COMPUTER NETWORKS

UNIT-4

Syllabus: The Medium Access Control Sublayer-The Channel Allocation Problem-Static Channel Allocation-Assumptions for Dynamic Channel Allocation, Multiple Access Protocols- Aloha- Carrier Sense Multiple Access Protocols-Collision-Free Protocols-Limited Contention Protocols- Wireless LAN Protocols, Ethernet-Classic Ethernet Physical Layer-Classic Ethernet MAC Sublayer Protocol-Ethernet Performance-Fast Ethernet Gigabit Ethernet-10-Gigabit Ethernet- Retrospective on Ethernet, Wireless Lans-The 802.11 Architecture and Protocol Stack- The 802.11 Physical Layer-The802.11 MAC Sublayer Protocol-The 805.11 Frame Structure- Services

~~~~~~

The Channel Allocation Problem:

How to allocate a single broadcast channel among competing users. The channel might be a portion of the wireless spectrum in a geographic region, or a single wire or optical fiber to which multiple nodes are connected. It does not matter. In both cases, the channel connects each user to all other users and any user who makes full use of the channel interferes with other users who also wish to use the channel.

Static Channel Allocation:

The traditional way of allocating a single channel, such as a telephone trunk, among multiple competing users is to chop up its capacity by using one of the multiplexing schemes such as FDM (Frequency Division Multiplexing). If there are N users, the bandwidth is divided into N equal-sized portions, with each user being assigned one portion. Since each user has a private frequency band, there is now no interference among users. When there is only a small and constant number of users, each of which has a steady stream or a heavy load of traffic, this division is a simple and efficient allocation mechanism. A wireless example is FM radio stations. Each station gets a portion of the FM band and uses it most of the time to broadcast its signal.

The basic problem is that when some users are quiescent, their bandwidth is simply lost. They are not using it, and no one else is allowed to use it either. A static allocation is a poor fit to most computer systems, in which data traffic is extremely bursty, often with peak traffic to mean traffic ratios of 1000:1. Consequently, most of the channels will be idle most of the time.

If we were to use time division multiplexing (TDM) and allocate each user every Nth time slot, if a user does not use the allocated slot, it would just lie fallow. The same would hold if we split up the networks physically.
~~~~~~

In principle, CDMA could be used but there are a number of constraints which makes it difficult to apply in the general case

e. g. , senders and receivers must know the chip sequence beforehand

Requires high synchronization

Assumptions for Dynamic Channel Allocation:

1 **Independent Traffic**. The model consists of N independent **stations** (e. g. , computers, telephones), each with a program or user that generates frames for transmission. The expected number of frames generated in an interval of length t is t, where is a constant (the arrival rate of new frames). Once a frame has been generated, the station is blocked and does nothing until the frame has been successfully transmitted.

2. **Single Channel**. A single channel is available for all communication. All stations can transmit on it and all can receive from it. The stations are assumed to be equally capable, though protocols may assign them different roles (e. g. , priorities).

3. **Observable Collisions**. If two frames are transmitted simultaneously, they overlap in time and the resulting signal is garbled. This event is called a **collision**. All stations can detect that a collision has occurred. A collided frame must be transmitted again later. No errors other than those generated by collisions occur.

4 **Continuous or Slotted Time**. Time may be assumed continuous, in which case frame transmission can begin at any instant. Alternatively, time may be slotted or divided into discrete intervals (called slots). Frame transmissions must then begin at the start of a slot. A slot may contain 0, 1, or more frames, corresponding to an idle slot, a successful transmission, or a collision, respectively.

5. **Carrier Sense or No Carrier Sense**. With the carrier sense assumption, stations can tell if the channel is in use before trying to use it. No station will attempt to use the channel while it is sensed as busy. If there is no carrier sense, stations cannot sense the channel before trying to use it. They just go ahead and transmit. Only later can they determine whether the transmission was successful.

MULTIPLE ACCESS PROTOCOLS:

When nodes or stations are connected and use a common link, called a multipoint or broadcast link, we need a multiple-access protocol to coordinate access to the link. The problem of controlling the access to the medium is similar to the rules of speaking in an assembly. Many

formal protocols have been devised to handle access to a shared link. We categorize them into three groups. Protocols belonging to each group are shown above.

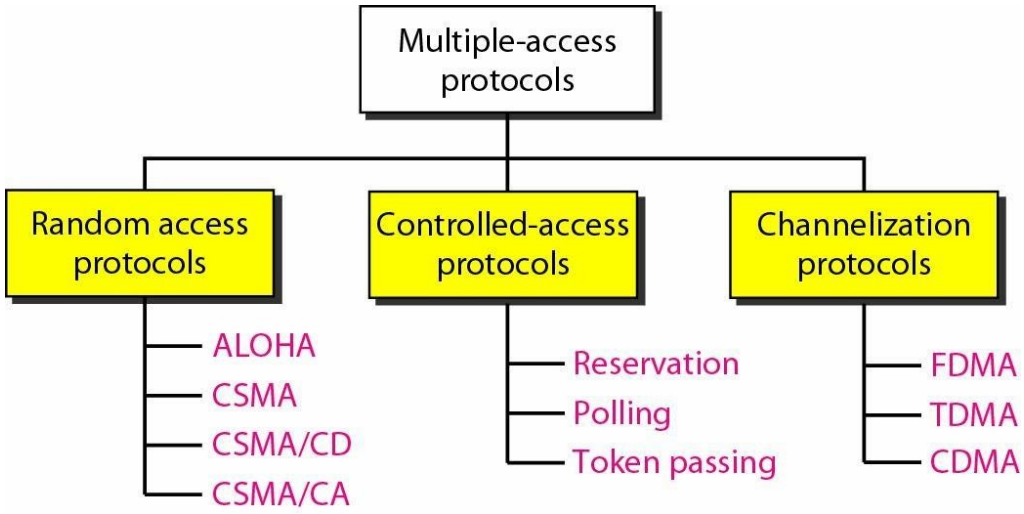

RANDOM ACCESS

In a random access method, each station has the right to the medium without being controlled by any other station. However, if more than one station tries to send, there is an access conflict- collision- and the frames will be either destroyed or modified. To avoid access conflict or to resolve it when it happens, each station follows a procedure to answer the following questions: When can the station access the medium?

What can the station do if the medium is busy?

How can the station determine the success or failure of the transmission?

What can the station do if there is an access conflict?

The random access methods have evolved from a very interesting protocol known as ALOHA, which used a very simple procedure called multiple access (MA). The method was improved with the addition of a procedure that forces the station to sense the medium before transmitting. This was called carrier sense multiple access. This method later evolved into two parallel methods: carrier senses multiple access with collision detection (CSMA/CD) and carrier sense multiple access with collision avoidance (CSMA/CA). CSMA/CD tells the station what to do when a collision is detected. CSMA/CA tries to avoid the collision.

ALOHA:

In the 1970s, Norman Abramson and his colleagues at the University of Hawaii devised a new and elegant method to solve the ***channel allocation problem***. Their work has been extended by

many researchers since then (Abramson, 1985). We will discuss two versions of ALOHA here: ***pure and slotted.***

Pure ALOHA

The basic idea of an ALOHA system is simple: let users transmit whenever they have data to be sent. There will be collisions, of course, and the colliding frames will be damaged. However, due to the feedback property of broadcasting, a sender can always find out whether its frame was destroyed by listening to the channel, the same way other users do.

If the frame was destroyed, the sender just waits a random amount of time and sends it again. The waiting time must be random or the same frames will collide over and over, in lockstep. Systems in which multiple users share a common channel in a way that can lead to conflicts are widely known ***as contention systems.***

A sketch of frame generation in an ALOHA system is given in figure. We have made the frames all the same length because ***the throughput of ALOHA systems*** is maximized by having a uniform frame size rather than by allowing variable length frames.

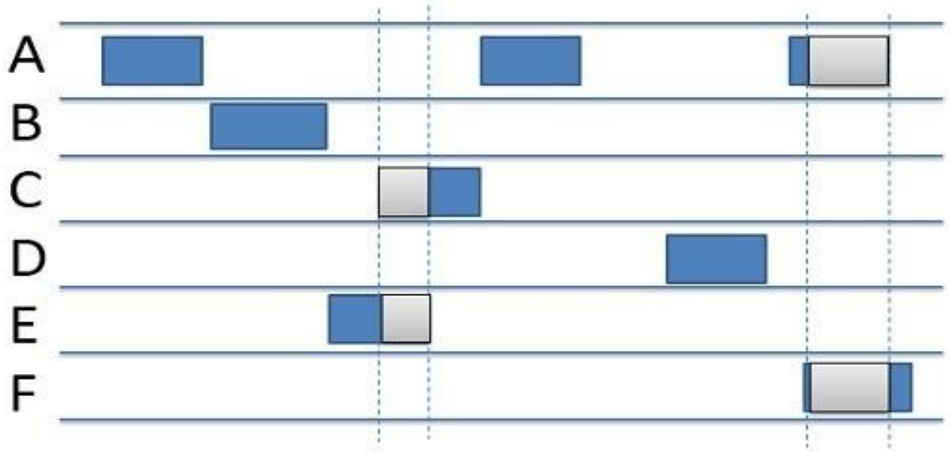

In pure ALOHA; frames are transmitted at completely arbitrary times.

Whenever two frames try to occupy the channel at the same time, there will be a collision and both will be garbled. If the first bit of a new frame overlaps with just the last bit of a frame almost finished, both frames will be totally destroyed and both will have to be retransmitted later. The checksum cannot (and should not) distinguish between a total loss and a near miss.

Let the "***frame time***" denote the amount of time needed to transmit the standard, fixed-length frame (i. e. , ***the frame length divided by the bit rate***). At this point we assume that the infinite

*population of users generates new frames according to a **Poisson distribution with mean N frames per frame time**. If **N > 1**, the user community is generating frames at a **higher rate** than the channel can handle, and nearly every frame will suffer a **collision**. For reasonable **throughput** we would expect **0 < N < 1.***

*In addition to the new frames, the stations also generate retransmissions of frames that previously suffered collisions. Let us further assume that the probability of **k transmission** attempts per frame time, old and new combined, is also Poisson, with **mean G per frame time**.*

*Clearly, **G ≥N**. At low load (i. e. , N ≈0), there will be few collisions, hence few retransmissions, so G≈N. At high load there will be many collisions, so G > N. Under all loads, the throughput, S, is just the offered load, G, times the probability, P_0, of a transmission succeeding—that is, $S = GP_0$, where P_0 is the probability that a frame does not suffer a collision.*

A frame will not suffer a collision if no other frames are sent within one frame time of its start, as shown in figure. Let t be the time required to send a frame. If any other user has generated a frame between time t_0 and $t_0 + t$, the end of that frame will collide with the beginning of the shaded one. In fact, the shaded frame's fate was already sealed even before the first bit was sent, but since in pure ALOHA a station does not listen to the channel before transmitting, it has no way of knowing that another frame was already underway. Similarly, any other frame started between $t_0 + t$ and $t_0 + 2t$ will bump into the end of the shaded frame.

Figure 4-2.

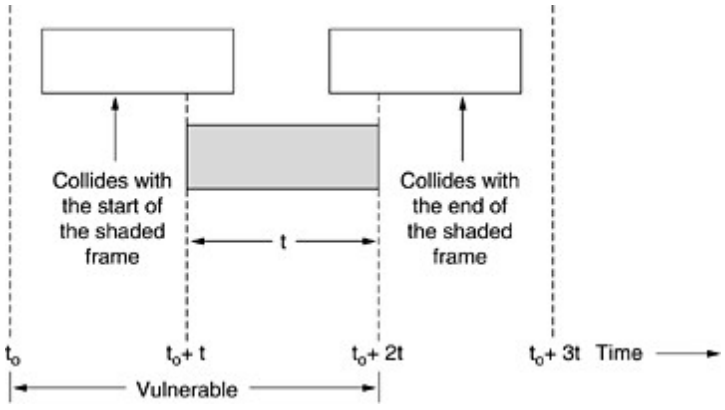

Vulnerable period for the shaded frame.

The probability that k frames are generated during a given frame time is given by the Poisson distribution:

$Pr[k] = (G^k e^{-G}) / k!$ ---------- Equation 4

so the probability of zero frames is just e^{-G}. In an interval two frame times long, the mean number of frames generated is 2G. The probability of no other traffic being initiated during the entire vulnerable period is thus given by $P_0 = e^{-2G}$. Using $S = GP_0$, we get

$S = Ge^{-2G}$

The relation between the offered traffic and the throughput is shown in figure.

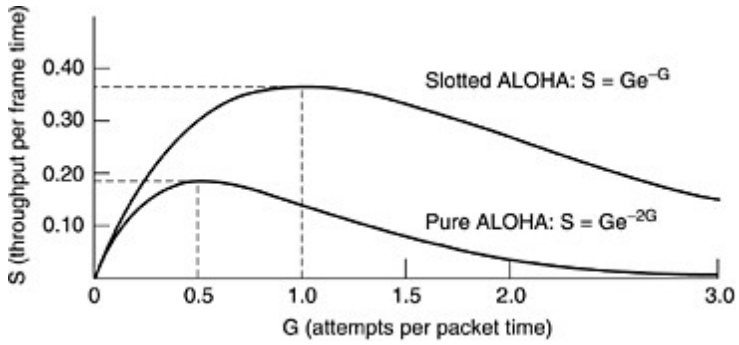

Throughput versus offered traffic for ALOHA systems.

