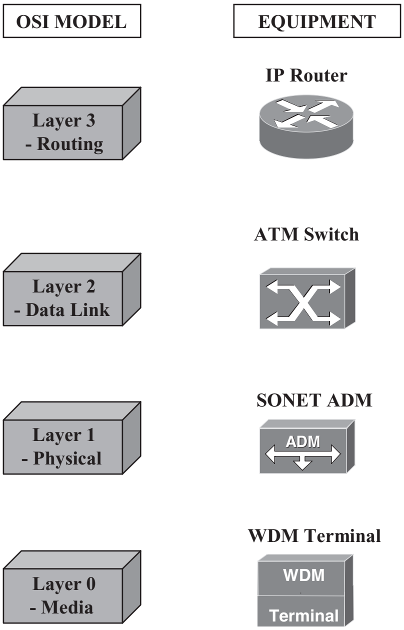

FIGURE 1.14 Networking elements and corresponding OSI layer hierarchy.

operate at routing layer 3, while switches operate at switching layer 2. Regenerators, modems, hubs, and add/drop multiplexers operate at layer 1.

1.5.1 Regenerators, Modems, Hubs, and Add–Drop Multiplexers

Regenerators A *regenerator* is a simple networking element operating on electrical bits. It does not add functionality in terms of traffic handling but is present in the network for transmission purposes. The regenerator regenerates a signal back to its original shape by processes of amplification and/or regeneration. In optical networks, a popular method of optical signal regeneration is to convert optical pulses (O) to electrical signals (E) and then back again to the optical domain (O). We discuss this O-E-O conversion process later in the book.

Modems *Modems* convert digital data into an analog signal and then back again. The conversion process involves *modulation* of the digital signal and, in the reverse direction, *demodulation* of the analog signal into the digital domain: hence the name *Modem*. A typical example of the modem is the device in your PC used to connect to the Internet for dial-up connections capable of 56-kb/s bandwidth. Similar devices for DSL networking are called *DSL modems* and for broadband cable connection are called *cable modems*. All modems use some form of compression and error correction. Compression algorithms

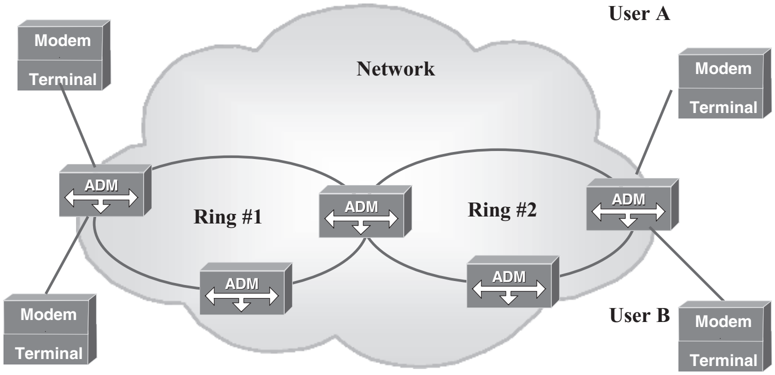

FIGURE 1.15 Add–drop multiplexer function in the network.

enable throughput to be enhanced two to four times over normal transmission. Error correction examines incoming data for integrity and requests retransmission of a packet when it detects a problem.

Hubs A *hub* is a device that aggregates data traffic from many low-bandwidth connections. In local area networks hubs are connected to PCs, workstations, and servers. *Passive hubs* simply provide cable connectivity, while *active hubs* also provide management capabilities.

Add–Drop Multiplexers A device that resembles a hub, called an *add–drop multiplexer* (ADM), is shown in Figure 1.15. ADMs are typically connected to form rings, but each device can add or drop some traffic coming from network points outside the ring. ADMs are used in very large quantities in metropolitan SONET/SDH networks.

1.5.2 Switches

A *switch*, as the name implies, switches frames between various ports. It assigns dedicated bandwidth to be designated to each device on the network connected to the switch. Switches split large networks into smaller segments, decreasing a number of users sharing the same resources. Typical examples of switches include ATM switches in wide area networks, Ethernet switches in local area networks, and fiber channel switches in storage area networks. A switch function is illustrated in Figure 1.16.

