Chapter 13

Wired LANs: Ethernet

Wired LANs: Ethernet

We learned that a local area network (LAN) is a computer network that is designed for a limited geographic area such as a building or a campus.

Although a LAN can be used as an isolated network to connect computers in an organization for the sole purpose of sharing resources, most LANs today are also linked to a wide area network (WAN) or the Internet.

The LAN market has several technologies such as Ethernet, Token Ring, Token Bus, FDDI, and ATM LAN. Some of these technologies survived for a while, but Ethernet is by far the dominant technology.

13-1 IEEE STANDARDS

In 1985, the Computer Society of the IEEE started a project, called Project 802, to set standards to enable intercommunication among equipment from a variety of manufacturers. Project 802 is a way of specifying functions of the physical layer and the data link layer of major LAN protocols.

Topics discussed in this section:

Data Link Layer
Physical Layer

The relationship of the 802 Standard to the traditional OSI (open source interconnection)model is shown in Figure 13.1. The IEEE has subdivided the data link layer into two sublayers: logical link control (LLC) and media access control (MAC). IEEE has also created several physicallayer standards for different LAN protocols.

Figure 13.1 *IEEE standard for LANs*

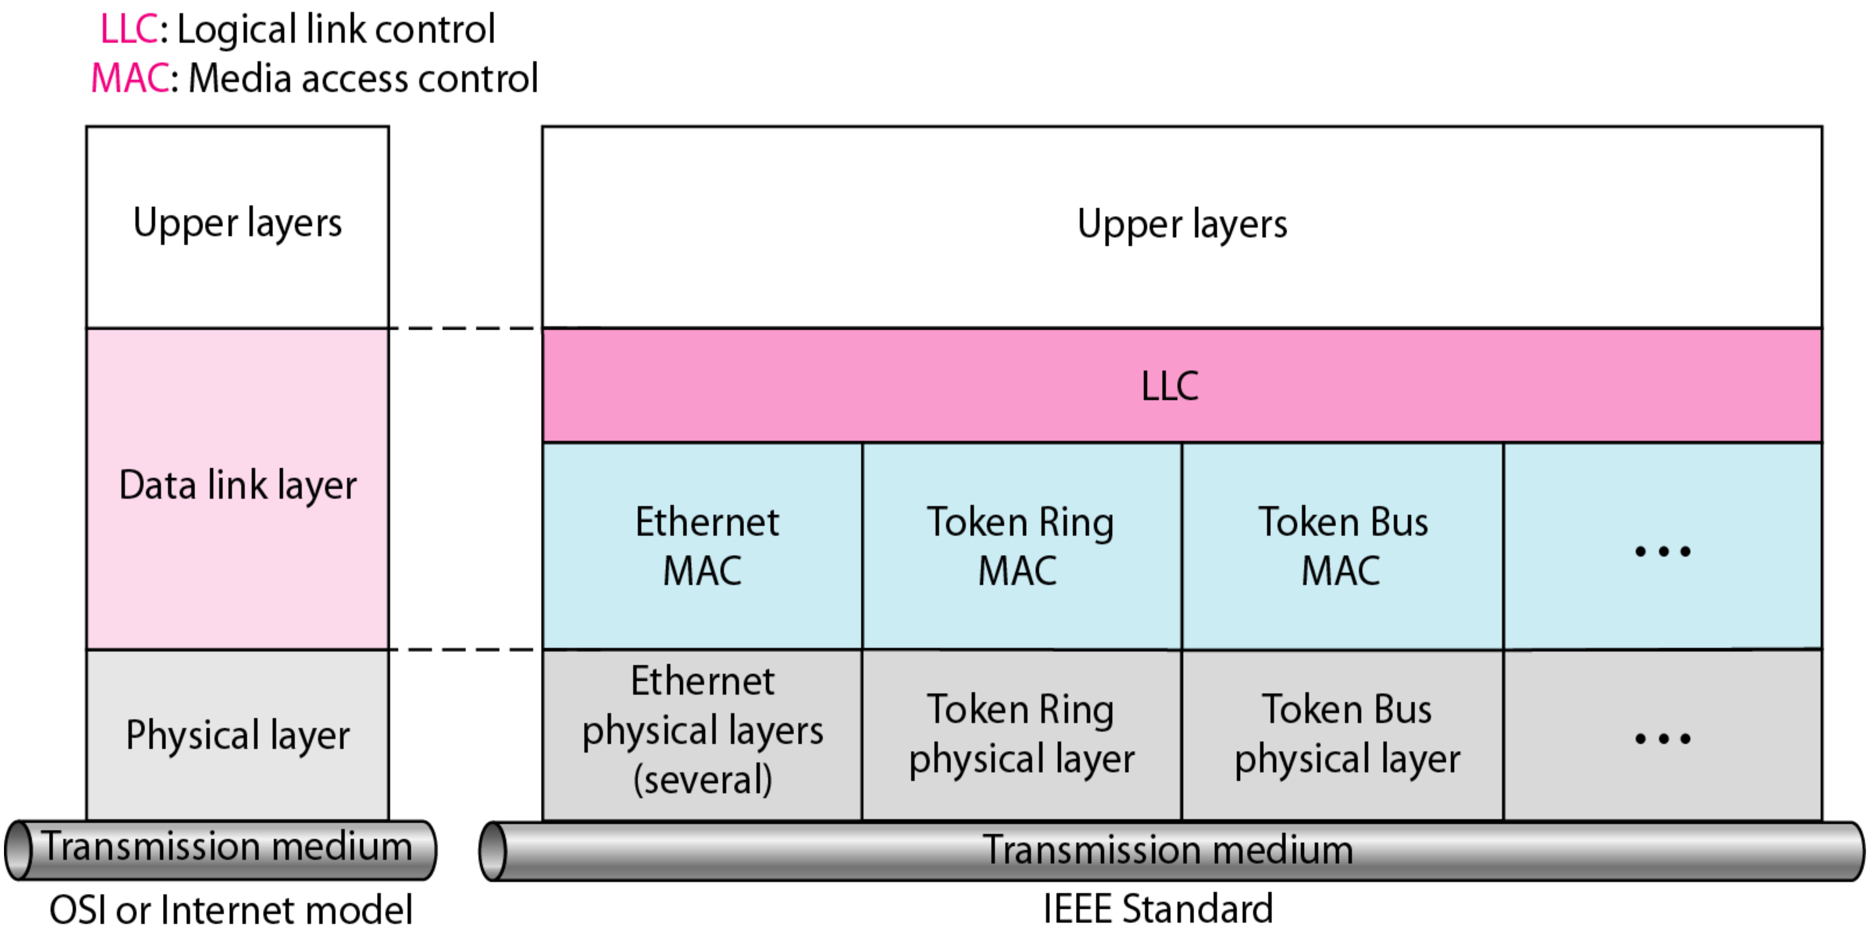

Data Link Layer:

As we mentioned before, the data link layer in the IEEE standard is divided into two sublayers: LLC and MAC.

- *Logical Link Control (LLC)* In Chapter 11, we discussed data link control. We said that data link control handles framing, flow control, and error control. In IEEE Project 802, flow control, error control, and part of the framing duties are collected into one sublayer called the logical link control. Framing is handled in both the LLC sublayer and the MAC sublayer.
- The LLC provides one single data link control protocol for all IEEE LANs. In this way, the LLC is different from the media access control sublayer MAC, which provides different protocols for different LANs

The purpose of the LLC is to provide flow and error control for the upper-layer protocols that actually demand these services. *Media Access Control (MAC)* In Chapter 12, we discussed multiple access methods including random access, controlled

access, and channelization. IEEE Project 802 has created a sublayer called media access control that defines the specific access method for each LAN.

For example, it defines *CSMA/CD as the media access method for Ethernet LANs and the token passing* method for Token Ring and Token Bus LANs. As we discussed in the previous section, part of the framing function is also handled by the MAC layer.

In contrast to the LLC sublayer, the MAC sublayer contains a number of distinct modules; each defines the access method and the framing format specific to the corresponding LAN protocol.

13-2 STANDARD ETHERNET

The original Ethernet was created in 1976 at Xerox's Palo Alto Research Center (PARC). Since then, it has gone through four generations. We briefly discuss the Standard (or traditional) Ethernet in this section.

Topics discussed in this section:

MAC Sublayer
Physical Layer

Figure 13.3 *Ethernet evolution through four generations*

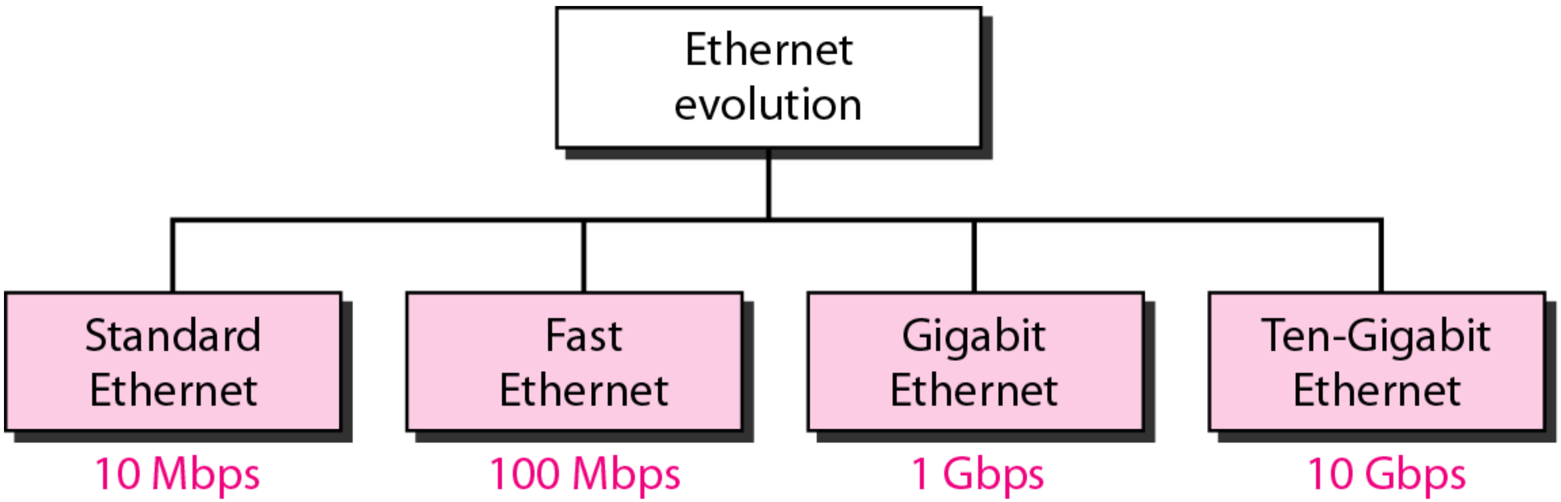

MAC Sublayer

In Standard Ethernet, the MAC sublayer governs the operation of the access method. It also frames data received from the upper layer and passes them to the physical layer.

