6.11 Advanced Topics: Networks

Networks are growing in popularity over time, and unlike other I/O devices, there are many books and courses on them. For readers who have not taken courses or read books on networking, this section gives a quick overview of the topics and terminology, including internetworking, the OSI model, protocol families such as TCP/IP, long-haul networks such as ATM, local area networks such as Ethernet, and wireless networks such as IEEE 802.11.

Networks are the major medium used to communicate between computers. Key characteristics of typical networks include the following:

- *Distance:* 0.01 to 10,000 kilometers
- *Speed:* 0.001 MB/sec to 1000 MB/sec
- *Topology:* bus, ring, star, tree
- *Shared lines:* none (switched point-to-point) or shared (multidrop)

We'll illustrate these characteristics with examples, starting from top down.

Internetworking

Undoubtedly one of the most important innovations in the communications community has been internetworking. It allows computers on independent and incompatible networks to communicate reliably and efficiently.

The low cost of internetworking is remarkable. For example, it is vastly less expensive to send electronic mail than to make a coast-to-coast telephone call and leave a message on an answering machine. This dramatic cost improvement is achieved using the same long-haul communication lines as the telephone call, which makes the improvement even more impressive.

The enabling technologies for internetworking are software standards that allow reliable communication without demanding reliable networks. The underlying principle of these successful standards is that they were composed as a hierarchy of layers, each layer taking responsibility for a portion of the overall communication task. Each computer, network, and switch implements its layer of the standards, relying on the other components to faithfully fulfill their responsibilities. These layered software standards are called *protocol families* or *protocol suites*. They enable applications to work with any interconnection without extra work by the application programmer.

The goal of a family of protocols is to simplify the standard by dividing responsibilities hierarchically among layers, with each layer offering services needed by the layer above. The application program is at the top, and at the bottom is the physical communication medium, which sends the bits. Just as abstract data types simplify the programmer's task by shielding the programmer from details of the implementation of the data type, this layered strategy makes the standard easier to understand.

Open Systems Interconnect (OSI) developed a model that popularized describing networks as a series of layers. Figure 6.11.1 shows the model. Although all protocols do not exactly follow this layering, the nomenclature for the different layers is widely used. Thus, you can hear discussions about a simple layer 3 switch versus a smart layer 7 switch.

The key to protocol families is that communication occurs *logically at the same level* of the protocol in both sender and receiver, but *services of the lower level implement it*. This style of communication is called *peer-to-peer*. As an analogy, imagine that General A needs to send a message to General B on the battlefield. General A writes the message, puts it in an envelope addressed to General B, and gives it to a colonel with orders to deliver it. This colonel puts it in an envelope, writes the name of the corresponding colonel who reports to General B, and gives it to a major with instructions for delivery. The major does the same thing and gives it to a captain, who gives it to a lieutenant, who gives it to a sergeant. The sergeant takes the envelope from the lieutenant, puts it into an envelope with the name of a sergeant who is in General B's division, and finds a private with orders to take the large envelope. The private borrows a motorcycle and delivers the envelope

Layer number	Layer name	Main function	Example protocol	Network component
7	Application	Used for applications specifically written to run over the network	FTP, DNS, NFS, http	Gateway, smart switch
6	Presentation	Translates from application to network format, and vice versa		Gateway
5	Session	Establishes, maintains, and ends sessions across the network	Named pipes, RPC	Gateway
4	Transport	Additional connection below the session layer	TCP	Gateway
3	Network	Translates logical network address and names to their physical address (e.g., computer name to MAC address)	IP	Router, ATM switch
2	Data link	Turns packets into raw bits and at the receiving end turns bits into packets	Ethernet	Bridge, network interface card
1	Physical	Transmits raw bit stream over physical cable	IEEE 802	Hub

FIGURE 6.11.1 The OSI model layers. Based on "OSI Models" by Grant Wilson, *www.geocities.com/SiliconValley/Monitor/3131/ne/osimodel.html.*

to the other sergeant. Once it arrives, it is passed up the chain of command, with each person removing an outer envelope with his name on it and passing on the inner envelope to his superior. As far as General B can tell, the note is from another general. Neither general knows who was involved in transmitting the envelope, nor how it was transported from one division to the other.

