CORE THEORY PAPER – 7

DATA COMMUNICATION & NETWORKS

Objective:
To equip students to basics of Data Communication and prepare them for better computer networking

UNIT II
Data Link Layer - Design issues - Channel allocation problem - Multiple access protocols - Ethernet - Wireless LAN - 802.11 architecture.

2.1 Introduction

The Data Link Layer is the second layer in the OSI model, above the Physical Layer, which ensures that the error free data is transferred between the adjacent nodes in the network. It breaks the datagram passed down by above layers and converts them into frames ready for transfer. This is called **Framing** .

It provides two main functionalities

? Reliable data transfer service between two peer network layers
? Flow Control mechanism which regulates the flow of frames such that data congestion is not there at slow receivers due to fast senders.

2.2 The Data Link Layer Design issues

Functions

? Providing a well-defined service interface to the network layer.
? Dealing with transmission errors.
? Regulating the flow of data so that slow receivers are not swamped by fast senders –flow control.

The two main functions of the data link layer are:

1. Data Link Control (DLC): It deals with the design and procedures for communication b/w node: Node-to-node communication.
2. Media Access Control (MAC): It explains how to share the link.

Data Link Control (DLC):

Data link control functions includes:

? Framing

? Error Control

? Flow Control

The Frame contains
1). Frame header
2). Payload field for holding packet
3). Frame trailer

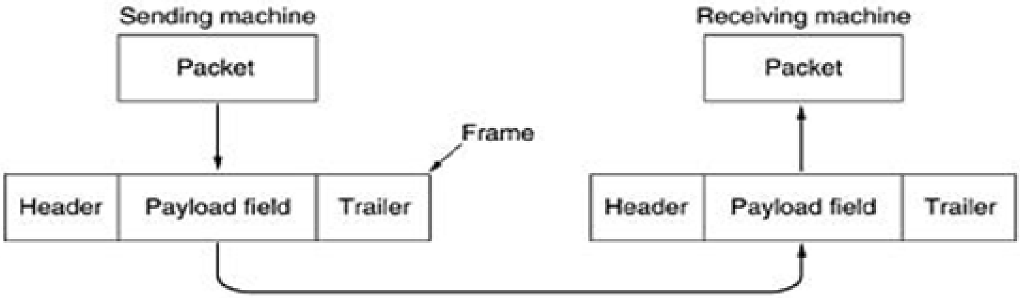

Fig 3.1 Relationships between Packets and Frames

2.2.1 Services provided to the network layer

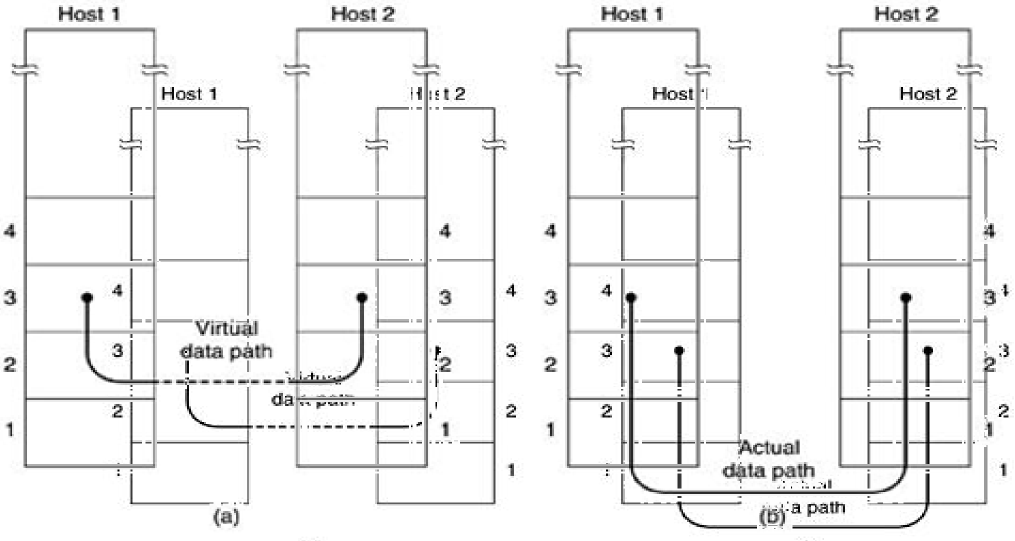

Figure 3.2 (a) Virtual communication. (b) Actual communication.

Transfering data from the network layer on the source machine to the network layer on the destination machine. The data link layer can be designed to offer various services. The actual services offered can vary from system to system. Three reasonable possibilities that are commonly provided are

1. **Unacknowledged connectionless service**

? Source machine sends independent frames to destination machine having destination machine acknowledge them.

? No logical connection.

? Used when error rate is very low.

? Good for real-time traffic (voice)

2. **Acknowledged connectionless service**

   ? No logical connection

   ? Each frame sent is individually acknowledged

   ? Useful over unreliable channels (i.e. wireless systems)

3. **Acknowledged connection-oriented service**

   ? Source and destination machines establish a connection before any data are transferred

   ? Each frame is numbered

   ? DLL guarantees that...