Figure 13.4 *802.3 MAC frame*

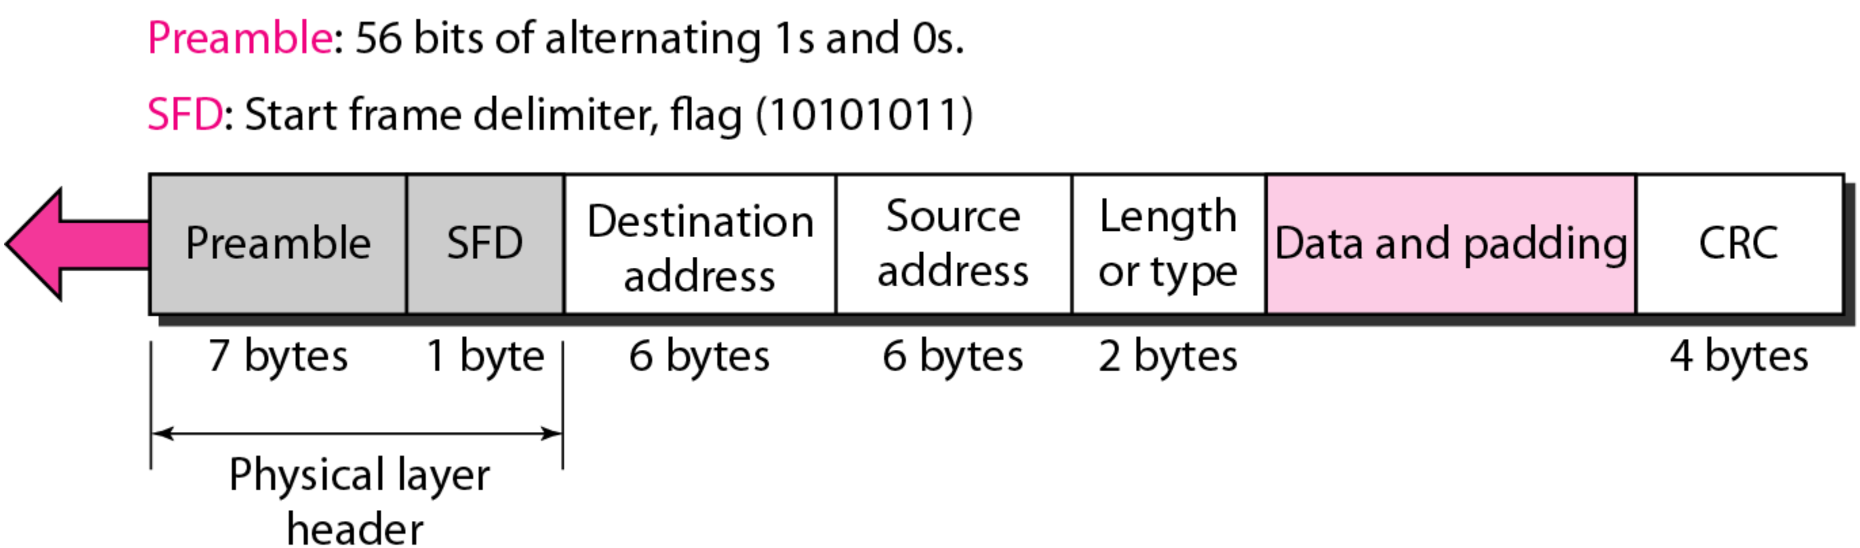

Frame Format

- Preamble. The first field of the 802.3 frame contains 7 bytes (56 bits) of alternating Os and 1s that alerts the receiving system to the coming frame and enables it to synchronize its input timing. The pattern provides only an alert and a timing pulse. The 56-bit pattern allows the stations to miss some bits at the beginning of the frame. The preamble is actually added at the physical layer and is not (formally) part of the frame.

- Start frame delimiter (SFD). The second field (l byte: 10101011) signals the beginning of the frame. The SFD warns the station or stations that this is the last chance for synchronization. The last 2 bits is 11 and alerts the receiver that the next field is the destination address.

- Destination address (DA). The DA field is 6 bytes and contains the physical address of the destination station or stations to receive the packet. We will discuss addressing shortly.
- Source address (SA). The SA field is also 6 bytes and contains the physical address of the sender of the packet. We will discuss addressing shortly.
- Length or type. This field is defined as a type field or length field. The original Ethernet used this field as the type field to define the upper-layer protocol using the MAC frame. The IEEE standard used it as the length field to define the number of bytes in the data field. Both uses are common today.
- Data. This field carries data encapsulated from the upper-layer protocols. It is a minimum of 46 and a maximum of 1500 bytes, as we will see later.
- CRC. The last field contains error detection information, in this case a CRC-32

Figure 13.5 *Minimum and maximum lengths*

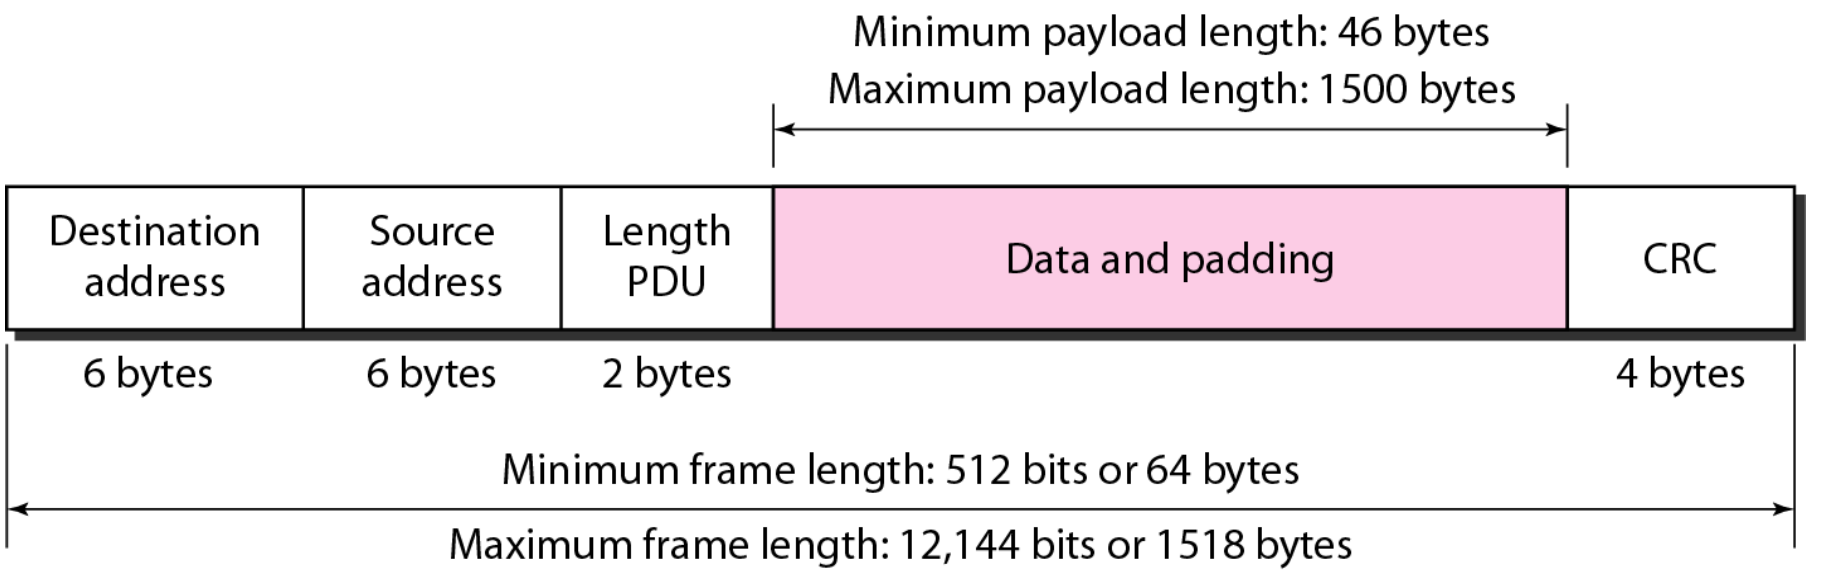

Note

Frame length:

Minimum: 64 bytes (512 bits)

Maximum: 1518 bytes (12,144 bits)

Addressing

- Each station on an Ethernet network (such as a PC, workstation, or printer) has its own network interface card (NIC). The NIC fits inside the station and provides the station with a 6-byte physical address. As shown in Figure 13.6, the Ethernet address is 6 bytes (48 bits), nonnally written in hexadecimal notation, with a colon between the bytes.
- Unicast, Multicast, and Broadcast Addresses A source address is always a unicast address-the frame comes from only one station. The destination address, however, can be unicast, multicast, or broadcast. Figure 13.7 shows how to distinguish a unicast address from a multicast address. If the least significant bit of the first byte in a destination address is 0, the address is unicast; otherwise, it is multicast.