Protocol families follow this analogy more closely than you might think, as Figure 6.11.2 shows. The original message includes a header and possibly a trailer sent by the lower-level protocol. The next lower protocol in turn adds its own header to the message, possibly breaking it up into smaller messages if it is too large for this layer. Reusing our analogy, a long message from the general is divided and placed in several envelopes if it could not fit in one. This division of the message and appending of headers and trailers continues until the message descends to the physical transmission medium. The message is then sent to the destination. Each level of the protocol family on the receiving end will check the message at its level and peel off its headers and trailers, passing it on to the next higher level and putting the pieces back together. This nesting of protocol layers for a specific message is called a *protocol stack*, reflecting the last-in, first-out nature of the addition and removal of headers and trailers.

As in our analogy, the danger in this layered approach is the considerable latency added to message delivery. Clearly, one way to reduce latency is to reduce the number of layers. Keep in mind that protocol families *define* a standard, but do not force how to *implement* the standard.

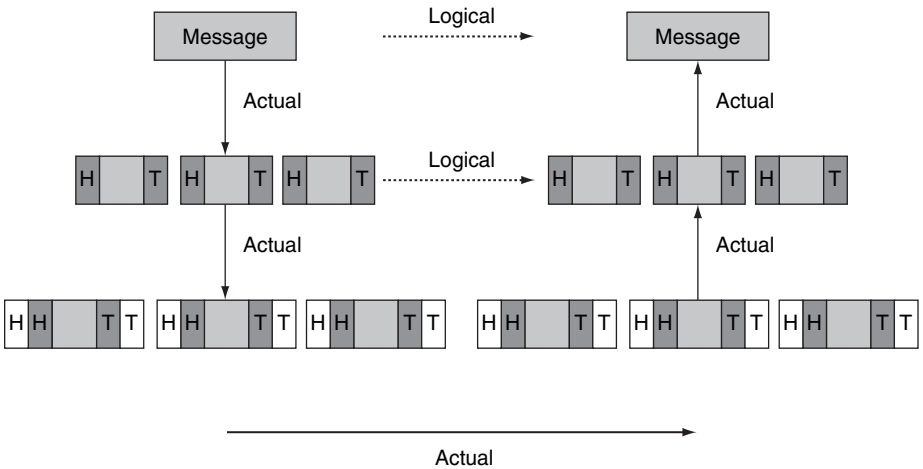

FIGURE 6.11.2 A generic protocol stack with two layers. Note that communication is peer-to-peer, with headers and trailers for the peer added at each sending layer and removed by each receiving layer. Each layer offers services to the one above to shield it from unnecessary details.

Hardware/ Software Interface

To allow communication across multiple networks with different characteristics, TCP/IP defines a standard packet format (see below). An IP packet, which contains Internet addressing information, encapsulates a TCP packet that contains both address information interpreted by the host and the data being communicated. The IP header specifies the length of the IP data (from 1 to 65,536 bytes). Since the TCP header uses 20 bytes of the IP data, the maximum size of the TCP data packet is 65,536 – 20 = 65,516 bytes.

