CS-204: COMPUTER NETWORKS

Lecture 3

Chapter 13: Wired LANs-Ethernet

Instructor: Dr. Vandana Kushwaha

1. INTRODUCTION

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 seen 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.

2. 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 does not seek to replace any part of the OSI or the Internet model. Instead, it is a way of specifying functions of the physical layer and the data link layer of major LAN protocols.

The original Ethernet was created in 1976 at Xerox's Palo Alto Research Center (PARC). Since then, it has gone through four generations:

- a. Standard Ethernet (10 Mbps),
- b. Fast Ethernet (100 Mbps),
- c. Gigabit Ethernet (1 Gbps), and
- **d.** Ten-Gigabit Ethernet (10 Gbps), as shown in Figure 13.1.



Figure 13.1 Ethernet evolution through four generations

2.1. Standard Ethernet(IEEE 802.3)

Standard Ethernet also known as IEEE 802.3 was the LAN standard proposed by IEEE. Data rate for standard Ethernet is 10 Mbps.

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.

• Frame Format

The Ethernet frame contains seven fields: preamble, SFD, DA, SA, length or type of protocol data unit (PDU), upper-layer data, and the CRC. Ethernet does not provide any mechanism for acknowledging received frames, making it what is known as an unreliable medium. Acknowledgments must be implemented at the higher layers. The format of the MAC frame is shown in Figure 13.2.

Preamble: 56 bits of alternating 1s and 0s.

SFD: Start frame delimiter, flag (10101011)



Figure 13.2 802.3 MAC frame

- i. **Preamble.** The first field of the 802.3 frame contains 7 bytes (56 bits) of alternating 0s 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.
- ii. **Start frame delimiter (SFD).** The second field (1 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.
- iii. **Destination address (DA).** The DA field is 6 bytes and contains the physical address of the destination station or stations to receive the packet.
- iv. **Source address (SA).** The SA field is also 6 bytes and contains the physical address of the sender of the packet.
- v. **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.
- vi. **Data.** This field carries data encapsulated from the upper-layer protocols. It is a minimum of 46 and a maximum of 1500 bytes.
- vii. **CRC.** The last field contains error detection information, in this case a CRC-32.
- viii. **Frame Length** Ethernet has imposed restrictions on both the minimum and maximum lengths of a frame, as shown in Figure 13.3.



Figure 13.3 Minimum and maximum lengths

The minimum length restriction is required for the correct operation of *CSMA/CD* as we will see shortly. An Ethernet frame needs to have a minimum length of 512 bits or 64 bytes. Part of this length is the header and the trailer. If we count 18 bytes of header and trailer (6 bytes of source address, 6 bytes of destination address, 2 bytes of length or type, and 4 bytes of CRC), then the minimum length of data from the upper layer is 64 - 18 = 46 bytes. If the upper-layer packet is less than 46 bytes, padding is added to make up the difference.

The standard defines the maximum length of a frame (without preamble and SFD field) as 1518 bytes. If we subtract the 18 bytes of header and trailer, the maximum length of the payload is 1500 bytes. The maximum length restriction has two historical reasons. First, memory was very expensive when Ethernet was designed: a maximum length restriction helped to reduce the size of the buffer. Second, the maximum length restriction prevents one station from monopolizing the shared medium, blocking other stations that have data to send.

2.2. MAC 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 (MAC) address. As shown in Figure 13.4, the Ethernet address is 6 bytes (48 bits), normally written in hexadecimal notation, with a colon between the bytes.





Unicast, Multicast, and Broadcast Addresses

Data is transmitted over a network by three simple methods i.e. Unicast, Broadcast, and Multicast Figure 13.5. So let's begin to summarize the difference between these three:

- Unicast: from one source to one destination i.e. One-to-One
- Broadcast: from one source to all possible destinations i.e. One-to-All

 Multicast: from one source to multiple destinations stating an interest in receiving the traffic i.e. One-to-Many



Figure 13.5 Unicasting, Multicasting and Broadcasting

- A source address is always a unicast address as the frame comes from only one station.
- The destination address, however, can be unicast, multicast, or broadcast.
- Figure 13.6 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 Unicast and multicast MAC addresses

- A unicast destination address defines only one recipient; the relationship between the sender and the receiver is one-to-one.
- A multicast destination address defines a group of addresses; the relationship between the sender and the receivers is one-to-many.
- The broadcast address is a special case of the multicast address; the recipients are all the stations on the LAN. A broadcast destination address is forty-eight 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 (even).
- b. This is a multicast address because 7 in binary is 0111 (odd).
- c. This is a broadcast address because all digits are F's.