The maximum throughput occurs at G = 0.5, with S = 1/2e, which is about 0.184. In other words, the best we can hope for is a channel utilization of 18 percent. This result is not very encouraging, but with everyone transmitting at will, we could hardly have expected a 100 percent success rate.

Slotted ALOHA

In Roberts' method, which has come to be known as slotted ALOHA, in contrast to pure ALOHA, a computer is not permitted to send whenever a carriage return is typed. Instead, it is required to wait for the beginning of the next slot.

As a result of the exponential dependence of E upon G, small increases in the channel load can drastically reduce its performance.

Slotted Aloha is important for a reason that may not be initially obvious. It was devised in the 1970s, used in a few early experimental systems, and then almost forgotten. When Internet access over the cable was invented, all of a sudden there was a problem of how to allocate a shared channel among multiple competing users, and slotted Aloha was pulled out of the garbage can to save the day.

MAC Address:

Medium access control MAC address or a hardware address refers to a unique value assigned to a network adapter on a LAN. This address is used to identify each node in a network in a distant manner. MAC addresses are 48 bit long and is represented as 12digit hexadecimal number as follows

AA:AA:AA:BB:BB:BB or AA-AA-AA-BB-BB-BB

Where the first half specifies the manufacturers id of an adapter and the second half specifies the serial number associated with the adapter. This serial number is assigned by the manufacturer to an adapter.

An OSI is the most widely used network architecture where the data link layer is divided into MAC and logical link control sub layers; MAC addresses are used to uniquely identify a computer on a network.

CARRIER SENSE MULTIPLE ACCESS PROTOCOLS:

With slotted ALOHA the best channel utilization that can be achieved is 1/e. In local area networks, however, it is possible for stations to detect what other stations are doing, and adapt their behavior accordingly. These networks can achieve a much better utilization than 1/e. Protocols in which stations listen for a carrier (i.e., a transmission) and act accordingly are called ***carrier sense protocols****.*

Persistent and Non-persistent CSMA:

1- persistent CSMA: *When a station has data to send, it first listens to the channel to see if anyone else is transmitting at that moment. If the channel is busy, the station waits until it becomes idle. When the station detects an idle channel, it transmits a frame. If a collision occurs, the station waits a random amount of time and starts all over again. The protocol is called 1- persistent because the station transmits with a probability of 1 when it finds the channel idle.*

The propagation delay has an important effect on the performance of the protocol. There is a small chance that just after a station begins sending, another station will become ready to send and sense the channel. If the first station's signal has not yet reached the second one, the latter will sense an idle channel and will also begin sending, resulting in a collision.

Even if the propagation delay is zero, there will still be collisions. If two stations become ready in the middle of a third station's transmission, both will wait politely until the transmission ends and then both will begin transmitting exactly simultaneously, resulting in a collision. If they were not so impatient, there would be fewer collisions. This approach will lead to a higher performance than pure ALOHA.

Non-persistent CSMA: In this protocol, a station senses the channel. If no one else is sending, the station begins doing so itself. However, if the channel is already in use, the station does not continually sense it for the purpose of seizing it immediately upon detecting the end of the previous transmission. Instead, it waits a random period of time and then repeats the algorithm. Consequently, this algorithm leads to better channel utilization but longer delays than 1- persistent CSMA.

P-persistent CSMA: It applies to slotted channels. When a station becomes ready to send, it senses the channel. If it is idle, it transmits with a probability p. With a probability q = 1 - p, it defers until the next slot. If that slot is also idle, it either transmits or defers again, with probabilities p and q. This process is repeated until either the frame has been transmitted or another station has begun transmitting. In the latter case, the unlucky station acts as if there had been a collision (i. e. , it waits a random time and starts again). If the station initially senses the channel busy, it waits until the next slot and applies the above algorithm. Figure shows the computed throughput versus offered traffic for all three protocols, as well as for pure and slotted ALOHA.

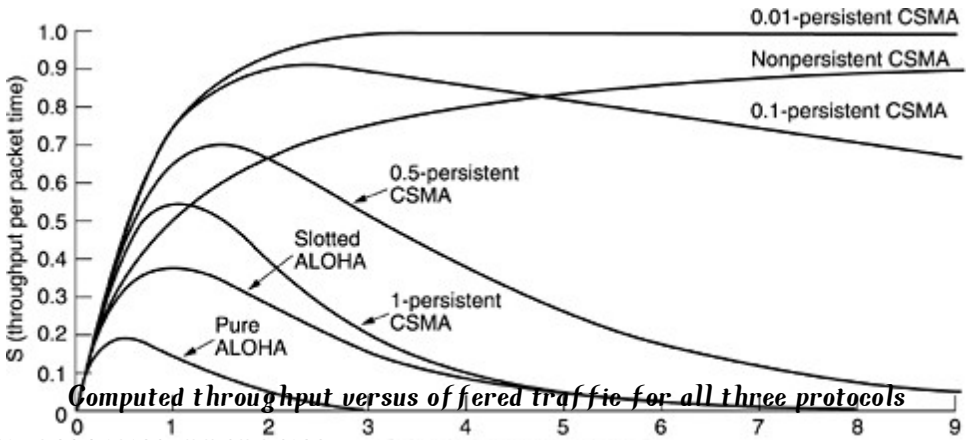

Computed throughput versus offered traffic for all three protocols

CSMA WITH COLLISION DETECTION:

Another improvement is for stations to abort their transmissions as soon as they detect a collision In other words, if two stations sense the channel to be idle and begin transmitting simultaneously, they will both detect the collision almost immediately. Quickly terminating damaged frames saves time and bandwidth. This protocol, known as **CSMA/CD (CSMA with Collision Detection)** is widely used on LANs in the MAC sub layer.

CSMA/CD, as well as many other LAN protocols, uses the conceptual model of figure. At the point marked t_0, a station has finished transmitting its frame. Any other station having a frame to send may now attempt to do so. If two or more stations decide to transmit simultaneously, there will be a collision. Collisions can be detected by looking at the power or pulse width of the received signal and comparing it to the transmitted signal.

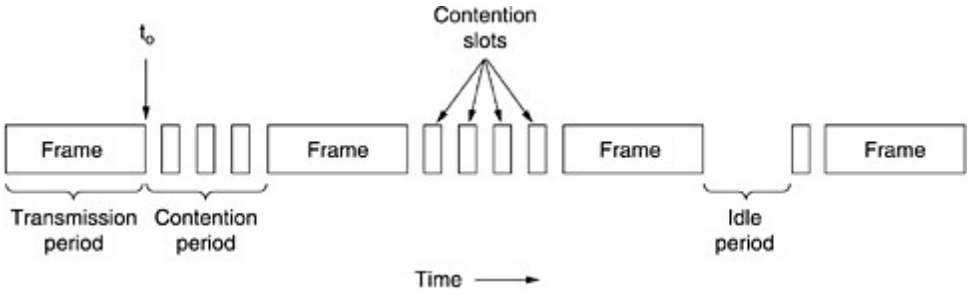

CSMA/CD Conceptual model

After a station detects a collision, it aborts its transmission, waits a random period of time, and then tries again, assuming that no other station has started transmitting in the meantime. Therefore, our model for CSMA/CD will consist of alternating contention and transmission periods, with idle periods occurring when all stations are quiet (e. g. , for lack of work).

Suppose that two stations both begin transmitting at exactly time t_0. How long will it take them to realize that there has been a collision? The answer to this question is vital to determining the length of the contention period and hence what the delay and throughput will be. The minimum time to detect the collision is then just the time it takes the signal to propagate from one station to the other.

Based on this reasoning, you might think that a station not hearing a collision for a time equal to the full cable propagation time after starting its transmission could be sure it had seized the cable. By "seized," we mean that all other stations knew it was transmitting and would not interfere. This conclusion is wrong. Consider the following worst-case scenario. Let the time for a signal to propagate between the two farthest stations be . At t_0, one station begins transmitting. At - , an instant before the signal arrive at the most distant station, that station also begins transmitting. Of course, it detects the collision almost instantly and stops, but the little noise burst caused by the collision does not get back to the original station until time 2 - . In other words, in the worst case a station cannot be sure that it has seized the channel until it has transmitted for 2 without hearing a collision. For this reason we will model the contention

interval as a slotted ALOHA system with slot width 2 . On a 1-km long coaxial cable, 5 μsec. For simplicity we will assume that each slot contains just 1 bit. Once the channel has been seized, a station can transmit at any rate it wants to, of course, not just at 1 bit per 2 sec.

It is important to realize that collision detection is an analog process. The station's hardware must listen to the cable while it is transmitting. If what it reads back is different from what it is putting out, it knows that a collision is occurring. The implication is that the signal encoding must allow collisions to be detected (e. g. , a collision of two 0-volt signals may well be impossible to detect). For this reason, special encoding is commonly used.

It is also worth noting that a sending station must continually monitor the channel, listening for noise bursts that might indicate a collision. For this reason, CSMA/CD with a single channel is inherently a half-duplex system. It is impossible for a station to transmit and receive frames at the same time because the receiving logic is in use, looking for collisions during every transmission.

To avoid any misunderstanding, it is worth noting that no MAC-sublayer protocol guarantees reliable delivery. Even in the absence of collisions, the receiver may not have copied the frame correctly for various reasons (e. g. , lack of buffer space or a missed interrupt).

CSMA with Collision Avoidance:

CSMA/CA protocol is used in wireless networks because they cannot detect the collision so the only solution is collision avoidance.

CSMA/CA avoids the collisions using three basic techniques.

(i) Interframe space

(ii) Contention window

(iii) Acknowledgements

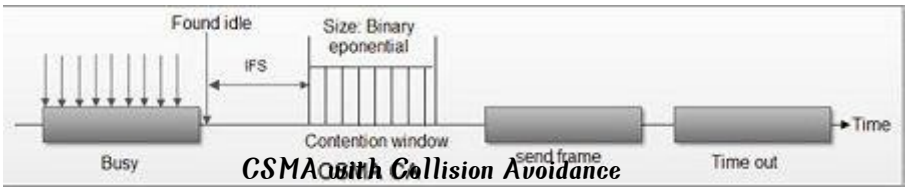

CSMA with Collision Avoidance

Whenever the channel is found idle, the station does not transmit immediately. It waits for a period of time called inter frame space (IFS).

When channel is sensed to be idle, it may be possible that same distant station may have already started transmitting and the signal of that distant station has not yet reached other stations.

Therefore the purpose of IFS time is to allow this transmitted signal to reach other stations.

If after this IFS time, the channel is still idle, the station can send, but it still needs to wait a time equal to contention time

A station that is ready to send chooses a random number of slots as its wait time. The number of slots in the window changes according to the binary exponential back-off strategy. It means that it is set of one slot the first time and then doubles each time the station cannot detect an idle channel after the IFS time. This is very similar to the p-persistent method except that a random outcome defines the number of slots taken by the waiting station . In contention window the station needs to sense the channel after each time slot. If the station finds the channel busy, it does not restart the process. It just stops the timer & restarts it when the channel is sensed as idle.

3. Acknowledgement

Despite all the precautions, collisions may occur and destroy the data.

The positive acknowledgment and the time-out timer can help guarantee that receiver has received the frame.

CSMA/CA Procedure:

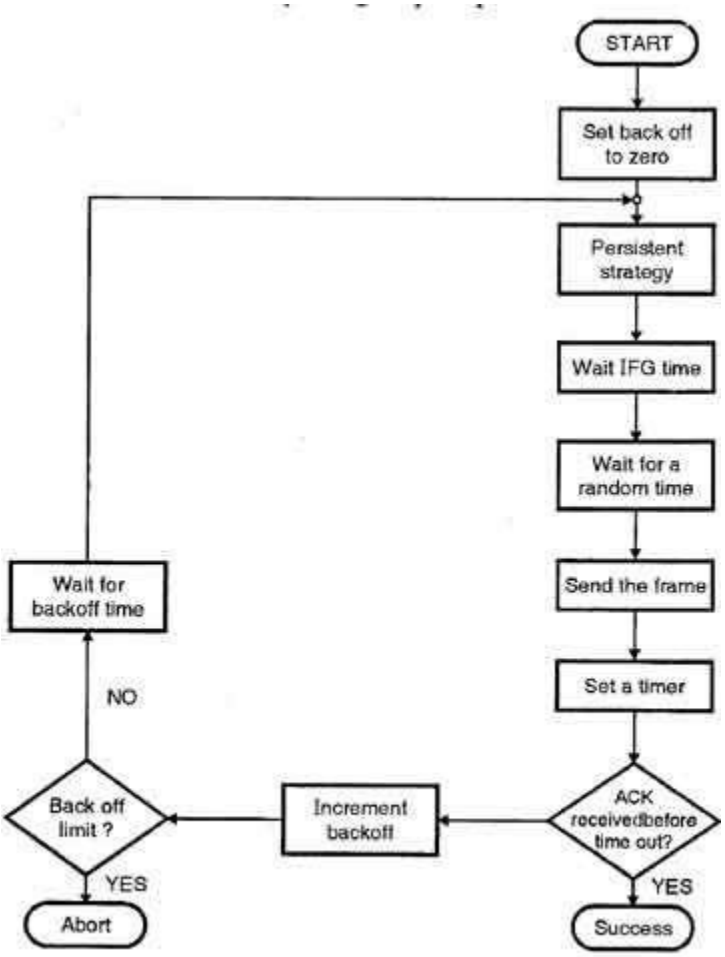

CSMA/CA procedure

Shows the flow chart explaining the principle of CSMA/CA.