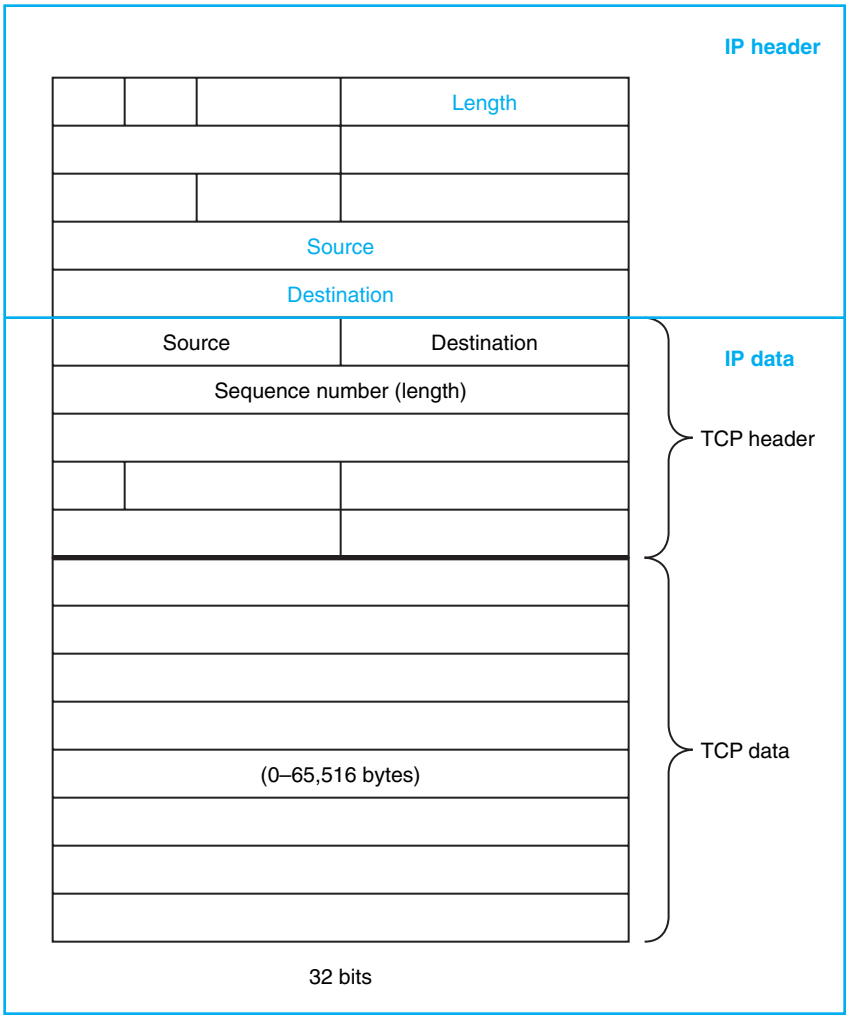

Long-Haul Networks

Long-haul networks cover distances of 10 to 10,000 kilometers. The first and most famous long-haul network was the ARPANET (named after its funding agency, the Advanced Research Projects Agency of the U.S. government). It transferred data at 56 Kbits/sec and used point-to-point dedicated lines leased from telephone companies. The host computer talked to an *interface message processor* (IMP), which communicated over the telephone lines. The IMP took information and broke it into 1 Kbit packets, which could take separate paths to the destination node. At each hop, a packet was stored (for recovery in case of failure) and then forwarded to the proper IMP according to the address in the packet. The destination IMP reassembled the packets into a message and then gave it to the host. Most networks today use this *packet-switched* approach, in which packets are individually routed from source to destination. The ARPANET was the precursor of the Internet. The key to interconnecting different networks was standardizing on a single protocol family, TCP/IP.

Long-haul networks are **switched networks**; switches allow multiple independent communications to occur, unlike shared media. In the limit, there is only one host per link and that host is directly connected to a switch. ATM (Asynchronous Transfer Method) is a scalable network technology (from 155 Mbits/sec to 10 Gbits/sec) that is popular for long-haul networks.

switched network A network of dedicated point-to-point links that are connected to each other with a switch.

The bandwidths of networks are probably growing faster than the bandwidth of any other type of device at present, with optical fibers offering bandwidths at 40 Gbits/sec and above.

Local Area Networks

The **local area network (LAN)** is what is commonly meant today when people mention a network, and **Ethernet** is what most people mean when they mention a LAN. (Ethernet has in fact become such a common term that it is often used as a generic term for LAN.) The basic Ethernet from 1978 was a 10 Mbit/sec, one-wire bus that had no central control. Messages, or *packets*, are sent over the Ethernet in blocks that vary from 64 bytes to 1518 bytes.

local area network (LAN) A network designed to carry data within a geographically confined area, typically within a single building.

An Ethernet is essentially a bus with multiple masters and a scheme for determining who gets bus control; we'll discuss how the distributed control is implemented in the exercises. Because the Ethernet was originally a bus, only one sender can be transmitting at any time; this limits the bandwidth.

Ethernet has been extraordinarily successful. The 100 Mbit/sec standard was proposed in 1994, and many classes of computers include it as a standard interface. Gigabit Ethernet is being deployed today and the 10 gigabit Ethernet standard was ratified in 2003. Ethernet is codified as IEEE standard 802.3.