1.5.3 Routers

Routers direct data packets, following IP rules, from one network to another. To accomplish this task they examine the content of data packets flowing

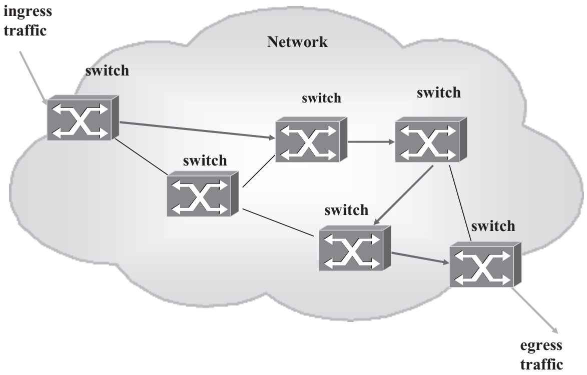

FIGURE 1.16 Switch function in the network.

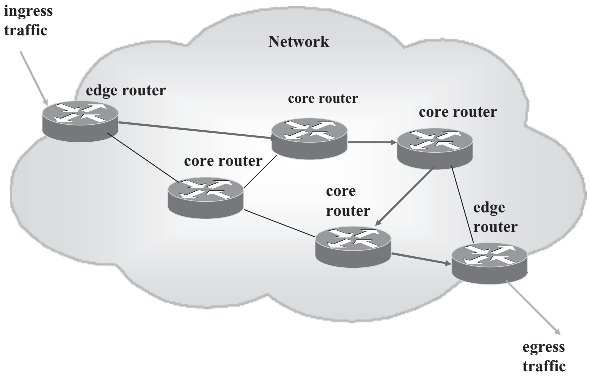

FIGURE 1.17 Router function in the network.

through them. Routers have to determine the most efficient path through the network using complex routing algorithms. After finding the most efficient path, routers switch frames between various ports. In this sense, routers perform the same switching function as that of switches. Routers are much more complex devices, however, as they have to deal with large routing tables for the global Internet and find appropriate routing addresses. Switches, on the other hand, merely switch frames and are aware only of their close network proximity. In terms of packet processing functions, routers perform an order of magnitude of many operations compared to switches.

Router equipment can be classified as core routers, and edge routers as shown in Figure 1.17. *Core routers* reside in the core of the network, so they typically have more capacity than edge routers. *Edge routers*, on the other hand, have more diversified functions than core routers, as they reside on the edge of the network and have to deal with traffic coming from access and MAN networks arriving in various protocols and speeds. Core routers have “big pipe” interfaces operating at Gb/s rates, whereas edge routers have to deal with many “smaller pipes” with bandwidths as low as a few Mb/s.

1.5.4 Networking Service Models

Having described briefly various pieces of networking equipment, we can now start putting together some simple networks and see what services can be offered to network customers. Let us utilize layer 1 equipment to form a transport network as shown in Figure 1.18. Customers with locations A and B will need to have some routers, as routers are always needed to find a path through the network. We can assume that routers belong to a customer and are installed on customer premises. As such, they belong to a customer premise equipment class. A network, on the other hand, belongs to a service provider. In the example shown, the network consists of SONET add–drop multiplexers. In a service model called a leased line, the service provider can lease one particular connection in the service provider network to be assigned permanently to locations A and B. The connection will be dedicated entirely to this customer, and nobody else will use any portion of it. The connection is secure, always available, but expensive, as service providers have to dedicate their equipment on that leased line entirely to one customer.

If a service provider has a layer 2 network available, which it normally does, it can offer a different service, as shown in Figure 1.19. This type of service is called a *virtual private network* (VPN). User traffic from point A is sent to a layer 2 network. After traversing through this switching network it reaches

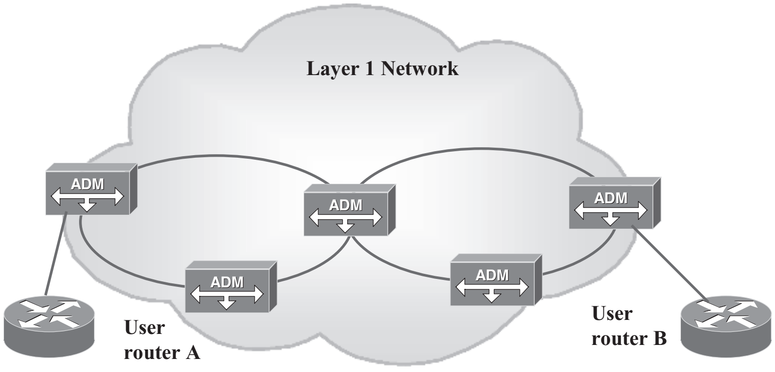

FIGURE 1.18 Leased line layer 1 transport network.