Figure 13.6 ***Example of an Ethernet address in hexadecimal notation***

06 : 01 : 02 : 01 : 2C : 4B

6 bytes = 12 hex digits = 48 bits

Figure 13.7 *Unicast and multicast addresses*

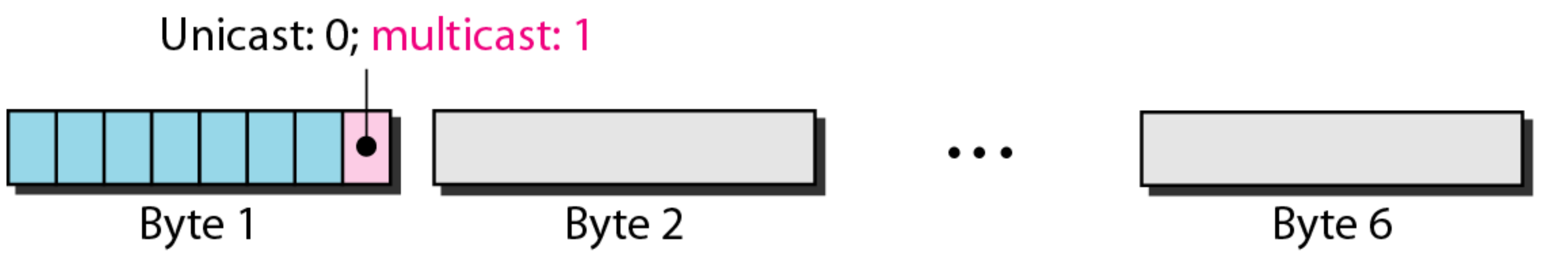

Note

The least significant bit of the first byte defines the type of address.
If the bit is 0, the address is unicast; otherwise, it is multicast.

Note

The broadcast destination address is a special case of the multicast address in which all bits are 1s.

Example 13.1

Define the type of the following destination addresses:

a. 4A:30:10:21:10:1A ***b. 47:20:1B:2E:08:EE***

c. FF:FF:FF:FF:FF:FF

Solution

To find the type of the address, we need to look at the second hexadecimal digit from the left. If it is even, the address is unicast. If it is odd, the address is multicast. If all digits are F's, the address is broadcast. Therefore, we have the following:

a. This is a unicast address because A in binary is 1010.

b. This is a multicast address because 7 in binary is 0111.

c. This is a broadcast address because all digits are F's.

Example 13.2

Show how the address 47:20:1B:2E:08:EE is sent out on line.

Solution

The address is sent left-to-right, byte by byte; for each byte, it is sent right-to-left, bit by bit, as shown below:

← 11100010 00000100 11011000 01110100 00010000 01110111

Physical Layer

- The Standard Ethernet defines several physical layer implementations; four of the most common, are shown in Figure 13.8.

Figure 13.8 *Categories of Standard Ethernet*

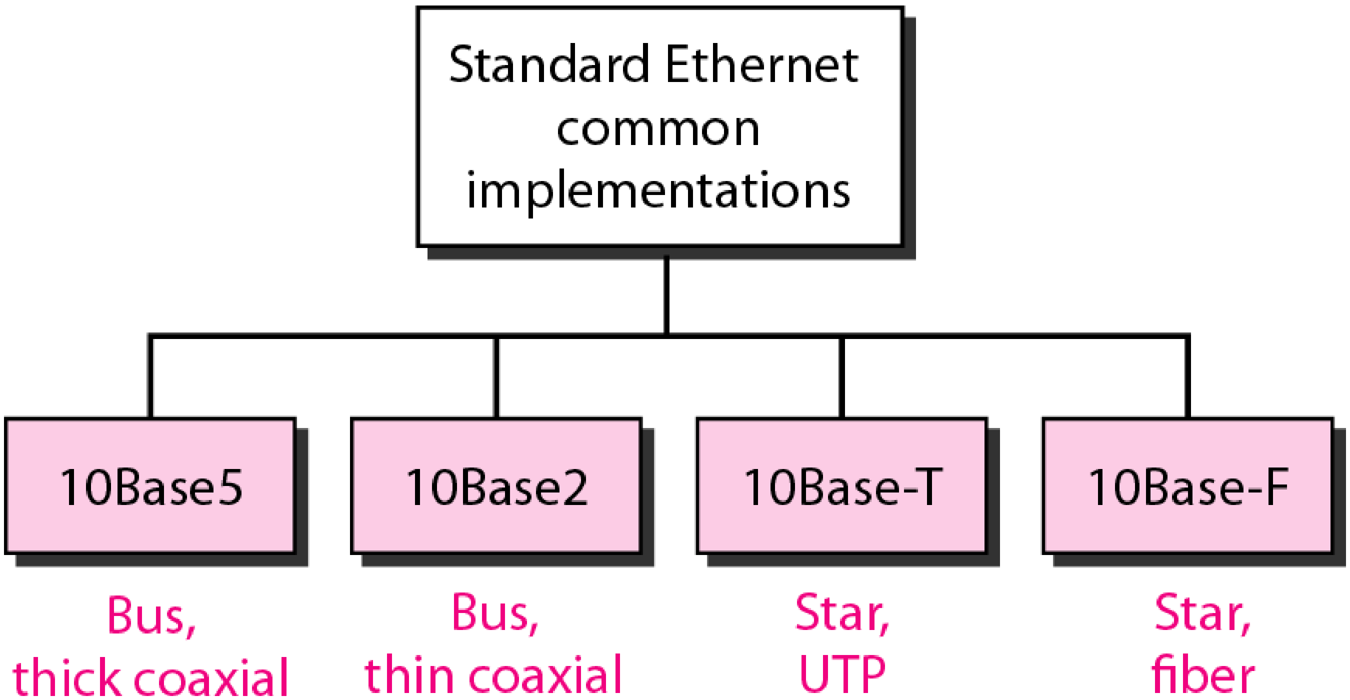

- *lOBase5: Thick Ethernet*

The first implementation is called **10Base5, thick Ethernet, or Thicknet. The nickname** derives from the size of the cable. lOBase5 was the first Ethernet specification to use a bus topology with an external **transceiver (transmitter/receiver) connected via a** tap to a thick coaxial cable. Figure 13.10 shows a schematic diagram of a lOBase5 implementation.

Figure 13.10 ***10Base5 implementation***

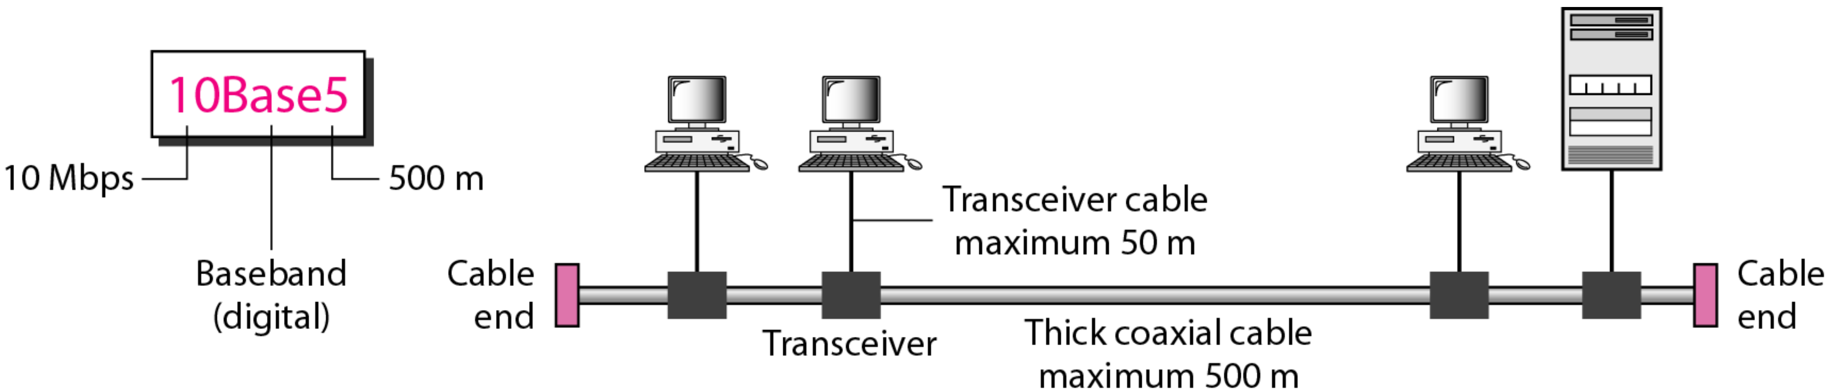

- *10Base2: Thin Ethernet*
- The second implementation is called lOBase2, **thin Ethernet, or Cheapernet. 1OBase2** also uses a bus topology, but the cable is much thinner and more flexible. The cable can be bent to pass very close to the stations. In this case, the transceiver is normally part of the network interface card (NIC), which is installed inside the station. Figure 13.11 shows the schematic diagram of a IOBase2 implementation.
- Note This implementation is more cost effective than 10Base5 because thin coaxial cable is less expensive than thick coaxial and the tee connections are much cheaper than taps. Installation is simpler because the thin coaxial cable is very flexible. However, the length of each segment cannot exceed *185 m (close to 200 m) due to the high level of thin coaxial* cable.