The way the addresses are sent out on line is different from the way they are written in hexadecimal notation. The transmission is left-to-right, byte by byte; however, for each byte, the least significant bit is sent first and the most significant bit is sent last. This means that the bit that defines an address as unicast or multicast arrives first at the receiver.

2.3. Categories of Standard Ethernet

The Standard Ethernet defines several physical layer implementations; four of the most common, are shown in Figure 13.7.



Figure 13.7 Categories of Standard Ethernet

Encoding and Decoding

All standard implementations use digital signaling (baseband) at 10 Mbps. At the sender, data are converted to a digital signal using the Manchester scheme; at the receiver, the received signal is interpreted as Manchester and decoded into data.

• 10Base5: Thick Ethernet

The first implementation is called **10Base5**, **thick Ethernet**, **or Thicknet**. The nickname derives from the size of the cable, which is roughly the size of a garden hose and too stiff to bend with your hands. 10Base5 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.



The transceiver is responsible for transmitting, receiving, and detecting collisions. The transceiver is connected to the station via a transceiver cable that provides separate paths for sending and receiving. This means that collision can only happen in the coaxial cable. The maximum length of the coaxial cable must not exceed *500* m, otherwise, there is excessive degradation of the signal. If a length of more than *500* m is needed, up to five segments, each a maximum of 500-meter, can be connected using repeaters.

• 10Base2: Thin Ethernet

The second implementation is called 10Base2, **thin** Ethernet, or Cheaper net. 10Base2 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.

Note that the collision here occurs in the thin coaxial cable. 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 attenuation in thin coaxial cable.

• 10Base-T: Twisted-Pair Ethernet

The third implementation is called 10Base-T or twisted-pair Ethernet. 10Base-T uses a physical star topology. The stations are connected to a hub via two pairs of twisted cable. 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. Compared to 10Base5 or 10Base2, we can see that the hub actually replaces the coaxial cable as far as a collision is concerned. The maximum length of the twisted cable here is defined as 100 m, to minimize the effect of attenuation in the twisted cable.

• 10Base-F: Fiber Ethernet

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

Summary

Table 13.1 shows a summary of Standard Ethernet implementations.

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

Table 13.1 Summary of Standard Ethernet implementations

2.4. 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. We discuss some of these changes in this section.

2.4.1. Bridged Ethernet

The first step in the Ethernet evolution was the division of a LAN by bridges. A Bridge is a two port switch used to connect two segments of a LAN. Bridges have two effects on an Ethernet LAN:

- They raise the bandwidth and
- They separate collision domains.

Raising the Bandwidth

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. For example, if two stations have a lot of frames to send, they probably alternate in usage. When one station is sending, the other one refrains from sending. We can say that, in this case, each station on average, sends at a rate of 5 Mbps. Figure 13.8 shows the situation.



Figure 13.8 Sharing bandwidth

A bridge divides the network into two or more networks. Bandwidth-wise, each network is independent. For example, in Figure 13.9, a network with 12 stations is divided into two networks, each with 6 stations. Now each network has a capacity of 10 Mbps. The 10-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. In a network with a heavy load, each station theoretically is offered 10/6 Mbps instead of 10/12 Mbps, assuming that the traffic is not going through the bridge. It is obvious that if we further divide the network, we can gain more bandwidth for each segment. For example, if we use a four-port bridge, each station is now offered 10/3 Mbps, which is 4 times more than an unbridged network.



b. With bridging

Figure 13.9 A network with and without a bridge

Separating Collision Domains

Another advantage of a bridge is the separation of the collision domain. Figure 13.10 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 tremendously. Without bridging, 12 stations contend for access to the medium; with bridging only 3 stations contend for access to the medium.



Figure 13.10 Collision domains in an unbridged network and a bridged network

2.4.2. 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.

A layer 2 switch is an N-port bridge with additional sophistication that allows faster handling of the packets. Evolution from a bridged Ethernet to a switched Ethernet was a big step that opened the way to an even faster Ethernet, as we will see. Figure 13.11 shows a switched LAN.



Figure 13.11 Switched Ethernet

2.4.3. Full-Duplex Ethernet

One of the limitations of 10Base5 and 10Base2 is that communication is half-duplex (10Base-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 13.12 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 13.12 Full-duplex switched Ethernet

No Need for CSMA/CD

In full-duplex switched Ethernet, there is no need for the *CSMAICD* method. In a full duplex switched Ethernet, each station is connected to the switch via two separate links. Each station or switch can send and receive independently without worrying about collision. Each link is a point-to-point dedicated path between the station and the switch. There is no longer a need for carrier sensing; there is no longer a need for collision detection. The job of the MAC layer becomes much easier. The carrier sensing and collision detection functionalities of the MAC sub-layer can be turned off.