This is the CSMA protocol with collision avoidance. The station ready to transmit, senses the line by using one of the persistent strategies. As soon as it find the line to be idle, the station waits for an IFG (Interframe gap) amount of time.

If then waits for some random time and sends the frame. After sending the frame, it sets a timer and waits for the acknowledgement from the receiver.

If the acknowledgement is received before expiry of the timer, then the transmission is successful. But if the transmitting station does not receive the expected acknowledgement before the timer expiry then it increments the back off parameter, waits for the back off time and resenses the line.

CONTROLLED ACCESS

In controlled access, the stations consult one another to find which station has the right to send. A station cannot send unless it has been authorized by other stations. We discuss three popular controlled-access methods.

Reservation

In the reservation method, a station needs to make a reservation before sending data. Time is divided into intervals. In each interval, a reservation frame precedes the data frames sent in that interval. If there are N stations in the system, there are exactly N reservation mini-slots in the reservation frame. Each mini-slot belongs to a station. When a station needs to send a data frame, it makes a reservation in its own mini-slot. The stations that have made reservations can send their data frames after the reservation frame.

Figure 12.18 shows a situation with five stations and a five-minislot reservation frame. In the first interval, only stations 1, 3, and 4 have made reservations. In the second interval, only station 1 has made a reservation.

Reservation

Polling

Polling works with topologies in which one device is designated as a primary station and the other devices are secondary stations. All data exchanges must be made through the primary device even when the ultimate destination is a secondary device. The primary device controls the link; the secondary devices follow its instructions. It is up to the primary device to determine which device is allowed to use the channel at a given time. The primary device, therefore, is always the initiator of a session (see Figure). If the primary wants to receive data, it asks the secondary's if they have anything to send; this is called poll function. If the primary wants to send data, it tells the secondary to get ready to receive; this is called select function.

Select

The *select* function is used whenever the primary device has something to send. Remember that the primary controls the link. If the primary is neither sending nor receiving data, it knows the

link is available. If it has something to send, the primary device sends it. What it does not know, however, is whether the target device is prepared to receive. So the primary must alert the secondary to the upcoming transmission and wait for an acknowledgment of the secondary's ready status. Before sending data, the primary creates and transmits a select (SEL) frame, one field of which includes the address of the intended secondary.

Poll

The poll function is used by the primary device to solicit transmissions from the secondary devices. When the primary is ready to receive data, it must ask (poll) each device in turn if it has anything to send. When the first secondary is approached, it responds either with a NAK frame if it has nothing to send or with data (in the form ofa data frame) if it does. If the response is negative (a NAK frame), then the primary polls the next secondary in the same manner until it finds one with data to send. When the response is positive (a data frame), the primary reads the frame and returns an acknowledgment (ACK frame), verifying its receipt.

Select Poll

Collision-Free Protocols:

Although collisions do not occur with CSMA/CD once a station has unambiguously captured the channel, they can still occur during the contention period. Some protocols that resolve the contention for the channel without any collisions at all, not even during the contention period. Most of these protocols are not currently used in major systems, but in a rapidly changing field, having some protocols with excellent properties available for future systems is often a good thing.

We assume that there are exactly N stations, each programmed with a unique address from 0 to N - 1. It does not matter that some stations may be inactive part of the time. We also assume that propagation delay is negligible.

A Bit-Map Protocol:

In our first collision-free protocol, the **basic bit-map method**, each contention period consists of exactly N slots. If station 0 has a frame to send, it transmits a 1 bit during the slot 0. No other station is allowed to transmit during this slot. Regardless of what station 0 does, station 1 gets the opportunity to transmit a 1 bit during slot 1, but only if it has a frame queued. In general, station j may announce that it has a frame to send by inserting a 1 bit into slot j. After all N slots have passed by, each station has complete knowledge of which stations wish to transmit. At that point, they begin transmitting frames in numerical order.

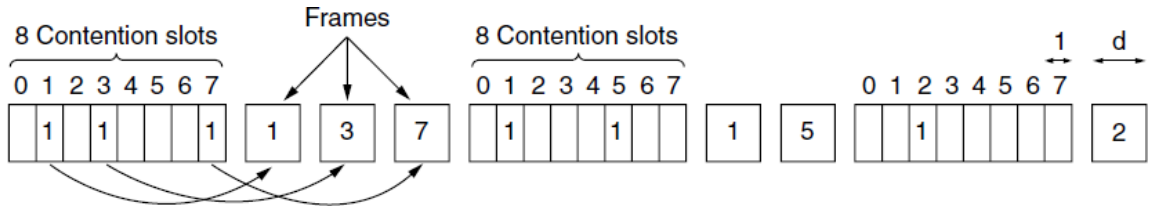

The basic bit-map protocol

Protocols like this in which the desire to transmit is broadcast before the actual transmission are called **reservation protocols** because they reserve channel ownership in advance and prevent collisions.

Token Passing

In the token-passing method, the stations in a network are organized in a logical ring. In other words, for each station, there is a *predecessor* and a *successor*. The predecessor is the station which is logically before the station in the ring; the successor is the station which is after the station in the ring. The current station is the one that is accessing the channel now. The right to this access has been passed from the predecessor to the current station. The right will be passed to the successor when the current station has no more data to send.

But how is the right to access the channel passed from one station to another? In this method, a special packet called a token circulates through the ring. The possession of the token gives the station the right to access the channel and send its data.

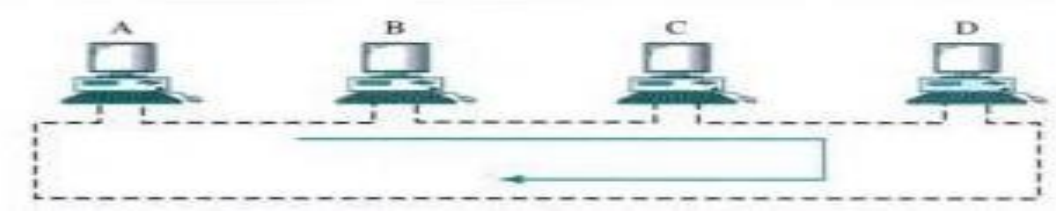

Token passing network

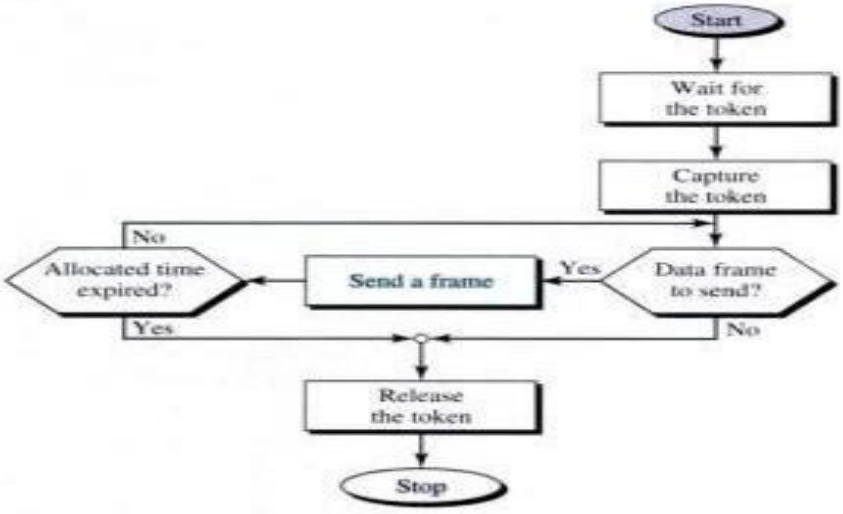

Token passing procedure

When a station has some data to send, it waits until it receives the token from its predecessor. It then holds the token and sends its data. When the station has no more data to send, it releases the token, passing it to the next logical station in the ring. The station cannot send data until it receives the token again in the next round. In this process, when a station receives the token and has no data to send, it just passes the data to the next station.

Token management is needed for this access method. Stations must be limited in the time they can have possession of the token. The token must be monitored to ensure it has not been lost or destroyed. For example, if a station that is holding the token fails, the token will disappear from the network. Another function of token management is to assign priorities to the stations and to the types of data being transmitted. And finally, token management is needed to make low- priority stations release the token to high priority stations.

Binary Countdown:

A problem with the basic bit-map protocol, and by extension token passing, is that the overhead is 1 bit per station, so it does not scale well to networks with thousands of stations. We can do better than that by using binary station addresses with a channel that combines transmissions. A

station wanting to use the channel now broadcasts its address as a binary bit string, starting with the highorder bit. All addresses are assumed to be the same length. The bits in each address position from different stations are BOOLEAN ORed together by the channel when they are sent at the same time. We will call this protocol ***binary countdown****.*

To avoid conflicts, an arbitration rule must be applied: as soon as a station sees that a high-order bit position that is 0 in its address has been overwritten with a 1, it gives up. For example, if stations 0010, 0100, 1001, and 1010 are all trying to get the channel, in the first bit time the stations transmit 0, 0, 1, and 1, respectively. These are ORed together to form a 1. Stations 0010 and 0100 see the 1 and know that a higher-numbered station is competing for the channel, so they give up for the current round. Stations 1001 and 1010 continue.

The next bit is 0, and both stations continue. The next bit is 1, so station 1001 gives up. The winner is station 1010 because it has the highest address. After winning the bidding, it may now transmit a frame, after which another bidding cycle starts. The protocol is illustrated in Fig. 4-8. It has the property that higher- numbered stations have a higher priority than lower- numbered stations, which may be either good or bad, depending on the context.

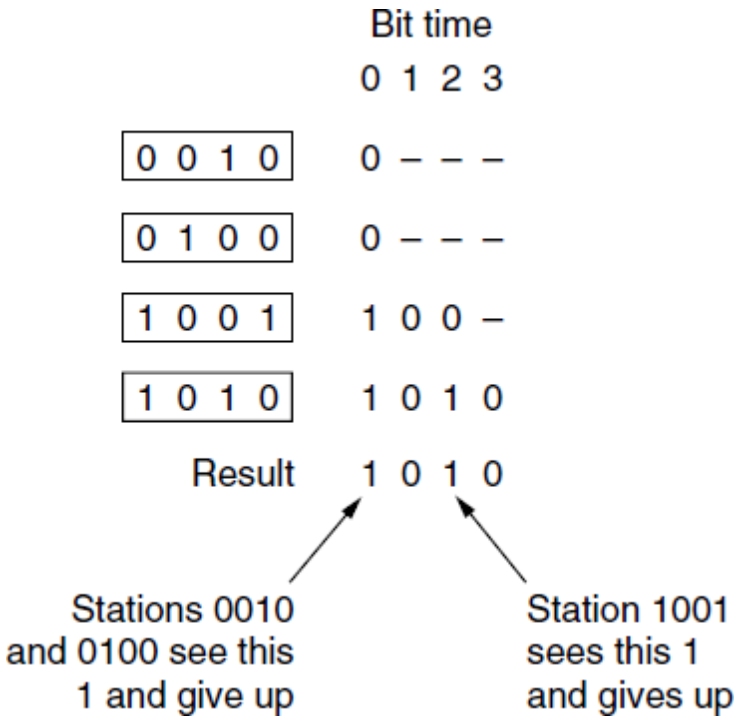

The binary countdown protocol. A dash indicates silence.

Limited-Contention Protocols:

It would be nice if we could combine the best properties of the contention and collision-free protocols, arriving at a new protocol that used contention at low load to provide low delay, but used a collision-free technique at high load to provide good channel efficiency. Such protocols, which we will call ***limited-contention protocols****.*

The only contention protocols we have studied have been symmetric. That is, each station attempts to acquire the channel with some probability, p, with all stations using the same p. Interestingly enough, the overall system performance can sometimes be improved by using a protocol that assigns different probabilities to different stations.

Wireless LAN Protocols:

A common configuration for a wireless LAN is an office building with access points (APs) strategically placed around the building. The APs are wired together using copper or fiber and provide connectivity to the stations that talk to them.

There is an even more important difference between wireless LANs and wired LANs. A station on a wireless LAN may not be able to transmit frames to or receive frames from all other stations because of the limited radio range of the stations. In wired LANs, when one station sends a frame, all other stations receive it. The absence of this property in wireless LANs causes a variety of complications.

A naive approach to using a wireless LAN might be to try CSMA: just listen for other transmissions and only transmit if no one else is doing so. The trouble is, this protocol is not really a good way to think about wireless because what matters for reception is interference at the receiver, not at the sender. To see the nature of the problem, consider Fig. 4-11, where four wireless stations are illustrated. For our purposes, it does not matter which are APs and which are laptops. The radio range is such that A and B are within each other's range and can potentially interfere with one another. C can also potentially interfere with both B and D, but not with A.