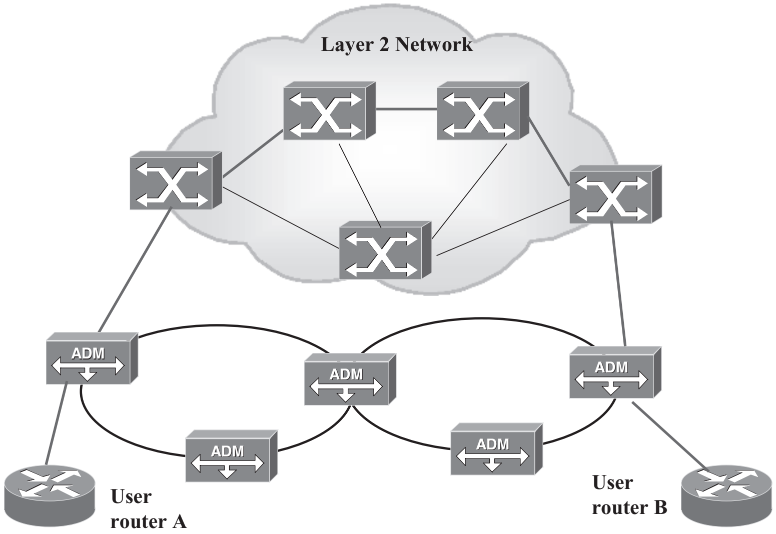

FIGURE 1.19 Virtual private network.

destination B. It is worthwhile to note that in this example, somewhere on the edges of the layer 2 network, layer 1 equipment is involved in the data transport as well, though only to a limited extent.

VPN service has some interesting advantages over leased lines. It uses the network resources much more efficiently, as resources can be shared across the entire customer base, and layer 2 networks can take advantage of statistical multiplexing. Some customers might have some apprehension regarding the security of the VPN connection, as the packets from one customer "mix" somewhere in the system with packets from another customer. In a properly engineered VPN system, though, security should not be a concern.

It is important to realize that from a customer's point of view, the statistical nature of a layer 2 network is not visible. From his or her point of view, the connection between A and B is his or her private connection, hence the term *virtual private network*. Finally, it has to be pointed out that to create a VPN in Figure 1.19, we used a layer 2 switched network. It is equally possible to create a layer 3 VPN, where the routing network is used to achieve the same function. VPN layers 2 and 3 are both offered commercially by service providers, in addition to traditional leased-line models based on layer 1 networks.

In closing, this chapter has shown for us that networking equipment comes in various sizes, "flavors," and functionality subsets. Figure 1.20 lists a number of network elements and their corresponding product naming, depending on whether they are deployed at LAN, MAN, or WAN networks.

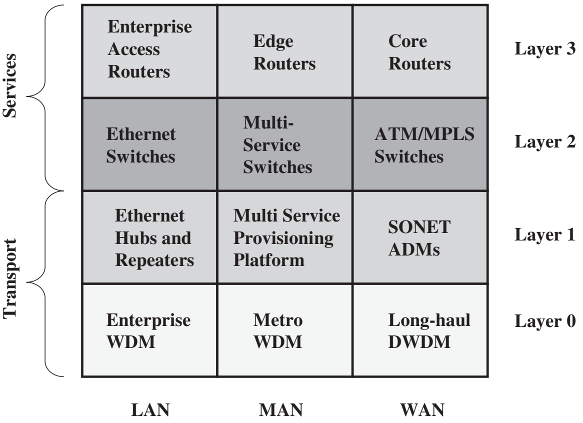

Services	Enterprise Access Routers	Edge Routers	Core Routers	Layer 3
	Ethernet Switches	Multi-Service Switches	ATM/MPLS Switches	Layer 2
Transport	Ethernet Hubs and Repeaters	Multi Service Provisioning Platform	SONET ADMs	Layer 1
	Enterprise WDM	Metro WDM	Long-haul DWDM	Layer 0
	LAN	MAN	WAN	

FIGURE 1.20 Various examples of networking equipment.

We discuss network elements' functionality in great detail in Chapters 12 and 13.

KEY POINTS

Transmission media:

- Four transmission media are available for data communication: copper wire, coaxial cable, optical fiber, and air (wireless connection).
- Optical fiber is the most efficient medium, as its has the lowest signal attenuation and is insensitive to electrical noise.