Ethernet A computer network whose length is limited to about a kilometer. Originally capable of transferring up to 10 million bits per second, newer versions can run up to 1000 million bits per second and even 10,000 million bits per second. It treats the wire like a bus with multiple masters and uses collision detection and a back-off scheme for handling simultaneous accesses.

Computers became thousands of times faster than they were in 1978, but the shared interconnection was no faster for almost 20 years. Hence, past engineers

invented temporary solutions until a faster Ethernet was available. One solution was to use multiple Ethernets to connect machines, and to connect these smaller Ethernets with devices that could take traffic from one Ethernet and pass it on to another as needed. These devices allow individual Ethernets to operate in parallel, thereby increasing the aggregate interconnection bandwidth of a collection of computers.

Figure 6.11.3 shows the potential parallelism. Depending on how they pass traffic and what kinds of interconnections they can put together, these devices have different names:

- *Routers* or *gateways:* These devices connect LANs to WANs, or WANs to WANs, and resolve incompatible addressing. Generally slower than switches, they operate at OSI layer 3, the network layer. Routers divide the interconnect into separate smaller subnets, which simplifies manageability and improves security. There might be one router or gateway per building.
- *Switches:* These devices connect LANs together, passing traffic from one side to another depending on the addresses in the packet. Switches map

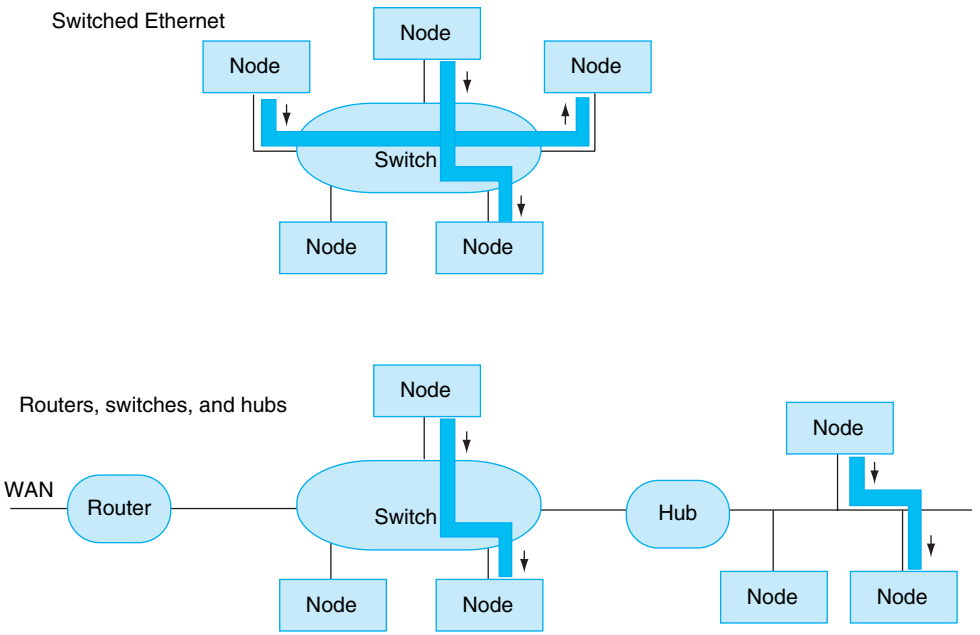

FIGURE 6.11.3 The potential parallelism due to switches, plus the role of routers and hubs. Note that in the lower drawing switches provide traffic isolation—that is, multiple domains—but the hub does not.

Ethernet addresses on each network segment and allow only necessary traffic to pass through them. When a switch receives a packet, it examines the destination and source addresses and compares them to a table of network segments and addresses. The packet is dropped if the segments are the same or forwarded if they are different. Switches operate at the Ethernet protocol level and are usually simpler and cheaper than routers. Using the notation of the OSI model (see Figure 6.11.1), bridges operate at layer 2, the data link layer. Whereas switches were optional for 10 Mbit/sec and 100 Mbit/sec Ethernet, they are required for Gigabit/sec and 10 Gigabit Ethernet. The media cannot be shared at those high speeds.

- *Hubs:* The final network devices are hubs, but they merely extend multiple segments into a single LAN. Thus, hubs do not help with performance, as only one message can transmit at a time. Hubs operate at OCI layer 1, called the physical layer.