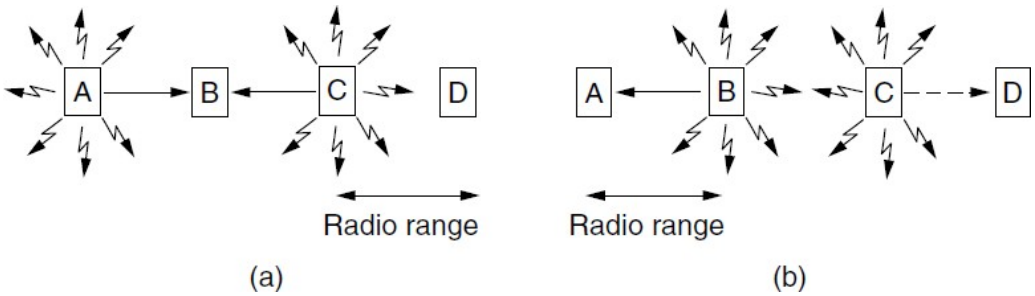

A wireless LAN. (a) A and C are hidden terminals when transmitting to B. (b) B and C are exposed terminals when transmitting to A and D.

First consider what happens when A and C transmit to B, as depicted in Fig. (a). If A sends and then C immediately senses the medium, it will not hear A because A is out of range. Thus C will falsely conclude that it can transmit to B. If C does start transmitting, it will interfere at B, wiping out the frame from A. (We assume here that no CDMA-type scheme is used to provide multiple channels, so collisions garble the signal and destroy both frames.) We want a MAC protocol that will prevent this kind of collision from happening because it wastes bandwidth. The problem of a station not being able to detect a potential competitor for the medium because the competitor is too far away is called the **hidden terminal problem**.

Now let us look at a different situation: B transmitting to A at the same time that C wants to transmit to D, as shown in Fig. 4-11(b). If C senses the medium, it will hear a transmission and falsely conclude that it may not send to D (shown as a dashed line). In fact, such a transmission would cause bad reception only in the zone between B and C, where neither of the intended receivers is located. We want a MAC protocol that prevents this kind of deferral from happening because it wastes bandwidth. The problem is called the **exposed terminal problem**.

An early and influential protocol that tackles these problems for wireless LANs is **MACA** (**Multiple Access with Collision Avoidance**) (Karn, 1990). The basic idea behind it is for the sender to stimulate the receiver into outputting a short frame, so stations nearby can detect this transmission and avoid transmitting for the duration of the upcoming (large) data frame. This technique is used instead of carrier sense.

MACA is illustrated in following Fig. Let us see how A sends a frame to B. A starts by sending an **RTS** (**Request To Send**) frame to B, as shown in Fig. 4-12(a). This short frame (30 bytes) contains the length of the data frame that will eventually follow. Then B replies with a **CTS** (**Clear To Send**) frame, as shown in Fig. 4-12(b). The CTS frame contains the data length (copied from the RTS frame). Upon receipt of the CTS frame, A begins transmission.

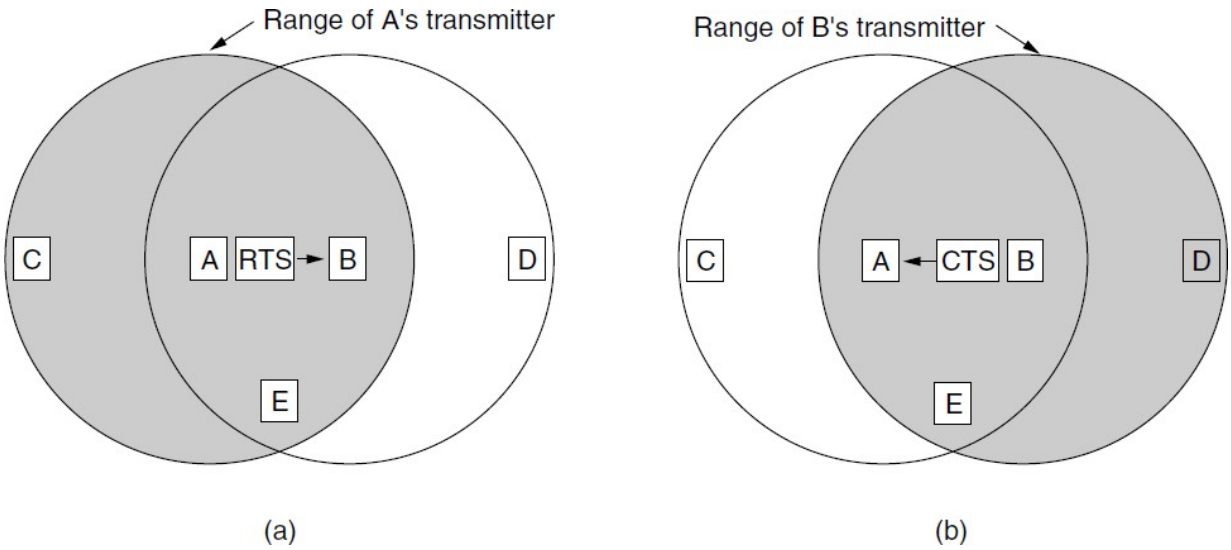

The MACA protocol. (a) A sending an RTS to B. (b) B responding with a CTS to A.

ETHERNET

- **LOCAL AREA NETWORK** (LAN) is a computer network that is designed for a limited geographic area such as a building or a campus. 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.
- The IEEE Standard Project 802, designed to regulate the manufacturing and interconnectivity between different LANs. Ethernet has changed to meet the market needs and to make use of the new technologies.
- E. g. ,

  IEEE 802. 3 → Ethernet

  IEEE 802. 11 → Wireless LAN (Wi-Fi)

  IEEE 802. 15 → Wireless PAN (Bluetooth, etc)

IEEE STANDARDS:

- In 1985, the Computer Society of the IEEE started a project, called **Project 802**, to define certain LAN standards. **Project 802**, consisting a variety of LAN standards deals primarily with physical layer and the data link layer issues.

- The standard was adopted by the American National Standards Institute (ANSI). In 1987, the International Organization for Standardization (ISO) also approved it as an international standard under the designation ISO 8802.
- 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 physical layer standards for different LAN protocols.

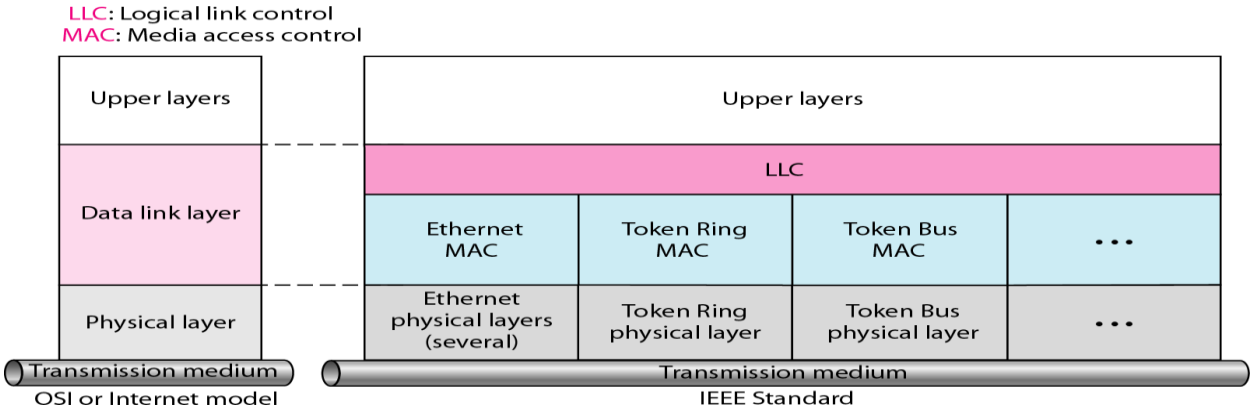

Fig: IEEE Standard of LANs

DATALINK LAYER:

- The data link layer in the IEEE standard is divided into two sublayers:

  Logical Link Control (LLC)

  Media Access Control (MAC)

Logical Link Control (LLC):

- 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 sub layer.
- The LLC provides one single data link control protocol for all IEEE LANs whereas, MAC sublayer provides different protocols for different LANs.
- A single LLC protocol can provide interconnectivity between different LANs because it makes the MAC sublayer transparent.

Framing: LLC defines a protocol data unit (PDU) that is somewhat similar to that of HDLC.

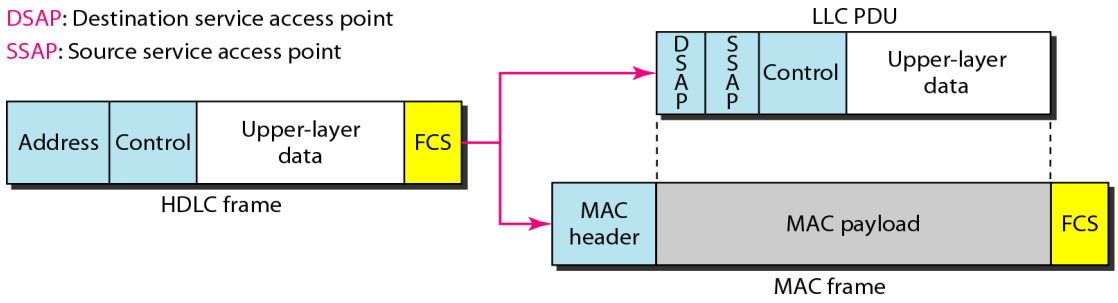

- *HDLC PDU is divided into PDU at the LLC and a frame at the MAC sublayer.*
- *The LLC PDU contains the address of the* ***destination service access point (DSAP)****,* ***source service access point (SSAP)****, control field and upper layer data.*
- *The control field is responsible for flow and error control.*
- ***Need for LLC*** *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):

- *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.*
- *Compared to 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.*

PHYSICAL LAYER:

The physical layer is dependent on the implementation and type of physical media used.

IEEE defines detailed specifications for each LAN implementation.

For example, although there is only one MAC sublayer for Standard Ethernet, there is a different physical layer specifications for each Ethernet implementations.

MANCHESTER ENCODING:

- *Manchester encoding is a synchronous clock encoding technique used by the physical layer to encode the clock and data of a synchronous bit stream. In this technique, the*

actual binary data to be transmitted over the cable are not sent as a sequence of logic 1's and 0's. Instead, the bits using straight binary encoding.

- *The data are represented NOT by logic 1 or 0, but with line transitions. A logic 0 is represented by a transition from* ***HIGH to LOW****, and a logic HIGH is represented by a transition from* ***LOW to HIGH****.*

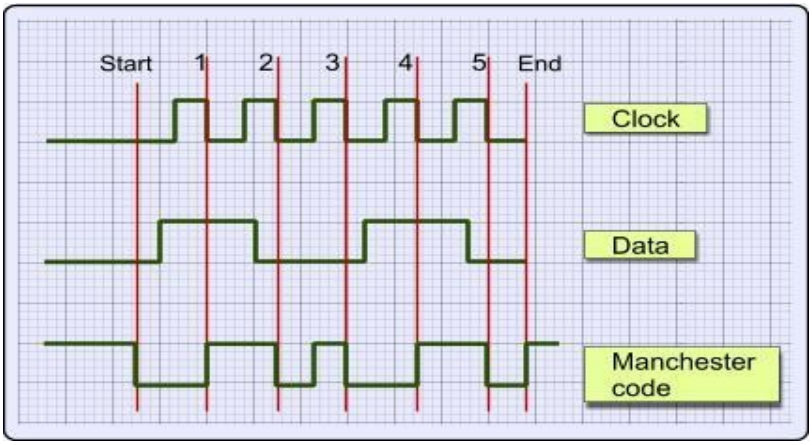

- *The data to encode is the binary number 10010, reading from left to right. The coding occurs on every* ***falling edge*** *of the clock. On the first falling edge of the clock, the coded signal has a LOW to HIGH transition, because the data is HIGH. On the second falling edge of the clock, the code has a HIGH to LOW transition because the data is LOW. The same algorithm is applied for the rest of the signal.*

The Differential Manchester Coding:

- *The Differential Manchester Code is a variation of the Manchester code*
- *The transmission line is kept HIGH when no data is sent. There are 2 encoding methods: The first is the "Transition on LOW" and the second is the "Transition on HIGH". For this example, to use the first method, "Transition on LOW".*
- *If the data bit is 0, then a polarity transition occurs (if was HIGH it goes LOW, and if it was LOW it goes HIGH), otherwise the line remains unchanged.*
- *In our example, the data that is transmitted is the binary '10011101' (starting from left to right). Each data bit is transmitted during negative transition of the clock. Between each bit transmission, the code line changes polarity. This is done* ***to help the receiver recreate the clock signal*** *and synchronize with the transmitter.*

- *Use the "Transition on LOW" method. This means that when a bit is transmitted, if this bit is ZERO the data line changes polarity. Otherwise, if the bit is 1 the data line polarity remains unchanged.*