Network classification:

- Five distinct classes of networks are present in global Internet infrastructure: access networks, local area networks, storage area networks, metropolitan area networks, and wide area networks.
  - Access networks connect consumers and corporate users with Internet infrastructure.
  - Local area networks connect multiple users in a contained environment like a corporate building or a campus. Typically, Ethernet is used as the protocol of choice.
  - Storage area networks are corporate data storage networks that connect backend storage disks via high-speed interfaces using primarily fiber channel protocol.

  - Metropolitan area networks connect data and voice traffic at the city level typically using SONET rings.
  - Wide area networks connect multiple corporate locations or cities across long distances, also known as core or long-haul networks.

Network topologies:

- Various network topologies are used in different networks: point to point, hub, ring, or mesh. Each has advantages, disadvantages, and a particular level of complexity.
  - Point-to-point links are the simplest form of network. Typically, they are used in long-haul connections: for example, under the Atlantic Ocean.
  - Hub networks aggregate multiple connections into one higher-throughput connection. Ethernet hub topology is frequently employed in local area networks.
  - Ring topology introduces the possibility of higher reliability at the cost of more hardware. SONET/SDH rings are frequently used in metropolitan area networks.
  - Mesh topology introduces the possibility of large redundancy, but it is more expensive in terms of hardware complexity. WDM links are used to create mesh in some core networks.

Switching concepts:

- To establish a connection between two points in the network, two techniques can be used: circuit switching and packet switching.
  - In circuit switching, a dedicated connection is established for the time needed. Telephone service is an example of the circuit-switching technique.
  - In packet switching, each packet travels across the network using a different path. Internet protocol routing is an example of the packet-switching technique.

Multiplexing concepts:

- There are three ways to multiplex data in optical networks: wavelength-division multiplexing (WDM), synchronous time-domain multiplexing (TDM), and asynchronous TDM.
  - WDM takes optical signals, each carrying information at a certain bit rate, gives them a specific wavelength, and sends them down the same fiber. As a result, each optical signal has the illusion of having its own fiber.

- Synchronous TDM takes multiple synchronized streams of data, each carrying information at a certain bit rate, and assigns each piece of data a precise time slot in the output stream.
- Asynchronous TDM takes multiple asynchronous variable-size packets and assigns them in the output stream. The length of the time slot is determined based on the relative needs of input data streams. Asynchronous TDM is more efficient but much more complex than synchronous TDM. This technique requires the use of packet buffering and queuing.

OSI model:

- An open system interconnect (OSI) model is a very useful conceptual model used to classify various networking functions.
- OSI has the following hardware layers: network (layer 3), data link (layer 2), transport (layer 1), and media (layer 0).
- The data encapsulation at layer 3 is usually referred to as a packet, at layer 2 as a frame or cell, and at layer 1 as a frame or electrical/optical signal.
- The typical protocol stack in today's optical networks is IP over ATM over SONET over WDM.

Networking equipment:

- Modems, regenerators, and hubs belong to the layer 1 class of networking equipment. They perform physical functions such as sending or receiving data at the terminal nodes (modem), regenerating signals on their way (regenerator) or combining signals into a larger data stream (hub). DSL modems, SONET regenerators, and Ethernet hubs are examples of this class of equipment.
- Switches belong to the layer 2 class of networking equipment. Switches perform switching functions by sending frames from inputs to desired output ports at the switch I/Os. Core ATM switches, enterprise Ethernet switches, and director fiber channel switches are examples of this class of equipment.
- Routers belong to the layer 3 class of the networking equipment. Routers perform the complex task of finding the most effective way of sending data packets through the global network. Core IP routers and edge routers are examples of this class of equipment.

REFERENCES

Agarwal, G. P., *Fiber-Optic Communication Systems*, Wiley, Hoboken, NJ, 1997.

Freeman, R., *Fiber-Optic Systems for Telecommunications*, Wiley, Hoboken, 2002.

Goralski, W., *SONET: A Guide to Synchronous Optical Networks*, McGraw-Hill, New York, 1997.

Kartalopoulos, S., *Next Generation Sonet/SDH*, IEEE Press, Piscataway, NJ, 2004.

Mukherjee, B., *Optical Communication Networks*, McGraw-Hill, New York, 1997.

Ramaswami, R., and K. Sivarajan, *Optical Networks: A Practical Perspective*, Academic Press, San Diego, CA, 1998.

Tomsu, P., and C. Schmutzer, *Next Generation Optical Networks*, Prentice Hall, Upper Saddle River, NJ, 2002.