   ? Each frame is received
   ? Each frame is received exactly once
   ? Each frame is received in the right order

3 PHASES

When connection-oriented service is used, transfers go through three distinct phases

1. Connection established
2. Frames are transmitted
3. Connection released

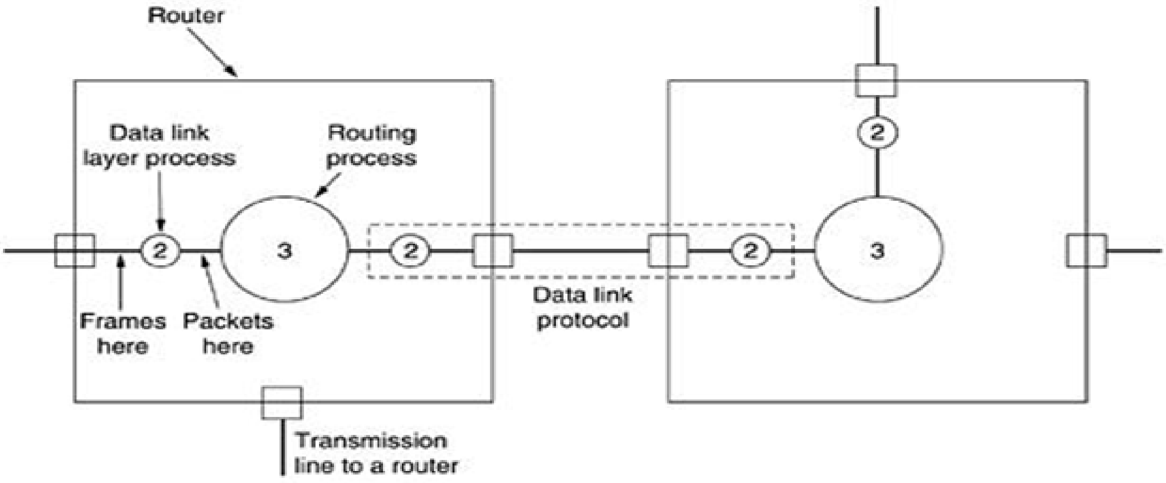

Figure 3.3 Placement of the data link protocol

? Consider a typical example: a WAN subnet consisting of routers connected by point-to-point leased telephone lines.
? When a frame arrives at a router, the hardware checks it for errors, and then passes the frame to the data link layer software.
? The data link layer software checks to see if this is the frame expected, and if so, gives the packet contained in the payload field to the routing software.
? The routing software then chooses the appropriate outgoing line and passes the packet back down to the data link layer software, which then transmits it. The flow over two routers is shown in Fig. 3.3.

2.2.2 Framing

Breaking the bit stream up into frames is more difficult than it at first appears. One way to achieve this framing is to insert time gaps between frames, much like the spaces between words in ordinary text. However, networks rarely make any guarantees about timing, so it is possible these gaps might be squeezed out or other gaps might be inserted during transmission.

There are four methods:

1. Character count.

2. Flag bytes with byte stuffing.

3. Starting and ending flags, with bit stuffing.

4. Physical layer coding violations.

Character count:

The first framing method uses a field in the header to specify the number of characters in the frame. When the data link layer at the destination sees the character count, it knows how many characters follow and hence where the end of the frame is. This technique is shown in Fig. 3-4(a) for four frames of sizes 5, 5, 8, and 8 characters, respectively.

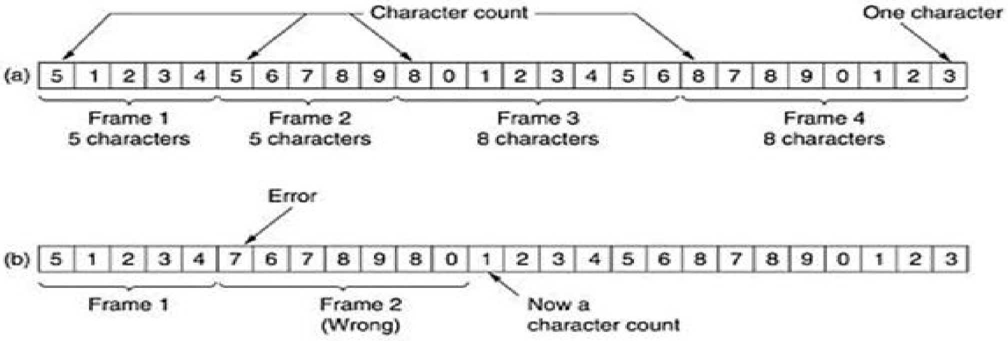

Figure 3-4. A character stream. (a) Without errors. (b) With one error.

Explanation (Figure 3-4.(a) A character stream Without errors.)

? The first framing method uses a field in the header to specify the number of characters in the frame.

? When the data link layer at the destination sees the character count, it knows how many characters follow and hence where the end of the frame.