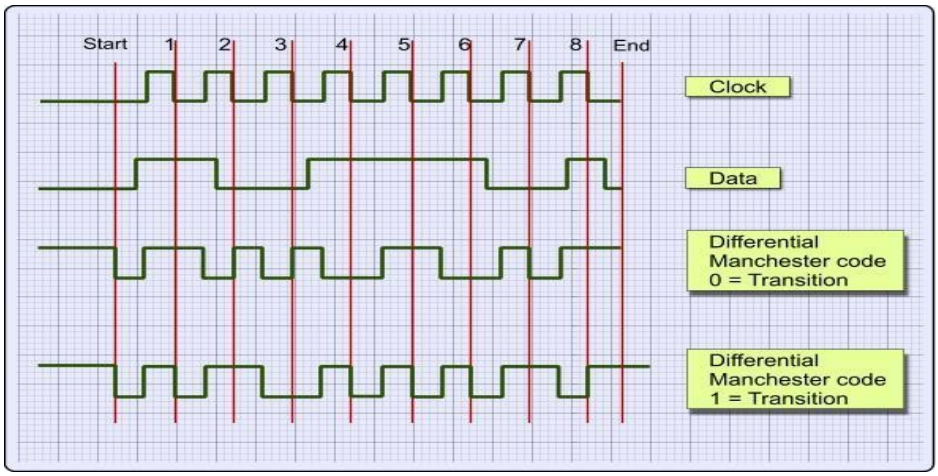

- *Suppose now that we want to transmit this byte (10011101). The code line is HIGH. The transmission is initiated by pulling the code line LOW. After half a pulse, the output is pulled HIGH. This is part of the synchronization transition that occurs every middle of a bit transmission. Half a pulse next, the first bit is transmitted. The first bit is 1 (starting from left), so the code line polarity remains unchanged. After half a pulse, the code line polarity changes state and goes LOW. After one full pulse, the second bit is about to be transmitted. This bit is the 0, so the code line changes polarity and goes HIGH. The same algorithm is used to transmit all 8 bits of the data. Finally, the code line is pulled HIGH and the transmission ends.*

STANDARD ETHERNET:

The original Ethernet was created in 1976 at Xerox's palo alto research center(PARC) The ETHERNET has gone through four generations:

- **Standard Ethernet (10 Mbps)**
- **Fast Ethernet (100 Mbps)**
- **Gigabit Ethernet (1 Gbps)**
- **Ten-Gigabit Ethernet (10 Gbps)**

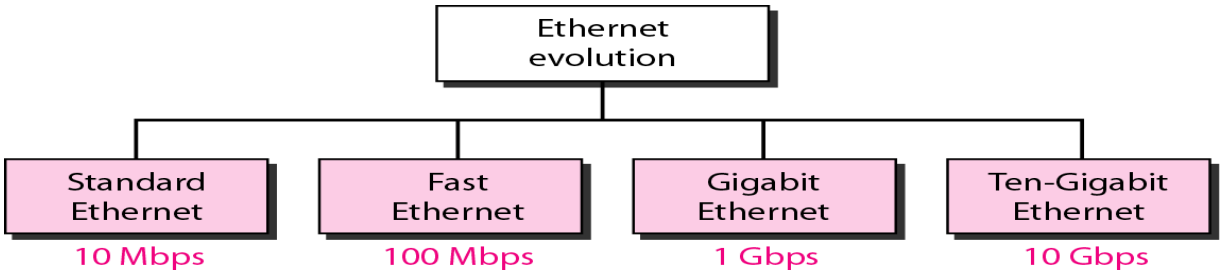

MAC Sublayer:

- *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*

Preamble: 56 bits of alternating 1s and 0s.
SFD: Start frame delimiter, flag (10101011)

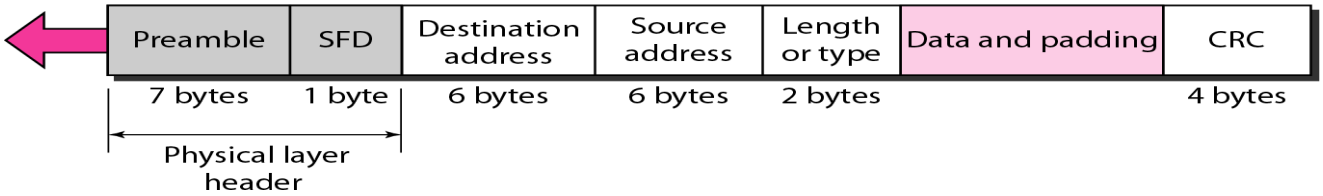

- ***Preamble.*** *The first field of the 802.3 frame contains 7 bytes (56 bits) of alternating 0s and 1s that alerts the receiving system to the coming frame and enables it to synchronize its input timing. The preamble is added at the physical layer and is not part of the frame.*
- ***Start frame delimiter (SFD).*** *The second field (1 byte: 10101011) signals the beginning of the frame. The last 2 bits is 11 and alerts the receiver that the next field is the destination address.*

- **Destination address (DA).** 6 bytes long and contains the physical address of the destination station.
- **Source address (SA).** 6 bytes long and contains the physical address of the sender station.
- **Length or type.** The length field to define the number of bytes in the data field.
- **Data.** This field carries data encapsulated from the upper-layer protocols. It is a minimum of 46 and a maximum of 1500 bytes.
- **CRC.** The last field contains error detection information, in this case a CRC-32.

Frame Length

Ethernet has imposed restrictions on both the minimum and maximum lengths of a frame.

Minimum length:

An Ethernet frame needs to have a minimum length of 512 bits or 64 bytes which includes **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.

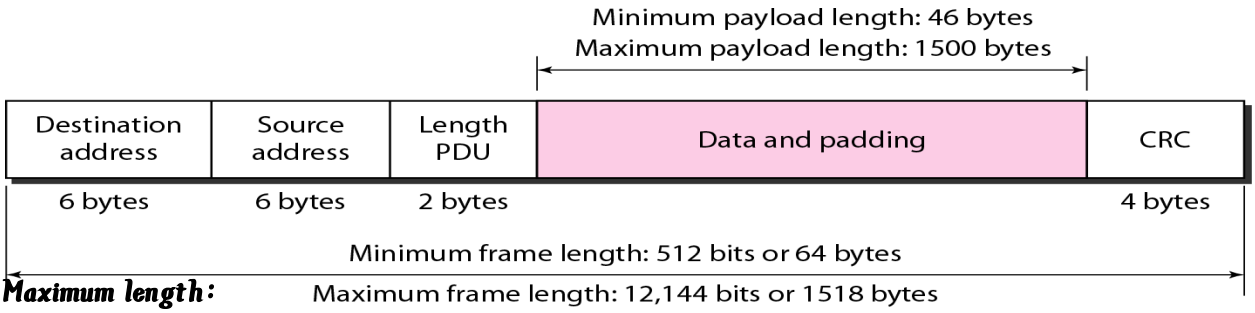

Maximum length:

The standard defines the maximum length of a frame 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 reasons.

1. To reduce the size of the buffer.
2. To prevents one station from monopolizing the shared medium, blocking others stations that have data to send.

Frame length:

Minimum: 64 bytes (512 bits) Maximum: 1518 bytes (12,144 bits)

Addressing

Each station on an Ethernet network has its own network interface card (NIC). The NIC provides 6-bytes physical address to the station. Hexadecimal notation only.

- **Unicast address**: *defines only one recipient; the relationship between the sender and the receiver is one-to-one.*
- **Multicast address**: *defines a group of addresses; the relationship between the sender and the receivers is one-to-many.*

The least significant bit of the 1st byte defines the type of address.

If the bit is 0, the address is UNICAST.

If the bit is 1, the address is MULTICAST.

- **Broadcast address**: *The recipients are all the stations on the LAN. A broadcast destination address is forty-eight 1's.*

Access Method: CSMA/CD:

Standard Ethernet uses 1-persistent CSMA/CD.

Slot Time: *In an Ethernet network, the round-trip time required for a frame to travel from one end of a maximum-length network to the other plus the time needed to send the jam sequence is called the slot time.*

Slot time =round-trip time + time required to send the jam sequence

The slot time in Ethernet is defined in bits. It is the time required for a station to send 512 bits. This means that the actual slot time depends on the data rate; for traditional 10-Mbps Ethernet it is 51.2 µs.

Slot Time and Collision: *It was chosen to allow the proper functioning of CSMA/CD. To understand the situation, let us consider two cases.*

1. **The sender sends a minimum-size packet of 512 bits.**

Before the sender can send the entire packet out, the signal travels through the network and reaches the end of the network. If there is another signal at the end of the network, a collision occurs. The sender has the opportunity to abort the sending of the frame and to send a jam sequence to inform other stations of the collision. The round-trip time plus the time required to send the jam sequence should be less than the time needed for the sender to send the minimum frame, 512 bits. The sender needs to be aware of the collision before it is too late, that is, before it has sent the entire frame.

2. **The sender sends a frame larger than the minimum size (between 512 and 1518 bits).**

If the station has sent out the first 512 bits and has not heard a collision, it is guaranteed that collision will never occur during the transmission of this frame. The reason is that the signal will reach the end of the network in less than one-half the slot time. If all stations follow the CSMA/CD protocol, they have already sensed the existence of the signal (carrier) on the line and have refrained from sending. If they sent a signal on the line before one-half of the slot time expired, a collision has occurred and the sender has sensed the collision.

Slot Time and Maximum Network Length: *There is a relationship between the slot time and the maximum length of the network.*

MaxLength= PropagationSpeedx SlotTime/2

MaxLength= $(2x10^8)$ X(51. 2 $X10^{-6}/2$)= 5120m

Of course, we need to consider the delay times in repeaters and interfaces, and the time required to send the jam sequence. These reduce the maximum-length of a traditional Ethernet network to 2500 m, just 48 percent of the theoretical calculation.

MaxLength=2500 m.

Physical Layer Implementations:

The Standard Ethernet defines several physical layer implementations.

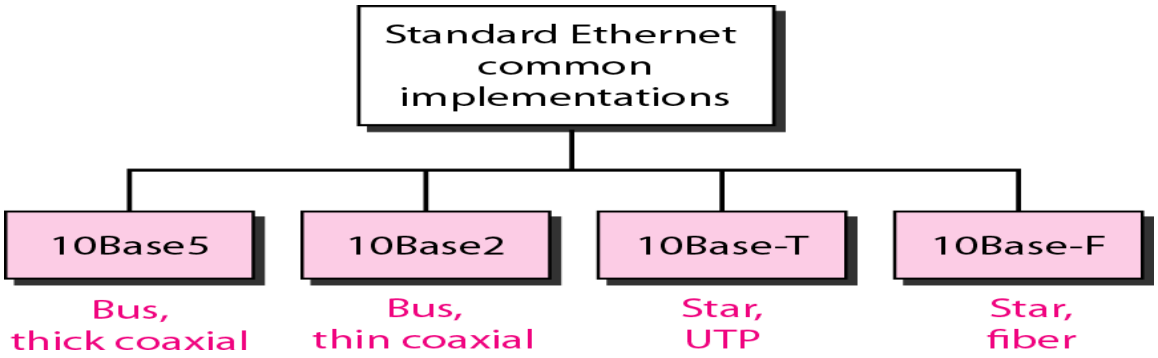

Encoding and Decoding:

- *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.*
- ***Manchester encoding is self-synchronous**, providing a transition at each bit interval.*

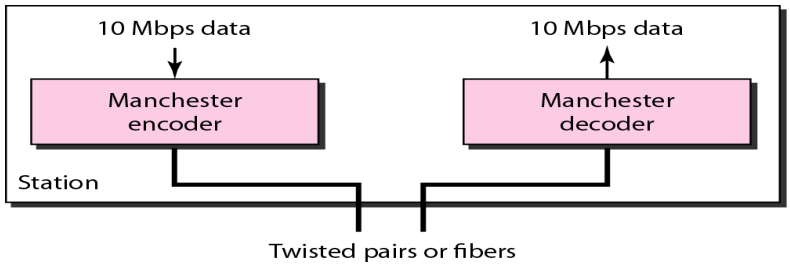

10Base5: Thick Ethernet:

- *The first implementation is called **10Base5, thick Ethernet, or Thicknet.***
- *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**.*

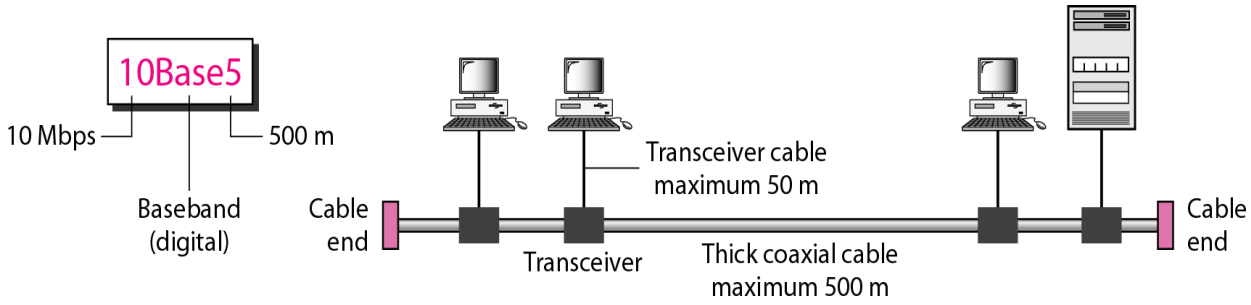

10Base2: Thin Ethernet:

- The second implementation is called 10Base2, **thin** Ethernet, or **Cheapernet**.
- 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.
- **collision** 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.