Figure 13.11 *10Base2 implementation*

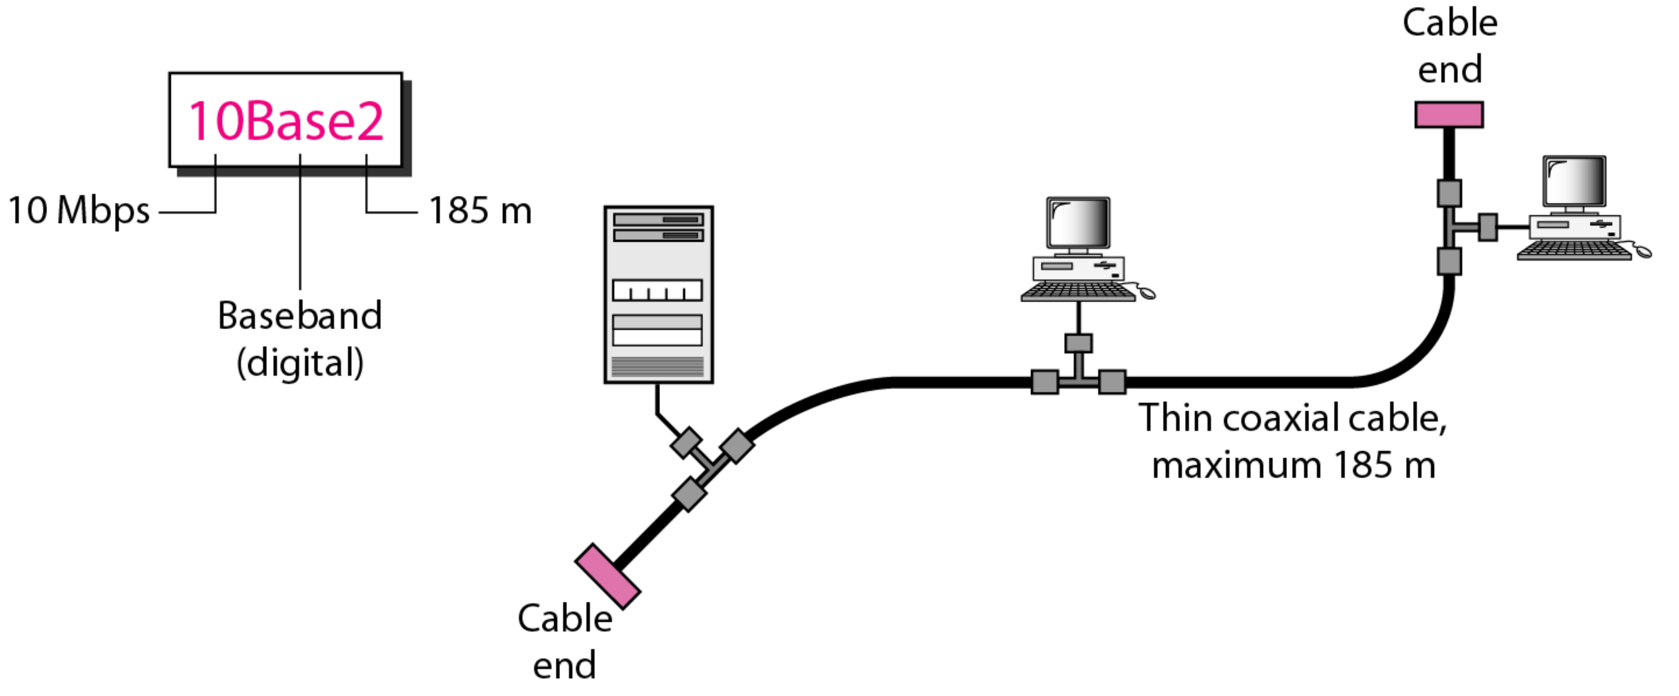

lOBase-T: Twisted-Pair Ethernet

- The third implementation is called lOBase-T or twisted-pair Ethernet. 1OBase-T uses a physical star topology. The stations are connected to a hub via two pairs of twisted cable, as shown in Figure 13.12.
- Note that two pairs of twisted cable create two paths (one for sending and one for receiving) between the station and the hub. Any collision here happens in the hub.
- The maximum length of the twisted cable here is defined as 100 m, to minimize the effect of attenuation in the twisted cable.

Figure 13.12 *10Base-T implementation*

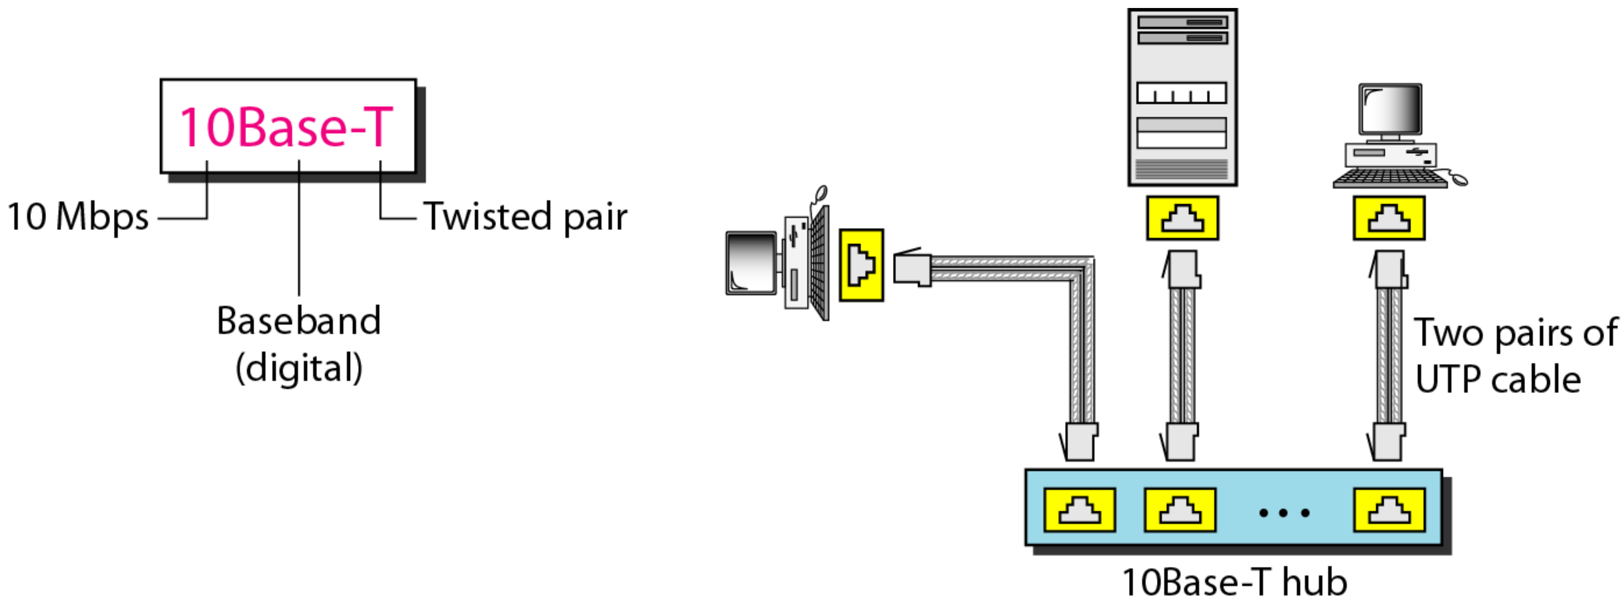

- *lOBase-F: Fiber Ethernet*

- Although there are several types of optical fiber lO Mbps Ethernet, the most common is called 10Base-F. lOBase-F uses a star topology to connect stations to a hub. The stations are connected to the hub using two fiber-optic cables, as shown in Figure 13.13.

Figure 13.13 ***10Base-F implementation***

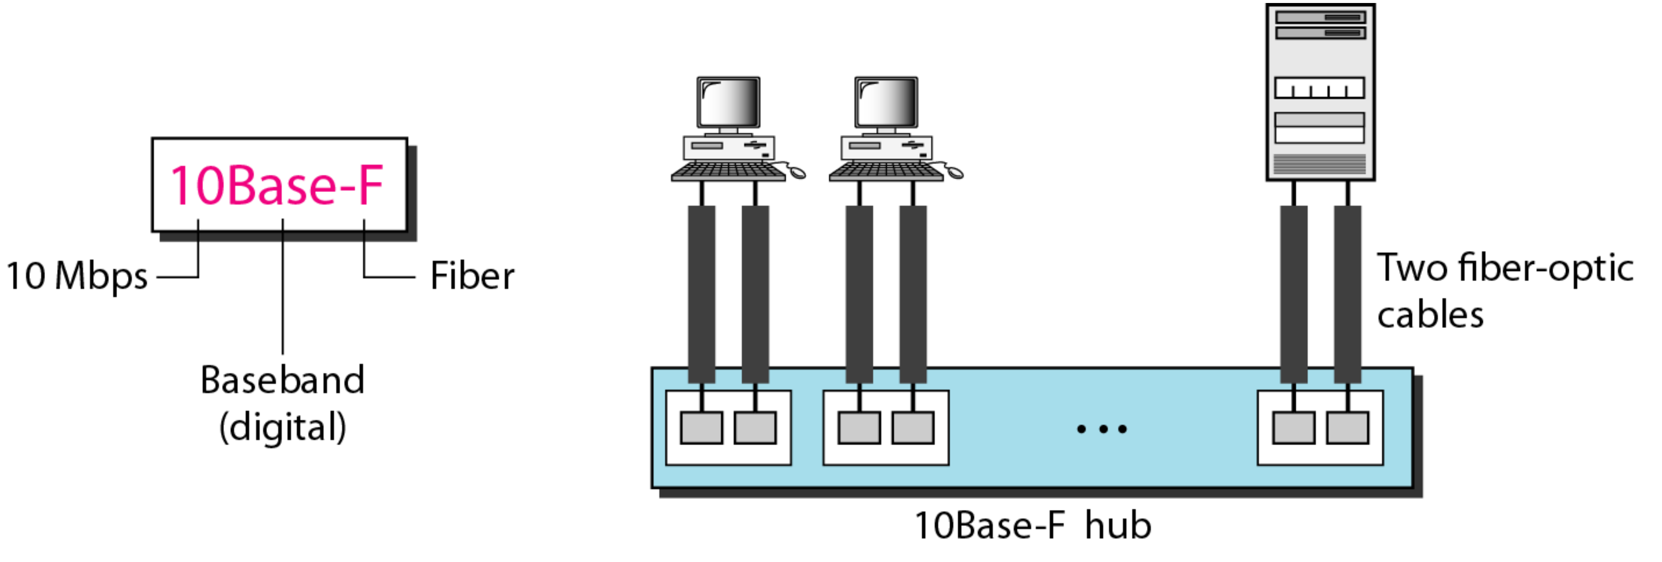

Table 13.1 *Summary of Standard Ethernet implementations*

Characteristics	*10Base5*	*10Base2*	*10Base-T*	*10Base-F*
Media	Thick coaxial cable	Thin coaxial cable	2 UTP	2 Fiber
Maximum length	500 m	185 m	100 m	2000 m
Line encoding	Manchester	Manchester	Manchester	Manchester

13-3 CHANGES IN THE STANDARD

The 10-Mbps Standard Ethernet has gone through several changes before moving to the higher data rates. These changes actually opened the road to the evolution of the Ethernet to become compatible with other high-data-rate LANs.