2.5. Fast Ethernet(IEEE 802.3u)

Fast Ethernet was designed to compete with LAN protocols such as FDDI or Fiber Channel (or Fibre Channel, as it is sometimes spelled). 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. The goals of Fast Ethernet can be summarized as follows:

- a. Upgrade the data rate to 100 Mbps.
- b. Make it compatible with Standard Ethernet.
- c. Keep the same 48-bit address.
- d. Keep the same frame format.
- e. Keep the same minimum and maximum frame lengths.

Topology

Fast Ethernet is designed to connect two or more stations together. If there are only two stations, they can be connected point-to-point. Three or more stations need to be connected in a star topology with a hub or a switch at the center, as shown in Figure 13.13.



Figure 13.13 Fast Ethernet topology

Implementation

Fast Ethernet implementation at the physical layer can be categorized as either two-wire or four-wire. The two-wire implementation can be either category 5 UTP (100Base-TX) or fiber-optic cable (100Base-FX). The four-wire implementation is designed only for category 3 UTP (100Base-T4). See Figure 13.14.



Figure 13.14 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	100m	100m	100m
Line encoding	MLT-3	NRZ-I	8B/6T

Fable 13.2 Sum	mary of Fast Eth	nernet implementations
----------------	------------------	------------------------

2.6. Gigabit Ethernet(IEEE 802.3z)

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. The goals of the Gigabit Ethernet design can be summarized as follows:

- a. Upgrade the data rate to 1 Gbps.
- b. Make it compatible with Standard or Fast Ethernet.
- c. Use the same 48-bit address.
- d. Use the same frame format.
- e. Keep the same minimum and maximum frame lengths.
- f. To support autonegotiation as defined in Fast Ethernet.

Gigabit Ethernet has two distinctive approaches for medium access: half-duplex and fullduplex. Almost all implementations of Gigabit Ethernet follow the full-duplex approach.

Topology

Gigabit Ethernet is designed to connect two or more stations. If there are only two stations, they can be connected point-to-point. Three or more stations need to be connected in a star topology with a hub or a switch at the center. Another possible configuration is to connect several star topologies or let a star topology be part of another as shown in Figure 13.15.



Figure 13.15 Topologies of Gigabit Ethernet

Implementation

Gigabit Ethernet can be categorized as either a two-wire or a four-wire implementation. The two-wire implementations use fiber-optic cable (1000Base-SX, short-wave, or 1000Base-LX, long-wave), or STP (1000Base-CX). The four-wire version uses category 5 twisted-pair cable (1000Base-T). In other words, we have four implementations, as shown in Figure 13.16.



Figure 13.16 Gigabit Ethernet implementations

Summary

Table 13.3 is a summary of the Gigabit Ethernet implementations.

Characteristics	1000Base-SX	1000Base-LX	1000Base-CX	1000Base-T
Media	Fiber	Fiber	STP	CAT 5 UTP
	Short wave	Long wave	511	
Number of wires	2	2	2	4
Maximum length	550m	5000m	25m	100m
Line encoding	NRZ	NRZ	NRZ	4D-PAM5

Table 13.3 Summary of Gigabit Ethernet implementations

2.7. Ten-Gigabit Ethernet(IEEE 802.3ae)

The IEEE committee created Ten-Gigabit Ethernet and called it Standard 802.3ae. The goals of the Ten-Gigabit Ethernet design can be summarized as follows:

- a. Upgrade the data rate to 10 Gbps.
- b. Make it compatible with Standard, Fast, and Gigabit Ethernet.
- c. Use the same 48-bit address.
- d. Use the same frame format.
- e. S. Keep the same minimum and maximum frame lengths.
- f. Allow the interconnection of existing LANs into a metropolitan area network (MAN) or a wide area network (WAN).
- g. Make Ethernet compatible with technologies such as Frame Relay and ATM.

Ten-Gigabit Ethernet operates only in **full duplex mode** which means there is no need for contention; *CSMA/CD* is not used in Ten-Gigabit Ethernet.

Implementation

Ten-Gigabit Ethernet is designed for using fiber-optic cable over long distances. Three implementations are the most common: 10GBase-S, 10GBase-L, and 10GBase-E. Table 13.4 shows a summary of the Ten-Gigabit Ethernet implementations:

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

Table 13.4 Summary of Ten-Gigabit Ethernet implementations

Reference:

1. B. A. Forouzan: Data Communications and Networking, Fourth edition, TMH .