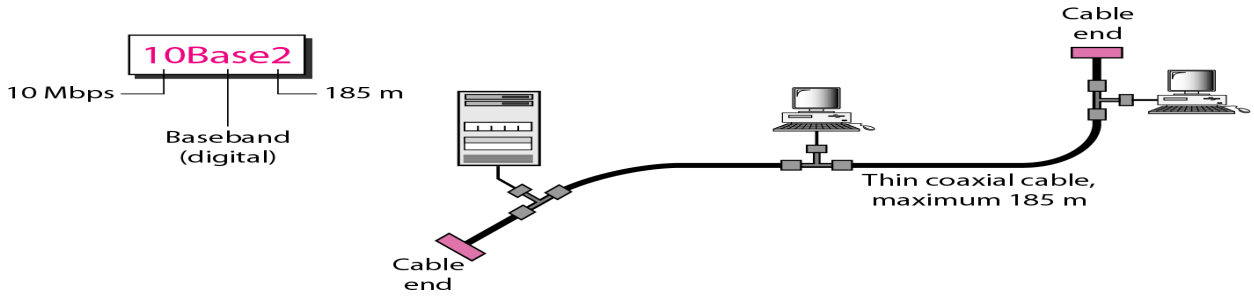

10Base-T: Twisted-Pair Ethernet:

- 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.

- *The maximum length of the twisted cable as* ***100 m****, to minimize the effect of attenuation in the twisted cable.*

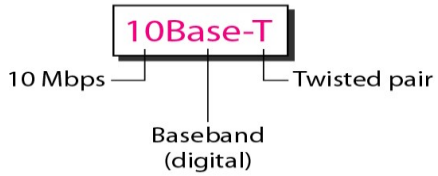

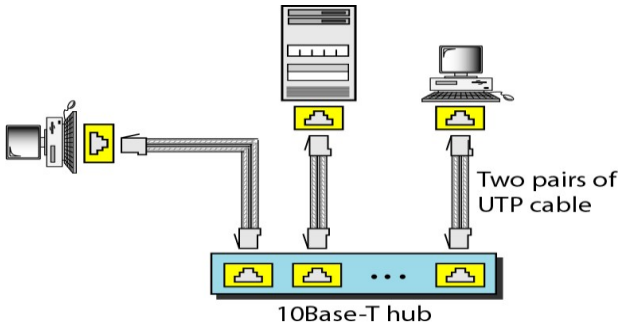

10Base-F: Fiber Ethernet

- *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.*

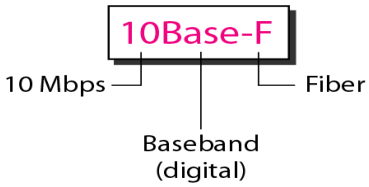

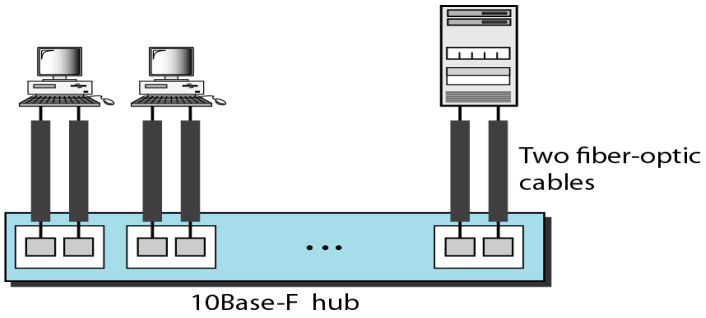

Cabling: *The types of Ethernet and their standards.*

Name	Cable	Max. segment	Nodes	Advantages
10Base5	Thick coax	500m	100	Good for backbones
10Base2	Thin coax	185m	30	cheapest
10Base-T	Twisted pair	100m	1024	Easy maintenance
10Base-F	Fiber optics	2000m	1024	Between building good

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

FAST ETHERNET:

- *Fast Ethernet was designed to compete with LAN protocols such as FDDI or Fiber Channel.*
- *The goals of Fast Ethernet as follows:*

1. *Upgrade the data rate to 100 Mbps.*
2. *Make it compatible with Standard Ethernet.*
3. *Uses the same 48-bit address.*
4. *Uses the same frame format.*
5. *Uses the same minimum and maximum frame lengths.*

MAC Sublayer:

MAC sublayer of the fast Ethernet is same as that of the traditional Ethernet.

1. **Access method:** *The Access method also remains the same for the half-duplex approach. It is same as CSMA/CD. for fullduplex Fast Ethernet, there is no need for CSMA/CD. But, the CSMA/CD is used for backward compatibility with Standard Ethernet.*
2. **Frame Format:** *Same as that of traditional Ethernet.*
3. **Minimum and maximum frame lengths:** *Same as that of traditional Ethernet.*
4. **Addressing:** *Same as that of traditional Ethernet.*

Auto negotiation:

A new feature added to Fast Ethernet is called autonegotiation. It allows a Autonegotiation allows two devices to negotiate the **mode or data rate of operation** *between two devices.*

It was designed particularly for the following purposes:

1. *To allow incompatible devices to connect to one another*
2. *To allow one device to have multiple capabilities.*
3. *To allow a station to check a hub's capabilities.*

VIEW

Physical Layer

The physical layer in Fast Ethernet is more complicated than the one in Standard Ethernet.

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.*

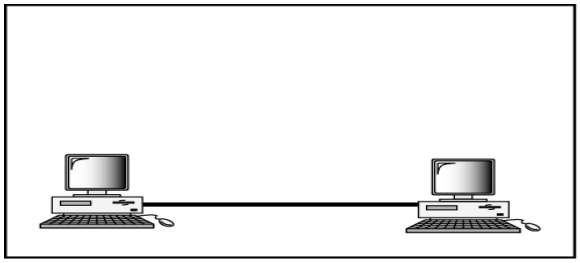

a. Point-to-point

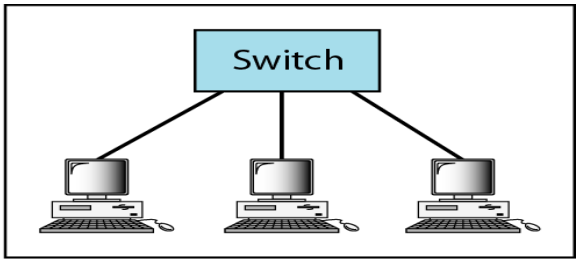

b. Star

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).*

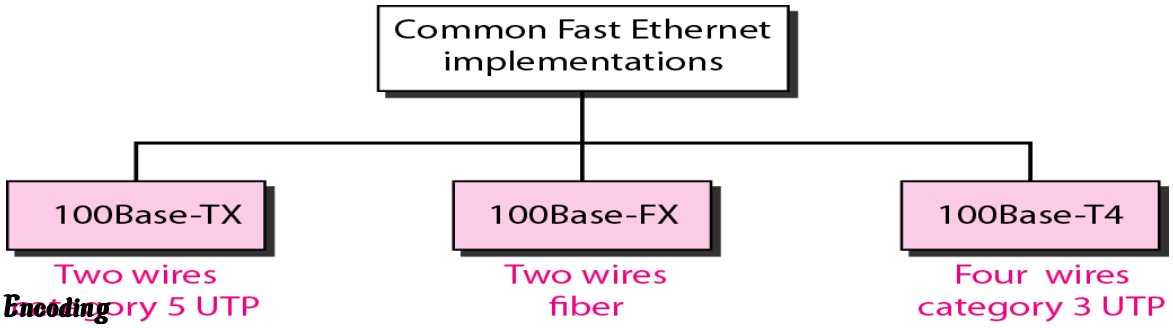

Encoding

- *Manchester encoding needs a 200-Mbaud bandwidth for a data rate of 100 Mbps, which makes it unsuitable for a medium such as twisted-pair cable.*
- *The Fast Ethernet designers try to discover some* ***alternative encoding/decoding scheme.***
- *Therefore, three different encoding schemes were chosen.*

100Base-TX:

- It uses **two pairs** of twisted-pair cable (either category 5 UTP or STP).
- For this implementation, the **MLT-3**(Multi-Level Transmit) scheme was selected since it has good bandwidth performance. However, since MLT-3 is not a self-synchronous line coding scheme, **4B/5B block coding** is used to provide bit synchronization by preventing the occurrence of a long sequence of 0s and 1s.
- This creates a data rate **of 125 Mbps**, which is fed into MLT-3 for encoding.

100Base-FX:

- It uses **two pairs** of fiber-optic cables. Optical fiber can easily handle high bandwidth requirements by using simple encoding schemes.
- The designers of 100Base-FX selected the **NRZ-I encoding scheme** for this implementation.
- However, NRZ-I has a bit synchronization problem for long sequences of 0s (or 1s, based on the encoding). To overcome this problem, the designers used **4B/5B block encoding** as we described for 100Base-TX.
- The block encoding increases the bit rate from **100 to 125 Mbps**, which can easily be handled by fiber-optic cable.

100Base-T4:

- The implementation **uses four pairs of UTP** for transmitting 100 Mbps.
- Encoding/decoding in 100Base-T4 is more complicated. As this implementation uses category 3 UTP, each twisted-pair cannot easily handle more than 25 Mbaud.
- In this design, one pair switches between sending and receiving. Three pairs of UTP category 3, however, can handle only **75 Mbaud** (25 Mbaud) each. We need to use an encoding scheme that converts 100 Mbps to a 75 Mbaud signal.
- In 8B/6T, eight data elements are encoded as six signal elements. This means that 100 Mbps uses only **(6/8) x 100** Mbps, or 75 Mbaud.

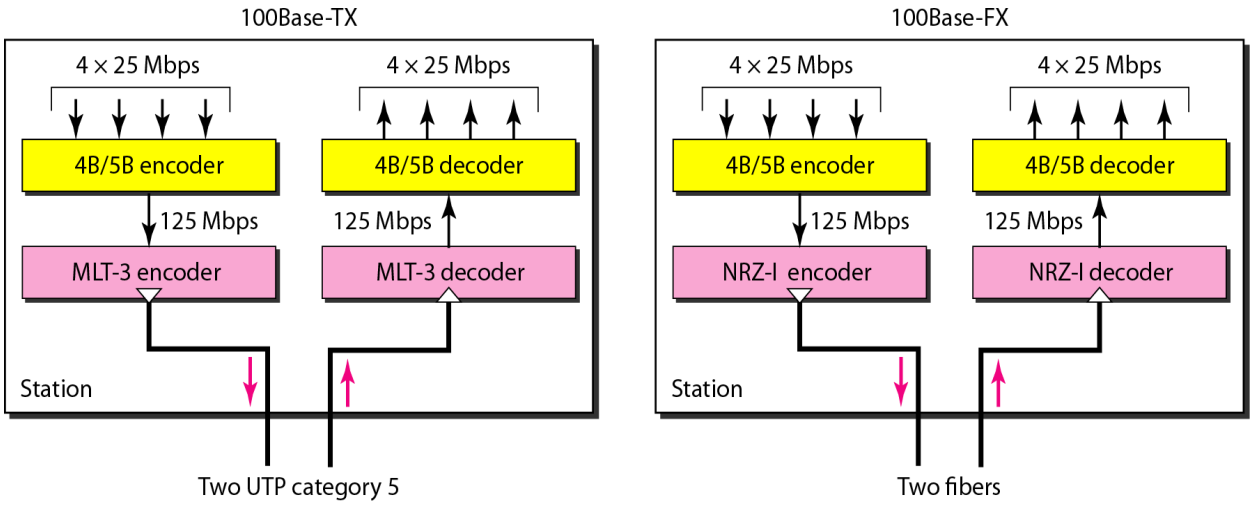

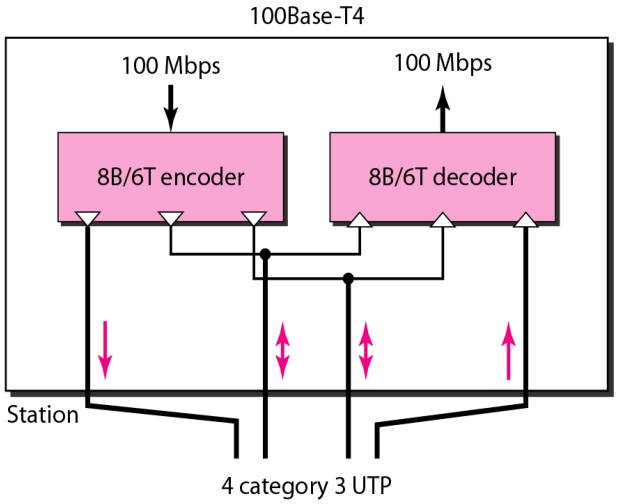

IEEE 802.11:

IEEE has defined the specifications for a wireless LAN, called IEEE 802.11, which covers the physical and data link layers.