Topics discussed in this section:

Bridged Ethernet
Switched Ethernet
Full-Duplex Ethernet

- Bridged Ethernet
- The first step in the Ethernet evolution was the division of a LAN by bridges.
- Bridges have two effects on an Ethernet LAN: They raise the bandwidth and they separate collision domains. In an unbridged Ethernet network, the total capacity (10 Mbps) is shared among all stations with a frame to send; the stations share the bandwidth of the network.
- If only one station has frames to send, it benefits from the total capacity (10 Mbps). But if more than one station needs to use the network, the capacity is shared. We can say that, in this case, each station on average, sends at a rate of 5 Mbps.

- A bridge divides the network into two or more networks.

  Bandwidth-wise, each network is independent. For

- example, in Figure 13.15, a network with 12 stations is divided

  into two networks, each with 6 stations.

- Now each network has a capacity of 10 Mbps. The lO-Mbps capacity in each segment is now shared between 6 stations (actually 7 because the bridge acts as a station in each segment), not 12 stations.

Figure 13.15 ***A network with and without a bridge***

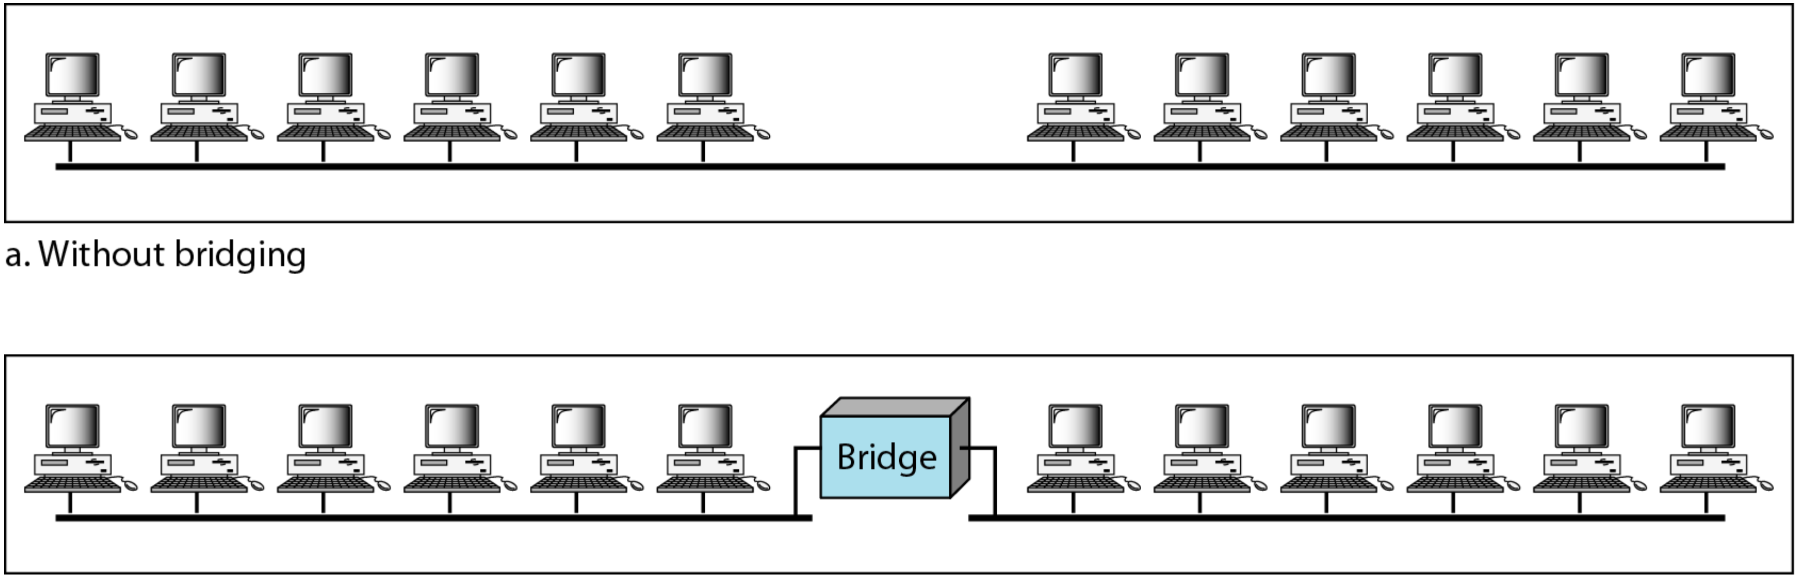

a. Without bridging

b. With bridging

- Another advantage of a bridge is the separation of the collision domain. Figure 13.16 shows the collision domains for an unbridged and a bridged network. You can see that the collision domain becomes much smaller and the probability of collision is reduced .

Figure 13.16 *Collision domains in an unbridged network and a bridged network*

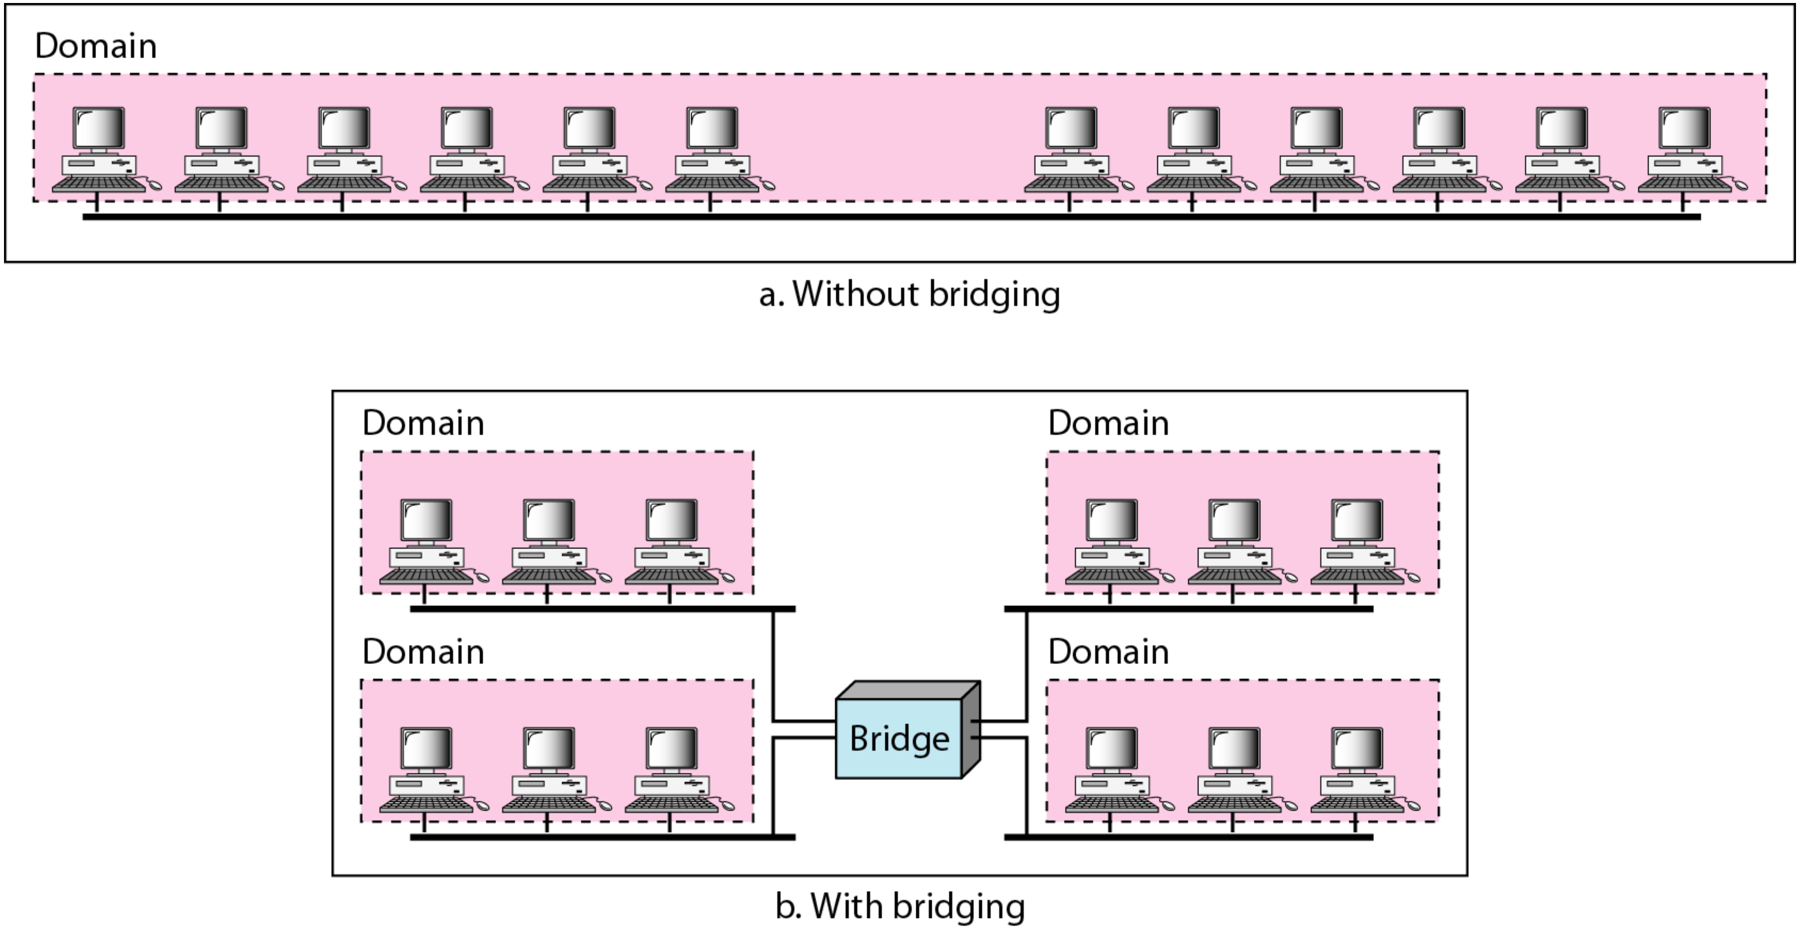

- Switched Ethernet
- The idea of a bridged LAN can be extended to a switched LAN. Instead of having two to four networks, why not have *N networks, where N is the number of stations on the* LAN? In other words, if we can have a multiple-port bridge, why not have an *N-port* switch? In this way, the bandwidth is shared only between the station and the switch (5 Mbps each). In addition, the collision domain is divided into *N domains.*