Since these devices were not planned as part of the Ethernet standard, their ad hoc nature has added to the difficulty and cost of maintaining LANs.

To see the importance of looking at performance from top to bottom, including both hardware and software, consider the following example.

Performance of Two Networks

EXAMPLE

Let's compare two LANs: 100 Mbit/sec and 1000 Mbit/sec Ethernet.

Characteristic	100 Mbit Ethernet	1000 Mbit Ethernet
Bandwidth from node to network	100 Mbit/sec	1000 Mbit/sec
Interconnect latency	10 μs	10 μs
HW latency to/from network	2 μs	2 μs
SW overhead sending to network	100 μs	100 μs
SW overhead receiving from network	100 μs	100 μs

Find the host-to-host latency for a 250-byte message using each network.

ANSWER

We can estimate the time required as the sum of the fixed latencies plus the time to transmit the message. The time to transmit the message is simply the message length divided by the bandwidth of the network.

The transmission times are

$$\text{Transmission time}_{100\text{Mbit}} = \frac{250 \times 8 \text{ bits}}{100 \times 10^6 \text{ bits/sec}} = 20 \text{ µs}$$

$$\text{Transmission time}_{100\text{Mbit}} = \frac{250 \times 8 \text{ bits}}{1000 \times 10^6 \text{ bits/sec}} = 2 \text{ µs}$$

So the transmission time for the ATM network is about a factor of 9 lower.

The total latency to send and receive the packet is the sum of the transmission time and the hardware and software overheads:

$$\text{Total time}_{100\text{Mbit}} = 10 + 2 + 80 + 2 + 100 + 222 = 214 \text{ µs}$$

$$\text{Total time}_{1000\text{Mbit}} = 10 + 2 + 80 + 2 + 100 + 22 = 196 \text{ µs}$$

The end-to-end latency of 1000 Mbit/sec Ethernet is only about 1.1 times faster, even though the transmission time is 10 times higher!

Wireless Local Area Networks

Thus far we have been assuming that packets travel through copper or fiber. IEEE 802.11, popularly known as WiFi, extended the Ethernet standard to communicate through the air, which comes closer to the original inspiration for its name. Three variations are available in 2006: 802.11b, with a peak bandwidth of 11 Mbits/sec; and 802.11a and 802.11g, both with a peak of 54 Mbits/sec. In practice, the delivered rates in the field are about a third of the peak rates in the lab. Although 802.11b has been widely deployed, the other two are vying to be its successor. In 2004, the sales of new 802.11g devices exceeded 802.11b.

IEEE 802.11 replaces the bottom layers of the OSI standard, which Ethernet labels the MAC layer and PHY layer. Not surprisingly, the physical layer is now radio. Before going further with the standard, let's review radio communication.

carrier signal A continuous signal of a single frequency capable of being modulated by a second data-carrying signal.

A radio wave is an electromagnetic wave propagated by an antenna. Radio waves are modulated, which means that the sound signal is superimposed on the stronger radio wave that carries the data, and hence is called the **carrier signal**. Radio waves have a particular wavelength or frequency: they are measured either as the length of the complete wave or as the number of waves per second. FM radio stations transmit on the band of 88 MHz–108 MHz using frequency modulations (FM) to record the data. By tuning into different frequencies, a radio receiver can pick up a specific signal. Both 802.11b and 802.11g use a band at the 2.4 GHz frequency carrier, and 802.11a uses a band at the 5 GHz frequency carrier. All actually use a small percentage of frequencies on either side of the normative amount, which gives them multiple channels on which to transmit. If transmitters collide, they hop to another channel and try again.

Radio signals are first received by the antenna, amplified, passed through a mixer, then filtered, demodulated, and finally decoded. A mixer accepts two signal inputs and forms an output signal at the sum and difference frequencies. Filters select a narrower band of frequencies to pass on to the next stage. Modulation encodes information in the amplitude, phase, or frequency of the signal to increase its robustness under impaired conditions. Decoding turns signals into information. The antenna acts as the interface between the medium through which radio waves travel and the electronics of the transmitter or receiver. Radio transmitters go through the same steps in the opposite order.