Architecture

- *The standard defines two kinds of services:*

Basic service set (BSS)

Extended service set (ESS)

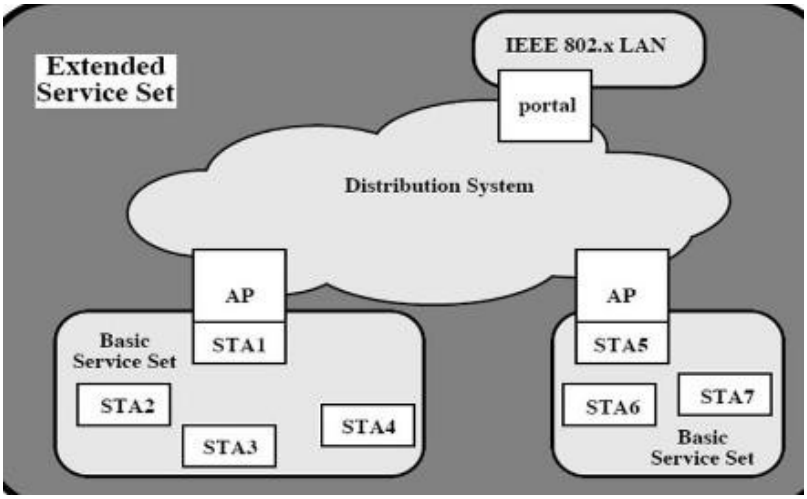

Basic Service Set:

- *The building block of a wireless LAN in IEEE 802.11 is the basic service set (BSS). The BSS contains several stations(STAs) and an optional central base station, known as the* ***access point (AP).***
- *The BSS without an AP is a stand-alone network and cannot send data to other BSSs. It is called an* ***ad hoc architecture,*** *stations can form a network without the need of an AP; they can locate one another and agree to be part of a BSS.*
- *A BSS with an AP is sometimes referred to as an* ***infrastructure network.***

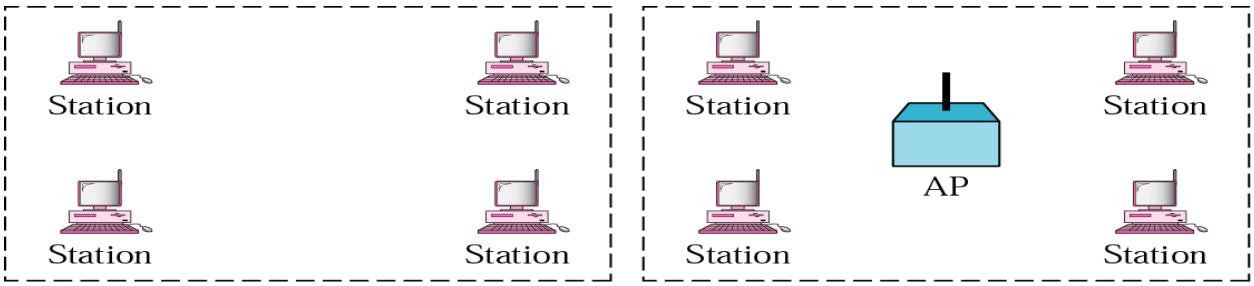

BSS without AP BSS with an AP

Extended Service Set

- *An extended service set (ESS) is made up of* ***two or more BSSs with APs.*** *The BSSs are connected through a* ***distribution system,*** *which is usually a wired LAN. The distribution system connects the APs in the BSSs.*
- *The extended service set uses two types of stations:* ***mobile and stationary.*** *The*

mobile stations *are normal stations inside a BSS.*

The ***stationary stations*** *are AP stations that are part of a wired LAN.*

- When BSSs are connected, the stations within reach of one another can communicate without the use of an AP. However, communication between two stations in two different BSSs usually occurs via two APs.

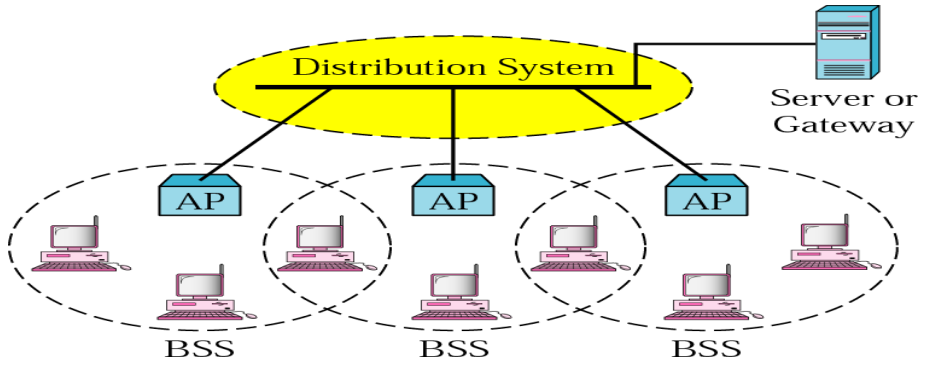

.

Station Types (STA):

- This component is used to connect to the wireless medium. Station can be any device, it can be a mobile device, a Network interface card etc. It provides the services of authentication, privacy and delivery of the data.

Access Point(AP):

- It provides same services as STA.
- The function of the AP is to provide both the communication to the wired LAN and the local relay function for the BSS.

MAC Sublayer:

- IEEE 802.11 defines two MAC sublayers:

Distributed coordination function (DCF) Point

coordination function (PCF).

The relationship between the two MAC sublayers, the LLC sublayer, and the physical layer.

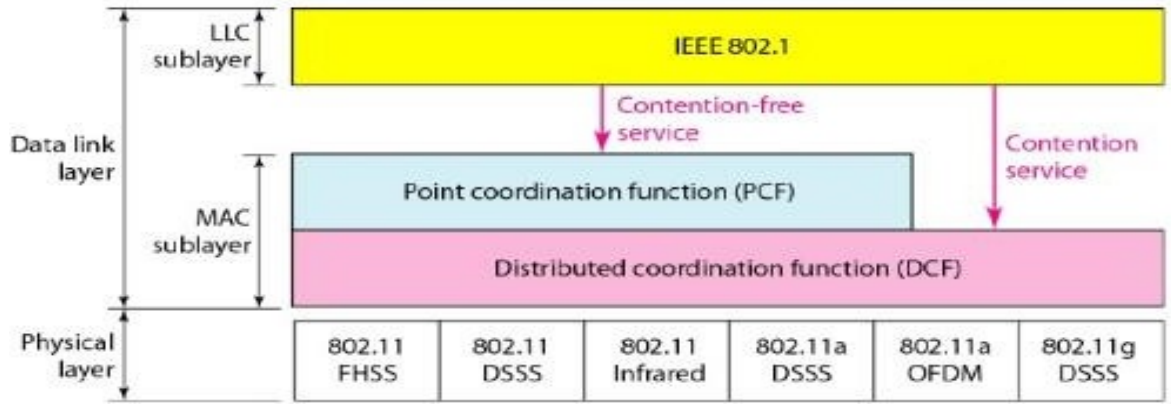

Distributed Coordination Function

Distributed coordination function (DCF) is the protocol defined at the MAC sublayer by IEEE.

DCF uses CSMA/CA as the access method.

Wireless LANs cannot implement **CSMA/CD** *for three reasons:*

1. *For collision detection a station must be able to send data and receive collision signals at the same time. This can mean costly stations and increased bandwidth requirements.*
2. *Collision may not be detected because of the hidden station problem.*
3. *The distance between stations can be great. Signal fading could prevent a station at one end from hearing a collision at the other end.*

Frame Exchange Timeline: *The data and control frames exchange in time.*

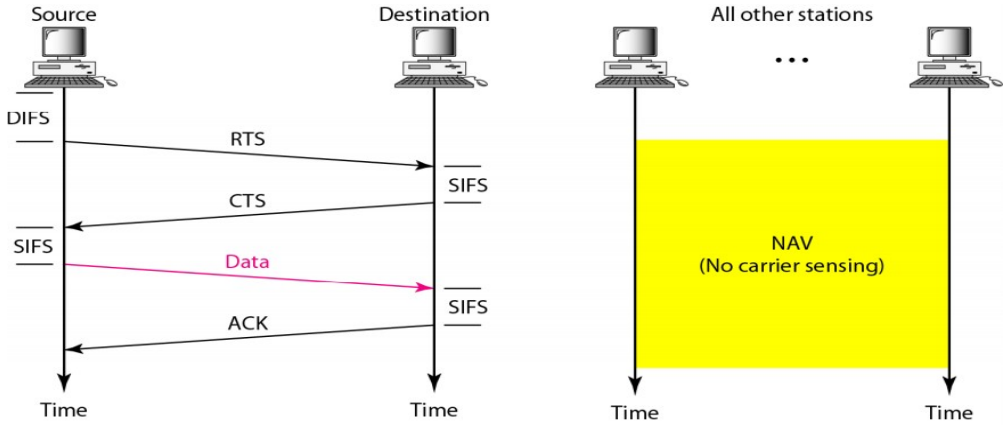

1. *Before sending a frame, the source station senses the medium by checking the energy level at the carrier frequency.*

a. The channel uses a persistence strategy with back-off until the channel is idle.

b. After the station is found to be idle, the station waits for a period of time called the distributed interframe space (DIFS); then the station sends a control frame called the request to send (RTS).

2. After receiving the RTS and waiting a period of time called the short interframe space (SIFS), the destination station sends a control frame, called the clear to send (CTS), to the source station. This control frame indicates that the destination station is ready to receive data.
3. The source station sends data after waiting an amount of time equal to SIFS.

4. The destination station, after waiting an amount of time equal to SIFS, sends an acknowledgment to show that the frame has been received. Acknowledgment is needed in this protocol because the station does not have any means to check for the successful arrival of its data at the destination.

Network Allocation Vector:

- When a station sends an RTS frame, it includes the duration of time that it needs to occupy the channel. The stations that are affected by this transmission create a timer called a **network allocation vector** (NAV) that shows how much time must pass before these stations are allowed to check the channel for idleness.
- Each time a station accesses the system and sends an RTS frame, other stations start their NAV.

Process Flowchart

- The process flowchart for CSMA/CA as used in wireless LANs

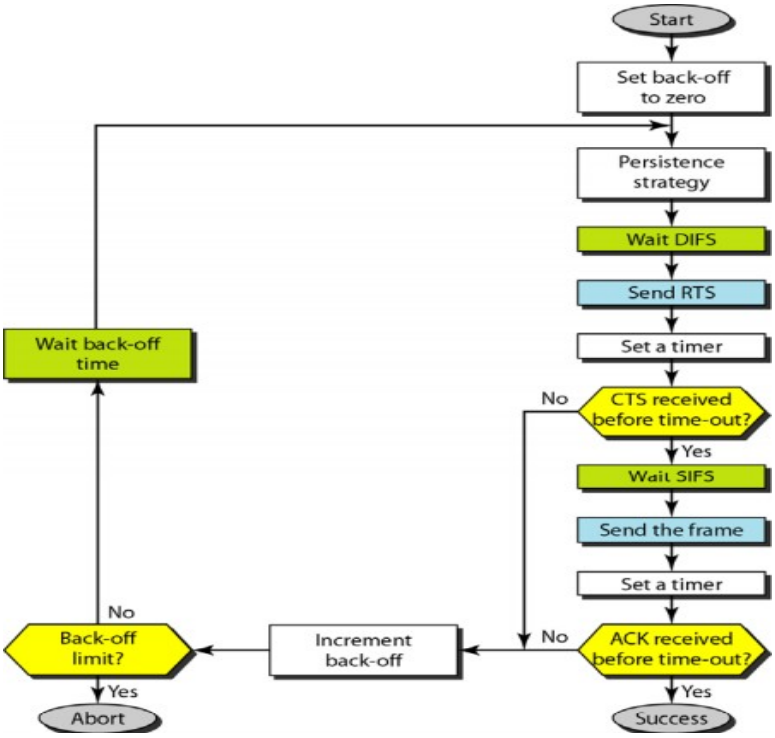

Collision During Handshaking

- Two or more stations may try to send RTS frames at the same time. These control frames may collide. However, because there is no mechanism for collision detection, the sender assumes there has been a collision if it has not received a CTS frame from the receiver. The back-off strategy is employed, and the sender tries again.

Point Coordination Function (PCF)

- The point coordination function (PCF) is an optional access method that can be implemented in **an infrastructure network**. It is implemented for **time sensitive transmission**.
- PCF has a **centralized, contention-free polling access method**. The AP performs polling for stations that are capable of being polled. The stations are polled one after another, sending any data they have to the AP.
- To give **priority to PCF over DCF**, another set of inter frame spaces has been defined: PIFS and SIFS. The SIFS is the same as that in DCF, but the PIFS (PCF IFS) is shorter than the DIFS. This means that if, at the same time, a station wants to use only DCF and an AP wants to use PCF, the AP has priority.

- Due to the priority of PCF over DCF, stations that only use DCF may not gain access to the medium. To prevent this, a repetition interval has been designed to cover both contention-free (PCF) and contention-based (DCF) traffic. The repetition interval, which is repeated continuously, starts with a special control frame, called a **beacon frame**. When the stations hear the beacon frame, they start their NAV for the duration of the contention-free period of the repetition interval. Figure 14. 6 shows an example of a repetition interval.