Figure 13.17 ***Switched Ethernet***

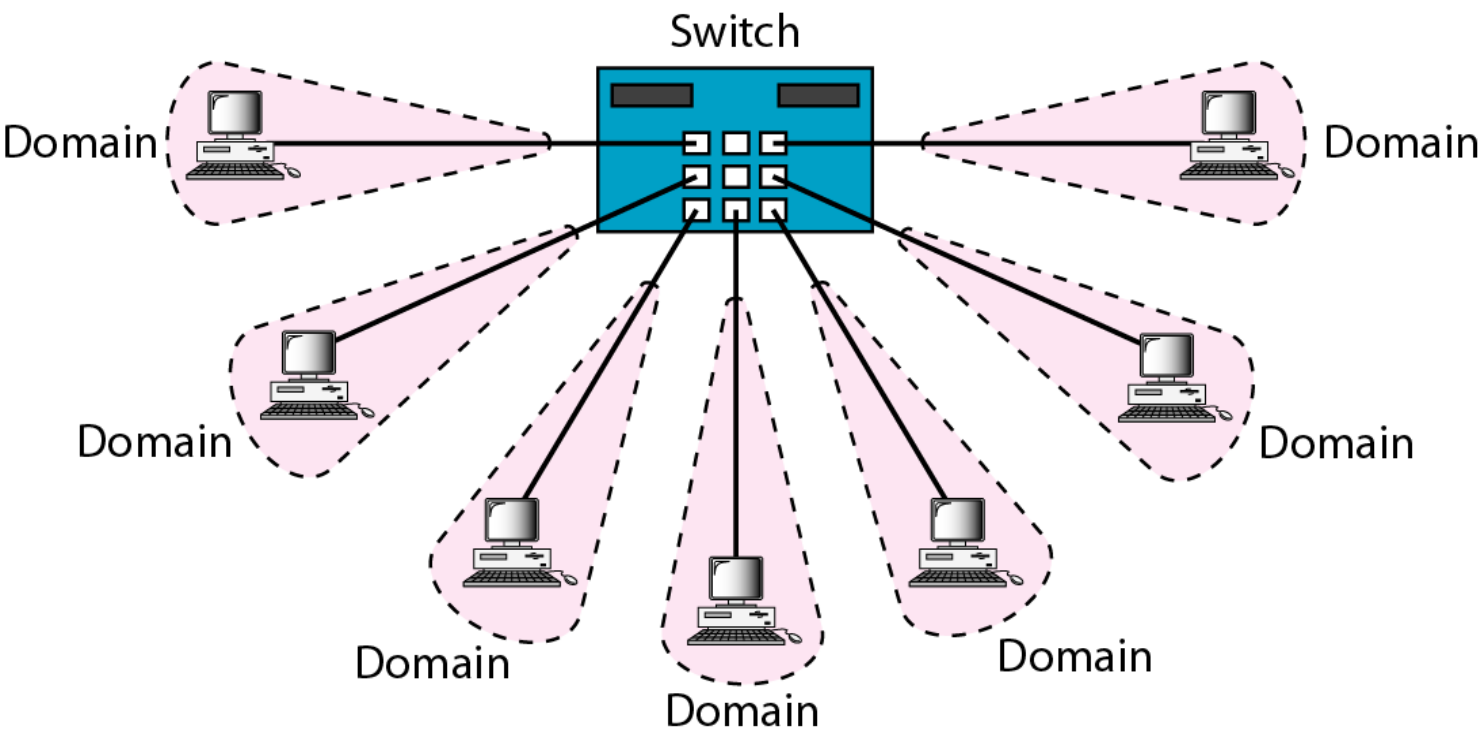

- Full-Duplex Ethernet
- One of the limitations of 10Base5 and lOBase2 is that communication is half-duplex (lOBase-T is always full-duplex); a station can either send or receive, but may not do both at the same time.
- The next step in the evolution was to move from switched Ethernet to full-duplex switched Ethernet. The full-duplex mode increases the capacity of each domain from 10 to 20 Mbps. Figure 13.18 shows a switched Ethernet in full-duplex mode.
- Note that instead of using one link between the station and the switch, the configuration uses two links: one to transmit and one to receive.

Figure 13.18 *Full-duplex switched Ethernet*

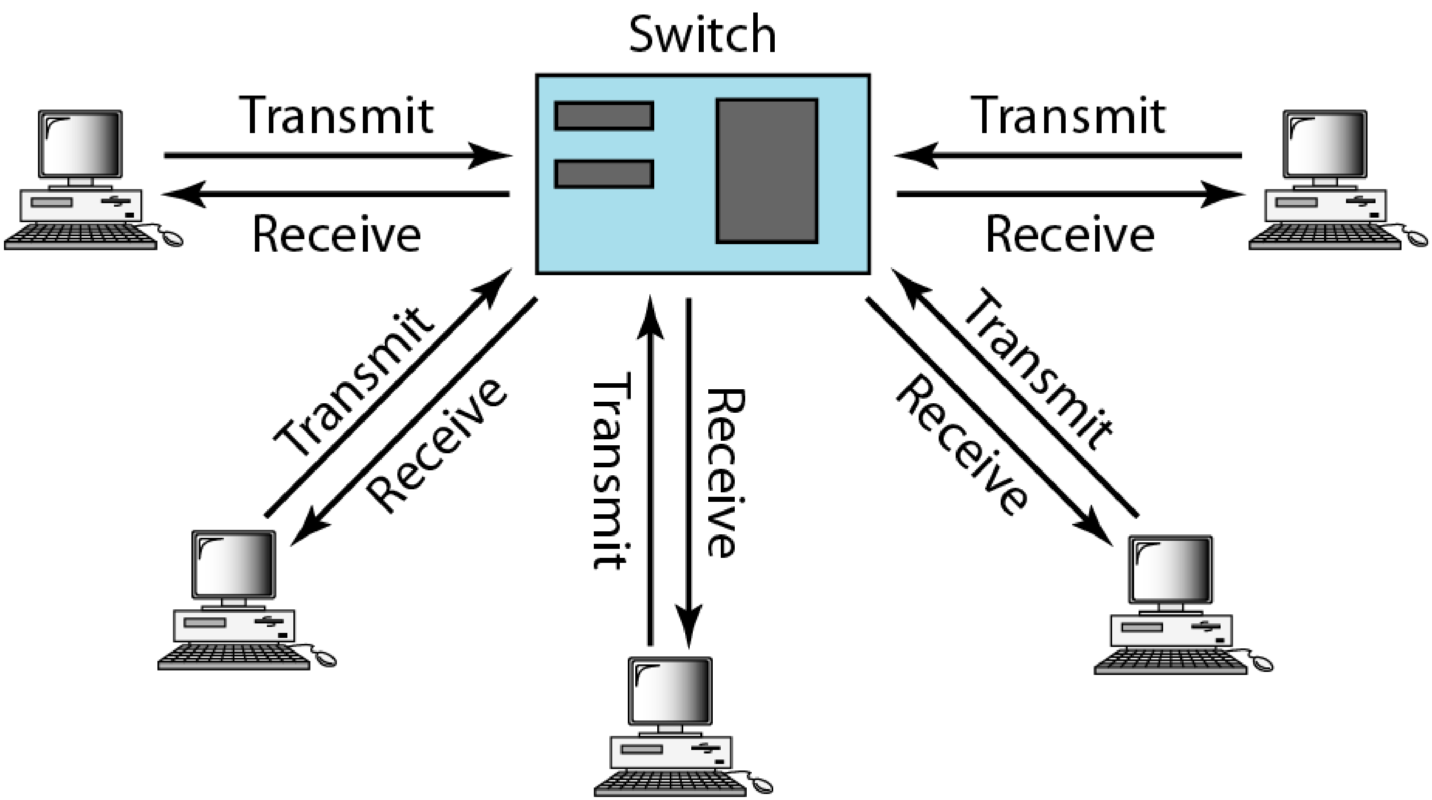

13-4 FAST ETHERNET

Fast Ethernet was designed to compete with LAN protocols such as FDDI or Fiber Channel. IEEE created Fast Ethernet under the name 802.3u. Fast Ethernet is backward-compatible with Standard Ethernet, but it can transmit data 10 times faster at a rate of 100 Mbps.