The **bit error rate** (BER) of a wireless link is determined by the received signal power and the noise due to interference caused by the receiver hardware and interference from other sources. Noise is typically proportional to the radio frequency bandwidth.

bit error rate The fraction in bits of a message or collection of messages that is incorrect.

Typically, wireless communication is selected because the communicating devices are mobile or because wiring is inconvenient, which means the wireless network must rearrange itself dynamically. Such rearrangement makes routing more challenging. A second challenge is that wireless signals are not protected and hence are subject to mutual interference, especially as devices move, and to eavesdropping. Power is another challenge for wireless communication, both because the mobile devices tend to be battery powered and because antennas radiate power to communicate and little of it reaches the receiver. As a result, raw bit error rates are typically a thousand to a million times higher than copper wire.

There are two primary architectures for wireless networks: *base station* architectures and *peer-to-peer* architectures. Base stations are connected by wire for longer-distance communication, and the mobile units communicate only with a single local base station. Peer-to-peer architectures allow mobile units to communicate with each other, and messages hop from one unit to the next until delivered to the desired unit. Although peer-to-peer is more reconfigurable, base stations tend to be more reliable, since there is only one hop between the device and the station.

The 802.11 standard is primarily used with base stations. Each base station, called an *access point*, controls a *cell*. The wireless devices in that cell that communicate with the access point are called *end-user stations*. The 802.11b standard advertises that an access point can be up to 1000 feet away from a station if there are no obstructions, and perhaps 100 feet when it is obstructed. A wired Ethernet connects multiple access points together as well as to the Internet.

Since wireless has a higher bit error rate than wired communication, the maximum Ethernet packet of 1538 bytes could be a problem to transmit successfully. Rather than modify the Ethernet format, 802.11 allows the MAC layer to fragment these large messages into several smaller messages. The MAC layer of the

receiving device then reassembles these smaller messages into the original full Ethernet message.

To try to protect against eavesdropping of the wireless broadcast, 802.11 offers Wired Equivalent Privacy. It uses a pseudorandom number generator initialized by a shared secret key. Operators initialize access points and end-user stations with the secret key. A pseudorandom sequence of bits equal to the largest packet is combined with the real packet to encode the packet transmitted in the air.

To save power on the end-user stations, which typically rely on batteries, 802.11 supports stations going into sleep mode without losing information. The access point maintains a list of stations in sleep mode and buffers packets sent to them. When the sleeping stations awake and start broadcasting again, buffered packets are then sent.

Finally, 802.11 supports limited peer-to-peer architecture by defining ad hoc networks. In this case, the end-user stations communicate without an access point and without some features, such as power saving.

Although 802.11 shares many ideas with cellular telephony, it is much simpler and thus much cheaper, because:

- 802.11 cells are 0.01 to 0.10 miles in diameter versus 2 to 10 miles for cellular telephony, so base stations are much cheaper and need much less power. Also, 802.11 end hosts have more power available than cell phones.
- There is no illusion of universal access, so there is no need for an elaborate infrastructure of base station towers that try to blanket urban areas.
- 802.11 is intended for data, so it can tolerate lower quality since it can simply retransmit or correct errors, and packet delays are more acceptable.
- Although both work with mobile end hosts, 802.11 does not need to work at automobile speeds.
- 802.11 leverages the existing Internet for connectivity between base stations versus the telephone system. The former typically relies on best-effort packet switching, while the latter offers connections with a quality of service guarantee for voice communication.

Check Yourself

Which of the following are true?

1. Protocol stacks are an example of using abstraction to hide complexity.
2. TCP/IP is used for WANs, but LANs use a protocol stack appropriate for the lower latency and higher bandwidths.
3. Although the 802.11 LAN standard is wireless like the cell phone, there is little commonality between the two technologies.

Elaboration: In addition to WAN, LAN, and Wireless LAN, there is SAN. It originally stood for *system area network,* but more recently it morphed to *storage area network*. It connects computers and storage devices in a machine room. Fibre Channel Arbitrated Loop (FC-AL) and Infiniband are SANs. Although FC-AL was originally a shared medium, to improve bandwidth there are FC-AL switches. The maximum distance of a SAN link is typically less than 100 meters, and it can connect up to hundreds of nodes.