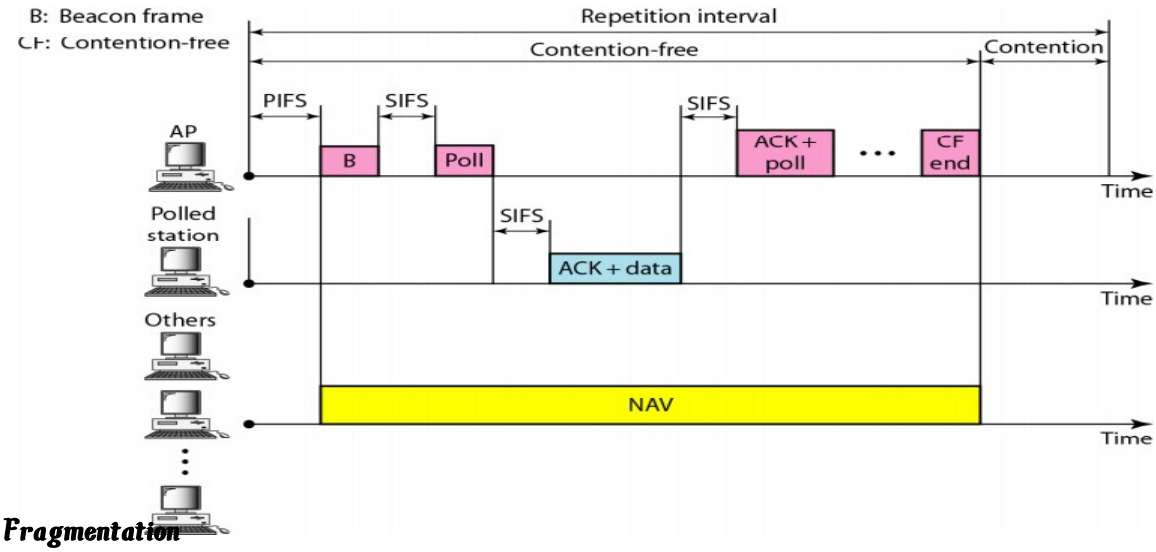

Fragmentation

- The wireless environment is very noisy; a corrupt frame has to be retransmitted. The protocol, therefore, recommends fragmentation-the division of a large frame into smaller ones. It is more efficient to resend a small frame than a large one.

Frame Format

The MAC layer frame consists of nine fields.

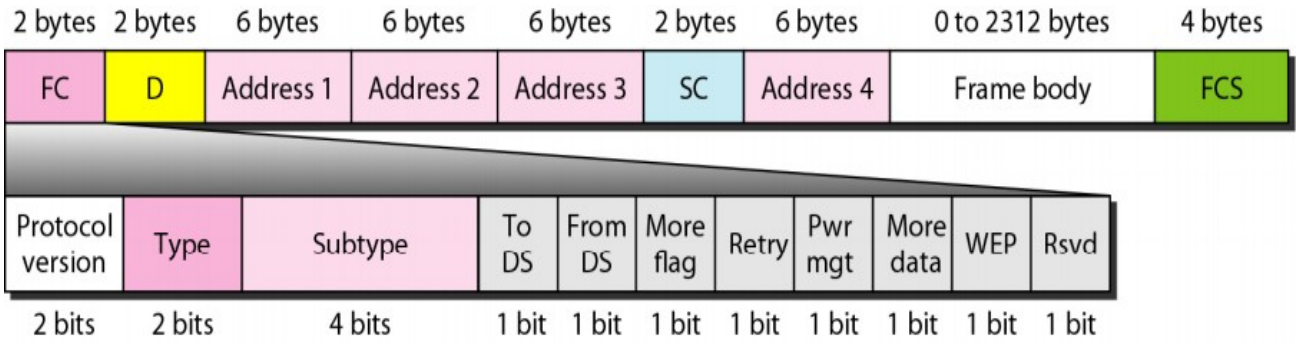

- ***Frame control field:*** *The FC field is 2 bytes long and defines the type of frame and some control information.*
- ***Duration:*** *In all frame types except one, this field defines the duration of the transmission that is used to set the value of NAV. In one control frame, this field defines the ID of the frame.*
- ***Addresses.*** *There are four address fields, each 6 bytes long. The meaning of each address field depends on the value of the To DS and From DS subfields.*
- ***Sequence control.*** *This field defines the sequence number of the frame to be used in flow control.*
- ***Frame body.*** *This field, which can be between 0 and 2312 bytes, contains information based on the type and the subtype defined in the FC field.*
- ***FCS.*** *The FCS field is 4 bytes long and contains a CRC-32 error detection sequence.*

Frame Types

- *A wireless LAN defined by IEEE 802.11 has three categories of frames:*

Management frames, control frames, and data frames.

Management Frames: *Management frames are used for the initial communication between stations and access points.*

Control Frames: *Control frames are used for accessing the channel and acknowledging frames.*

Subtype	*Meaning*
1011	Request to send (RTS)
1100	Clear to send (CTS)
1101	Acknowledgment (ACK)

Data Frames: Data frames are used for carrying data and control information.

Addressing Mechanism:

- The IEEE 802.11 addressing mechanism specifies four cases, defined by the value of the two flags in the FC field, To DS and From DS. Each flag can be either 0 or 1, resulting in four different situations.
- The interpretation of the four addresses (address 1 to address 4) in the MAC frame depends on the value of these flags.

To DS	From DS	Address 1	Address 2	Address 3	Address 4
0	0	Destination	Source	BSS ID	N/A
0	1	Destination	Sending AP	Source	N/A
1	0	Receiving AP	Source	Destination	N/A
1	1	Receiving AP	Sending AP	Destination	Source

Case 1: 00 To DS = 0 and From DS = 0.

This means that the frame is not going to a distribution system (To DS = 0) and is not coming from a distribution system (From DS =0). The frame is going from one station in a BSS to another without passing through the distribution system. The ACK frame should be sent to the original sender.

Case 2: 01 To DS = 0 and From DS = 1.

This means that the frame is coming from a distribution system (From DS = 1). The frame is coming from an AP and going to a station. The ACK should be sent to the AP.

Case 3: 10 To DS =1 and From DS =0.

This means that the frame is going to a distribution system (To DS = 1). The frame is going from a station to an AP. The ACK is sent to the original station.

Case 4:11 To DS =1 and From DS =1.

The frame is going from one AP to another AP in a wireless distribution system. Here, we need four addresses to define the original sender, the final destination, and two intermediate APs.

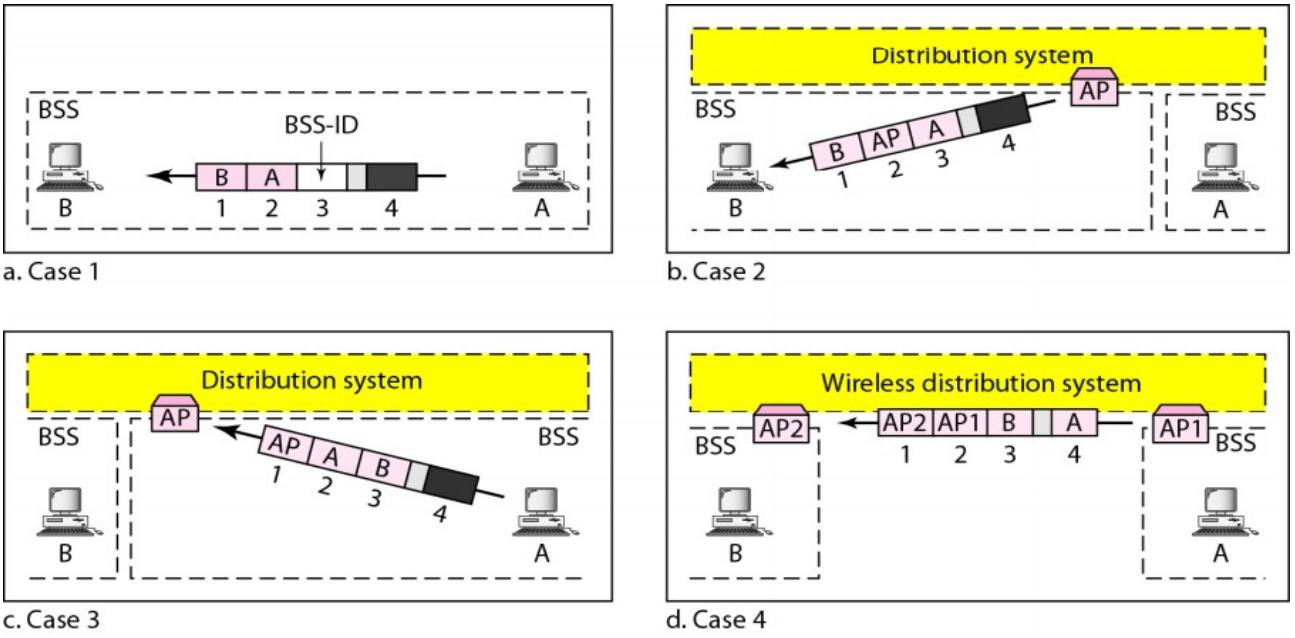

Hidden and Exposed Station Problems

- **Hidden Station Problem**: Station B has a transmission range shown by the left oval; every station in this range can hear any signal transmitted by station B. Station C has a transmission range shown by the right oval (sphere in space); every station located in this range can hear any signal transmitted by C.
- Station C is outside the transmission range of B; likewise, station B is outside the transmission range of C. Station A, however, is in the area covered by both Band C; it can hear any signal transmitted by B or C.

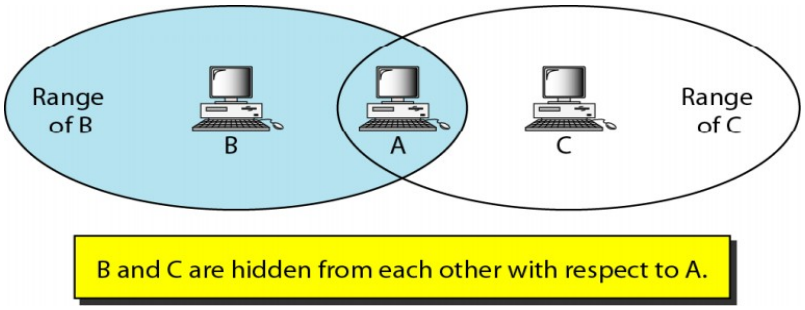

VIEW

- *Assume that station B is sending data to station A. In the middle of this transmission, station C also has data to send to station A. However, station C is out of B's range and transmissions from B cannot reach C. Therefore C thinks the medium is free.*
- *Station C sends its data to A, which results in a collision at A because this station is receiving data from both B and C. In this case, we say that stations Band C are hidden from each other with respect to A. Hidden stations can reduce the capacity of the network because of the possibility of collision.*
- ***The solution to the hidden station problem** is the use of the **handshake frames** (RTS and CTS). the RTS message from B reaches A, but not C. However, because both Band C are within the range of A, the CTS message, which contains the duration of data transmission from B to A reaches C. Station C knows that some hidden station is using the channel and refrains from transmitting until that duration is over.*

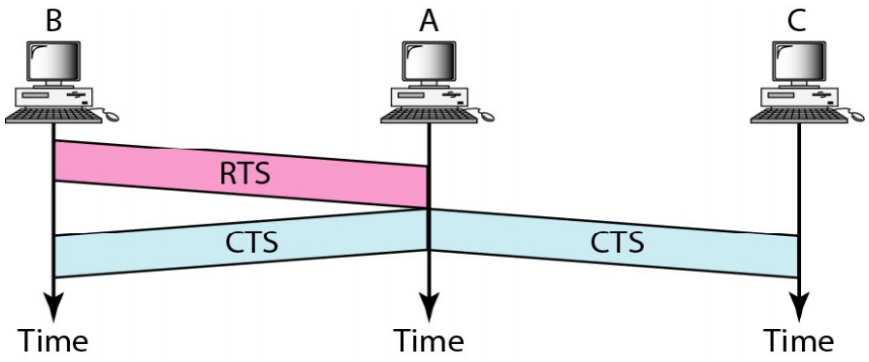

handshake frames

Exposed Station Problem

- *In this problem a station refrains from using a channel when it is, in fact, available.*
- *Station A is transmitting to station B.*
- *Station C has some data to send to station D, which can be sent without interfering with the transmission from A to B. However, station C is exposed to transmission from A; it hears what A is sending and thus refrains from sending. In other words, C is too conservative and wastes the capacity of the channel.*

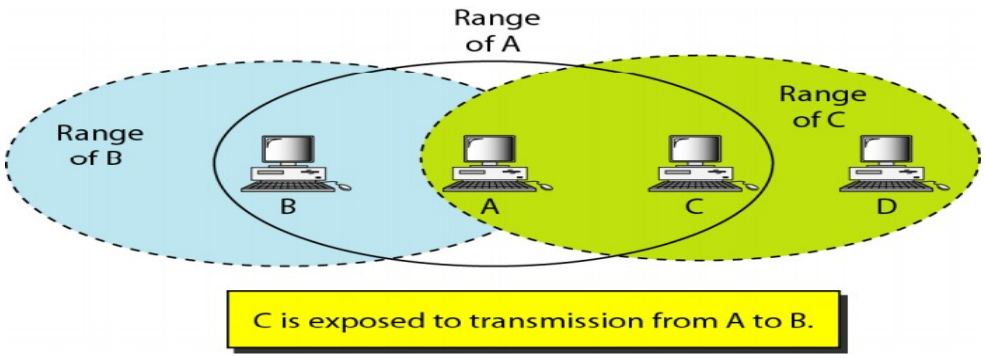

- *The handshaking messages RTS and CTS cannot help in this case, despite what you might think. Station C Station C hears the RTS from A, but does not hear the CTS from B. Station C, after hearing the RTS from A, can wait for a time so that the CTS from B reaches A; it then sends an RTS to D to show that it needs to communicate with D. Both stations B and A may hear this RTS, but station A is in the sending state, not the receiving state. Station B, however, responds with a CTS. The problem is here. If station A has started sending its data, station C cannot hear the CTS from station D because of the collision; it cannot send its data to D. It remains exposed until A finishes sending its data.*