Topics discussed in this section:

MAC Sublayer
Physical Layer

Figure 13.19 *Fast Ethernet topology*

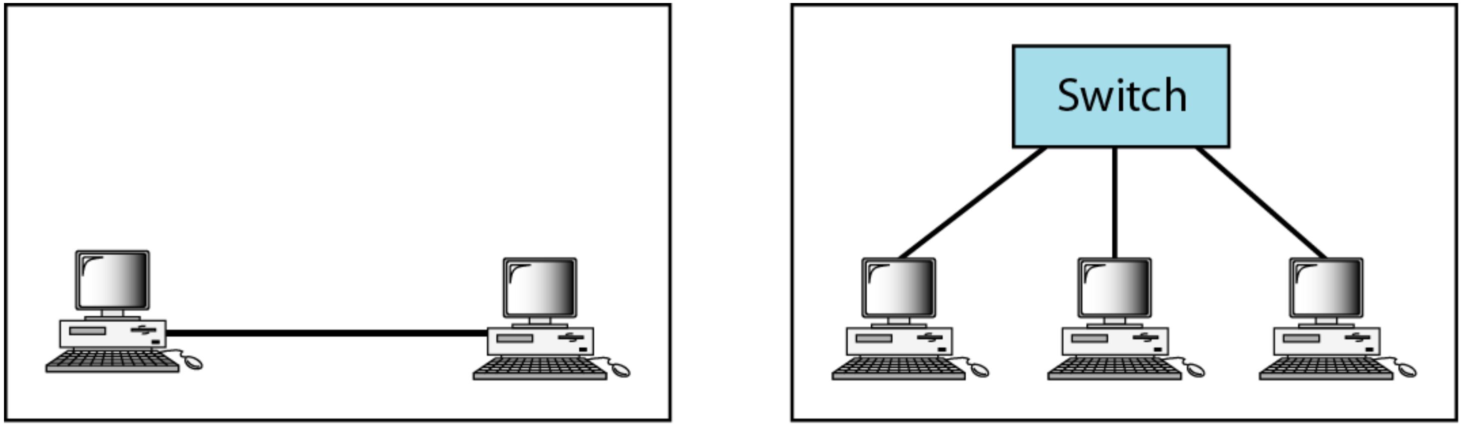

a. Point-to-point

b. Star

Figure 13.20 *Fast Ethernet implementations*

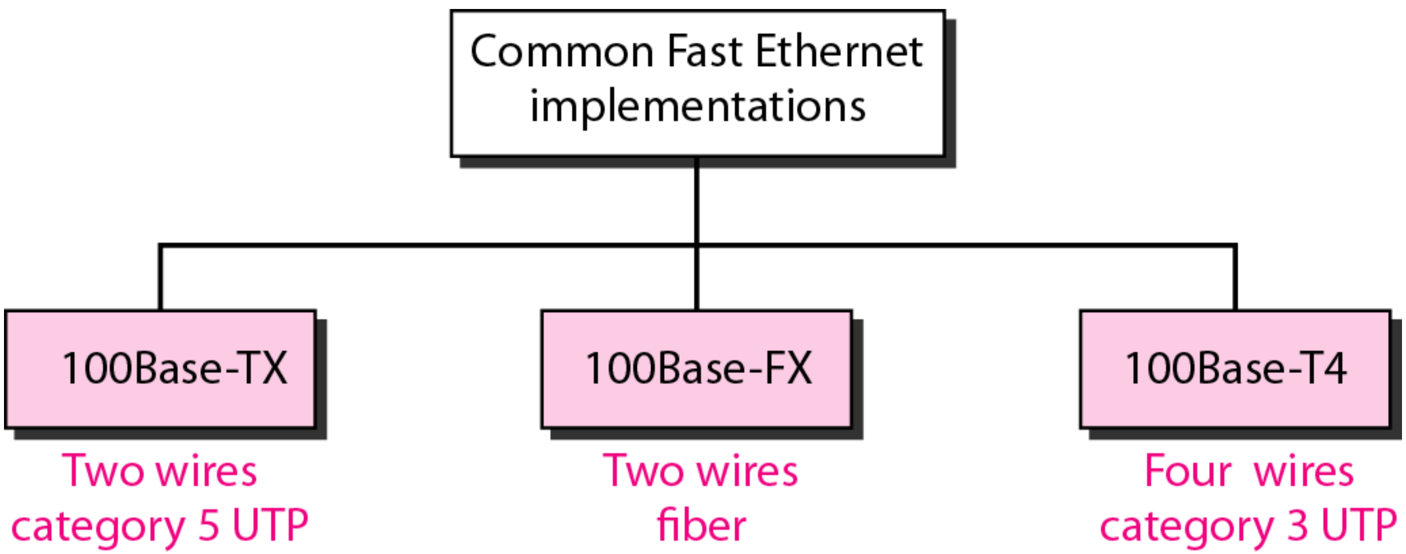

Figure 13.21 *Encoding for Fast Ethernet implementation*

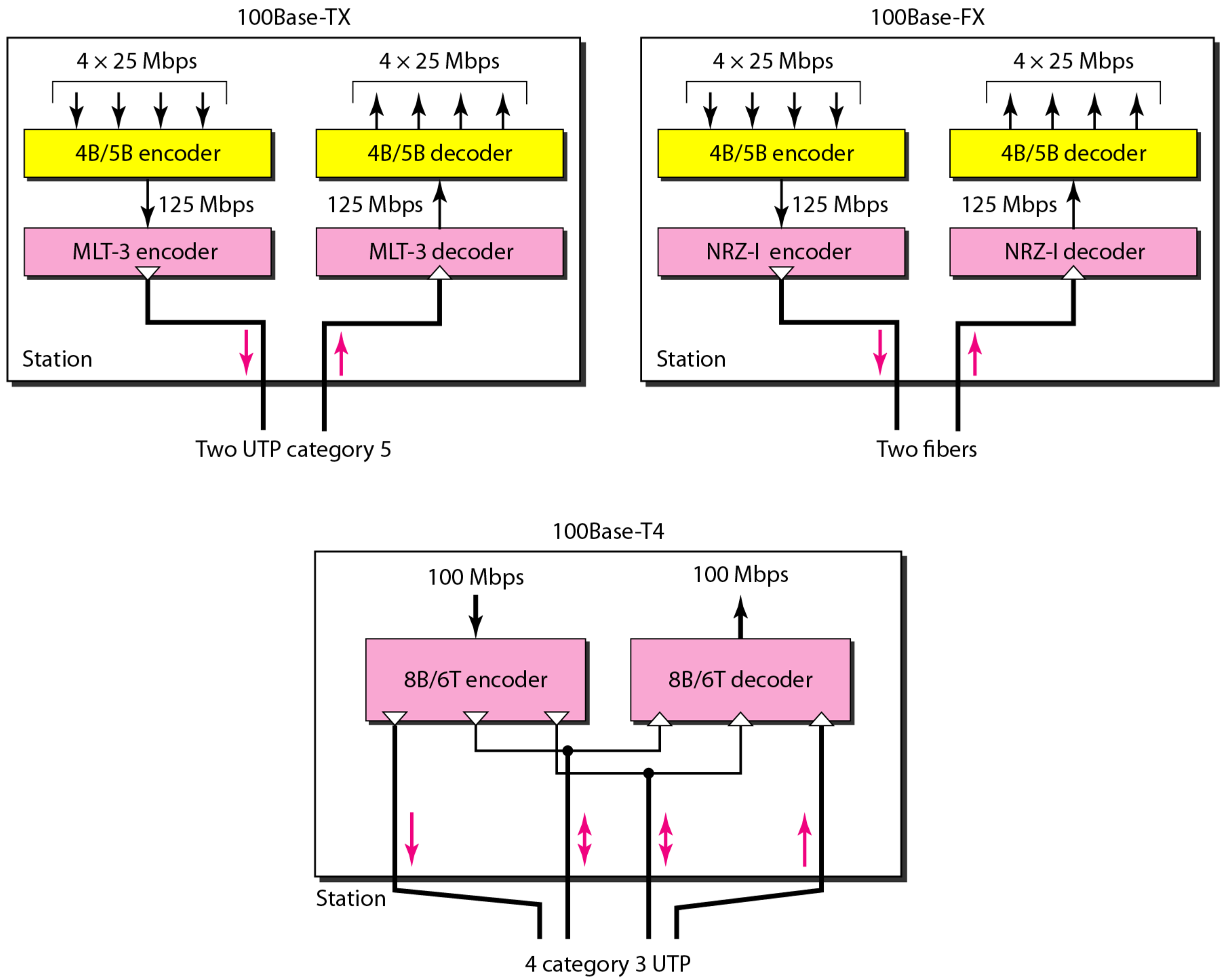

Table 13.2 ***Summary of Fast Ethernet implementations***

Characteristics	*100Base-TX*	*100Base-FX*	*100Base-T4*
Media	Cat 5 UTP or STP	Fiber	Cat 4 UTP
Number of wires	2	2	4
Maximum length	100 m	100 m	100 m
Block encoding	4B/5B	4B/5B	
Line encoding	MLT-3	NRZ-I	8B/6T

13-5 GIGABIT ETHERNET

The need for an even higher data rate resulted in the design of the Gigabit Ethernet protocol (1000 Mbps). The IEEE committee calls the standard 802.3z.

Topics discussed in this section:

MAC Sublayer
Physical Layer
Ten-Gigabit Ethernet

Note

**In the full-duplex mode of Gigabit Ethernet, there is no collision;
the maximum length of the cable is determined by the signal attenuation in the cable.**

Figure 13.22 *Topologies of Gigabit Ethernet*

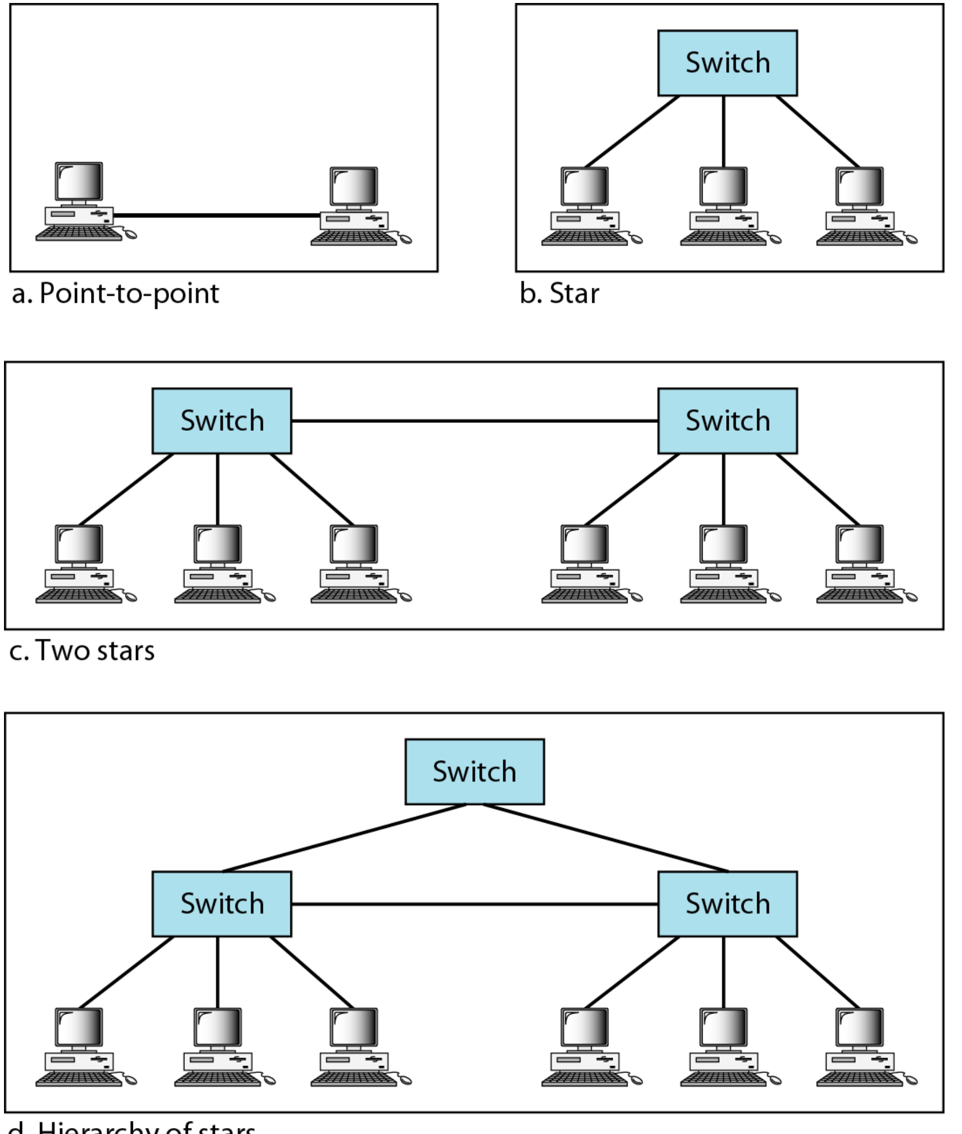

a. Point-to-point
b. Star
c. Two stars
d. Hierarchy of stars

Figure 13.23 *Gigabit Ethernet implementations*

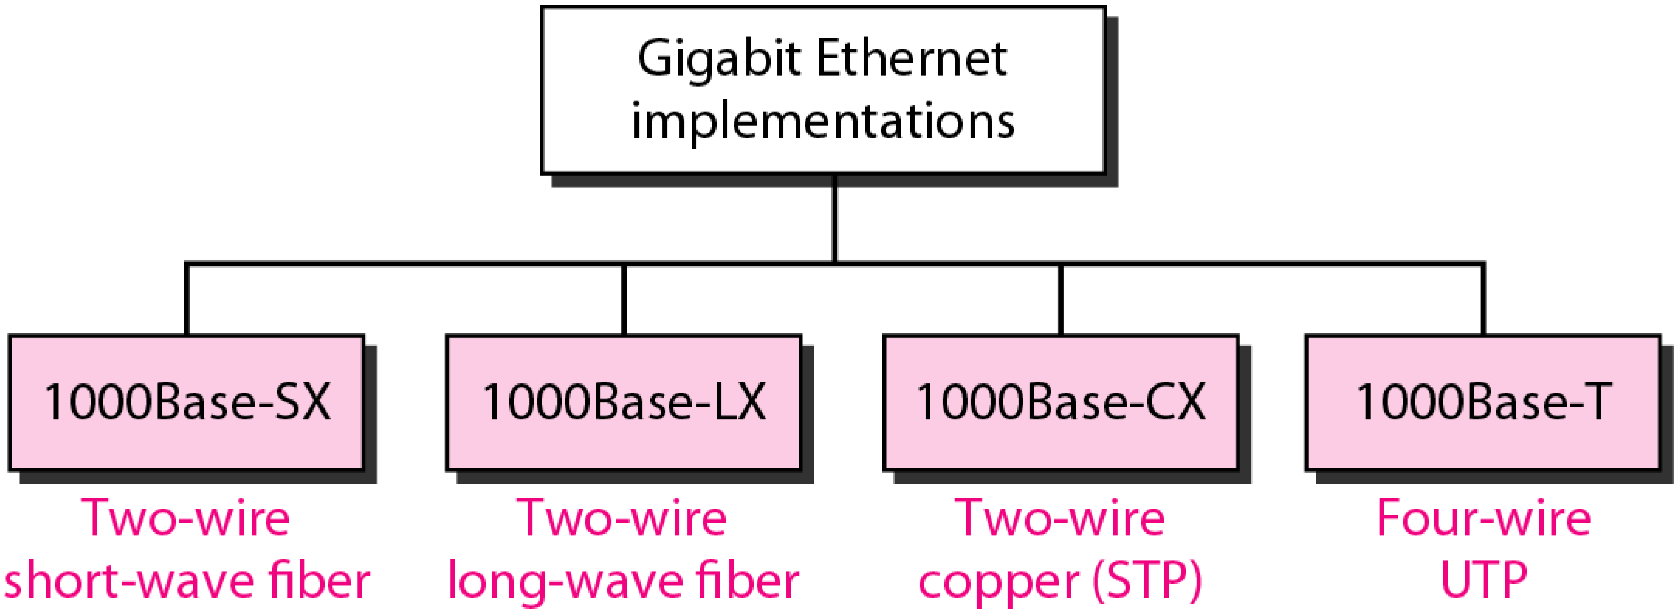

Figure 13.24 ***Encoding in Gigabit Ethernet implementations***

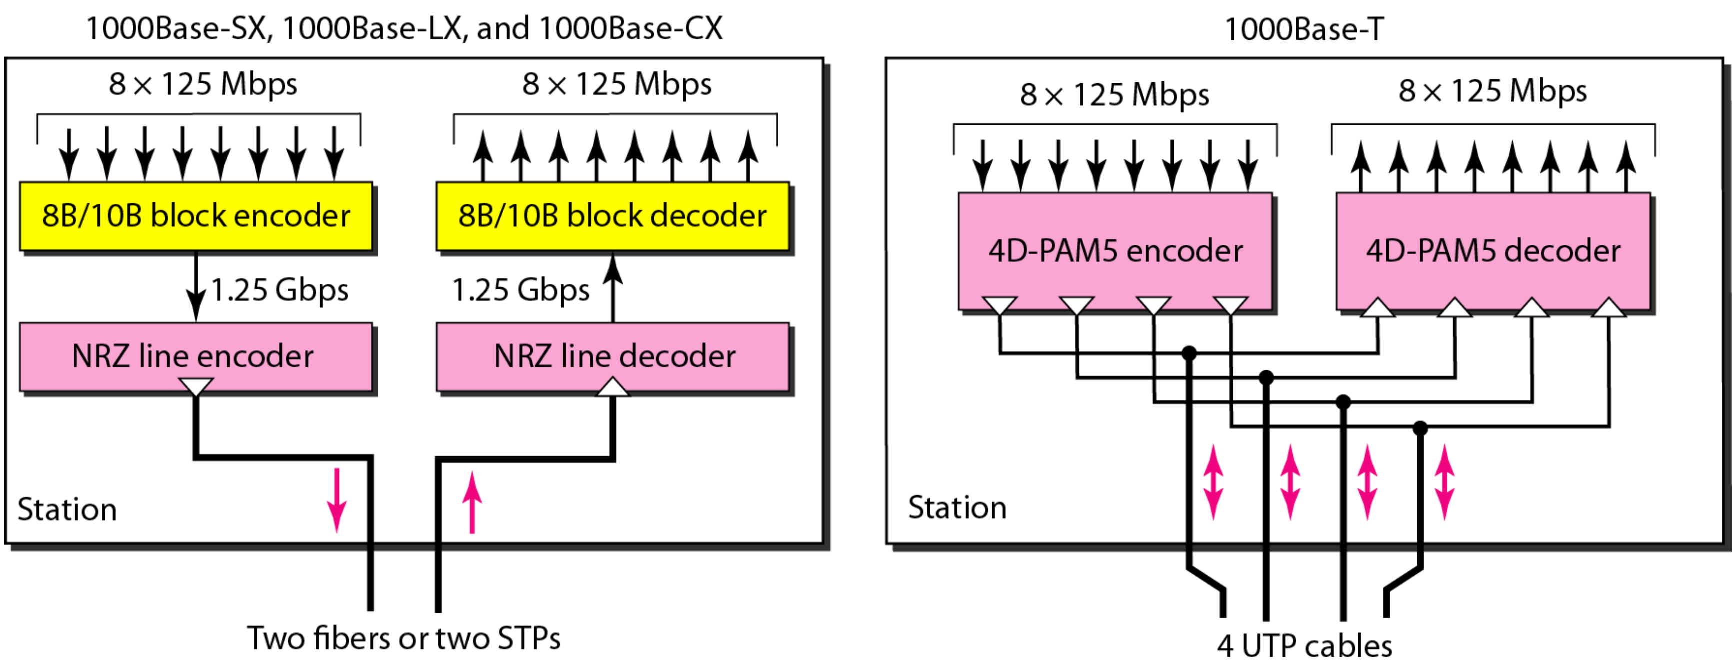

Table 13.3 *Summary of Gigabit Ethernet implementations*

Characteristics	*1000Base-SX*	*1000Base-LX*	*1000Base-CX*	*1000Base-T*
Media	Fiber short-wave	Fiber long-wave	STP	Cat 5 UTP
Number of wires	2	2	2	4
Maximum length	550 m	5000 m	25 m	100 m
Block encoding	8B/10B	8B/10B	8B/10B	
Line encoding	NRZ	NRZ	NRZ	4D-PAM5

Table 13.4 *Summary of Ten-Gigabit Ethernet implementations*

Characteristics	*10GBase-S*	*10GBase-L*	*10GBase-E*
Media	Short-wave 850-nm multimode	Long-wave 1310-nm single mode	Extended 1550-mm single mode
Maximum length	300 m	10 km	40 km