? This technique is shown in Fig. 3-4(a) for four frames of sizes 5, 5, 8, and 8 characters, respectively.

? The trouble with this algorithm is that the count can be garbled by a transmission error.

Explanation (Figure 3-4.(b) A character stream with errors.)

? For example, if the character count of 5 in the second frame of Fig. 3-4(b) becomes a 7, the destination will get out of synchronization and will be unable to locate the start of the next frame.

? Even if the checksum is incorrect so the destination knows that the frame is bad, it still has no way of telling where the next frame starts.

? Sending a frame back to the source asking for a retransmission does not help either, since the destination does not know how many characters to skip over to get to the start of the retransmission. For this reason, the character count method is rarely used anymore.

Flag bytes with byte stuffing:

Character-oriented framing approach

- ✍ In a character-oriented approach, data to be carried are 8-bit characters.
- ✍ The header, which normally carries the source and destination addresses and other control information.
- ✍ Trailer carries error detection or error correction redundant bits, are also multiples of 8 bits.
- ✍ To separate one frame from the next, an 8-bit (1-byte) flag is added at the beginning and the end of a frame.
- ✍ The flag, composed of protocol-dependent special characters, signals the start or end of a frame.

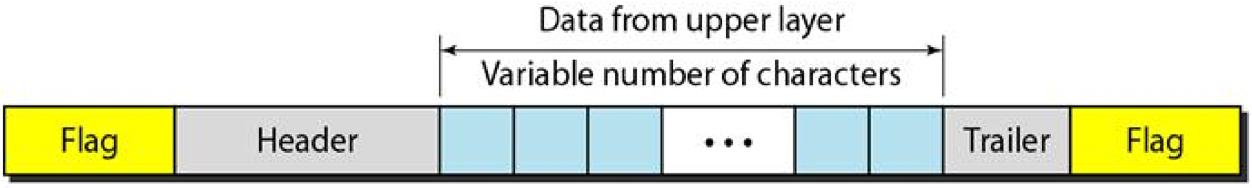

Figure: shows the format of a frame in a character-oriented protocol

Advantage:

1. Simple framing method.
2. Character-oriented framing was popular when only text was exchanged by the data Link layers.
3. The flag could be selected to be any character not used for text communication.

Disadvantage:

1. Even if with checksum, the receiver knows that the frame is bad there is no way to tell where the next frame starts.
2. Asking for retransmission doesn't help either because the start of the retransmitted frame is not known.
3. Hence No longer used.

Starting and ending character with byte stuffing

Byte stuffing is the process of adding 1 extra byte whenever there is a flag or escape character in the text.

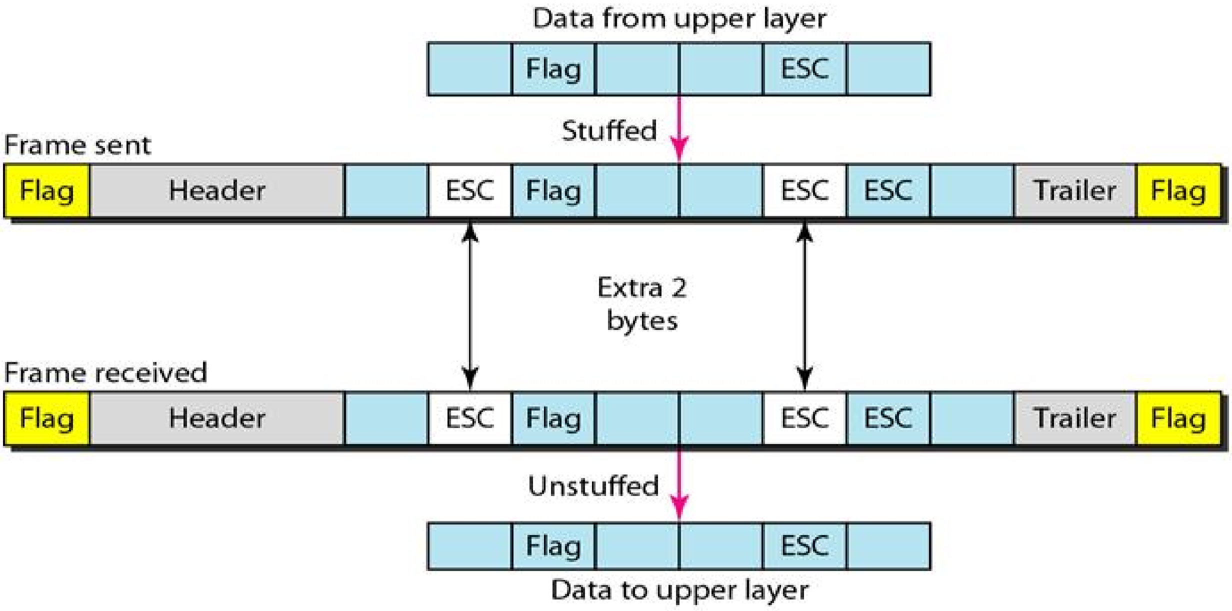

Figure : Byte stuffing and unstuffing

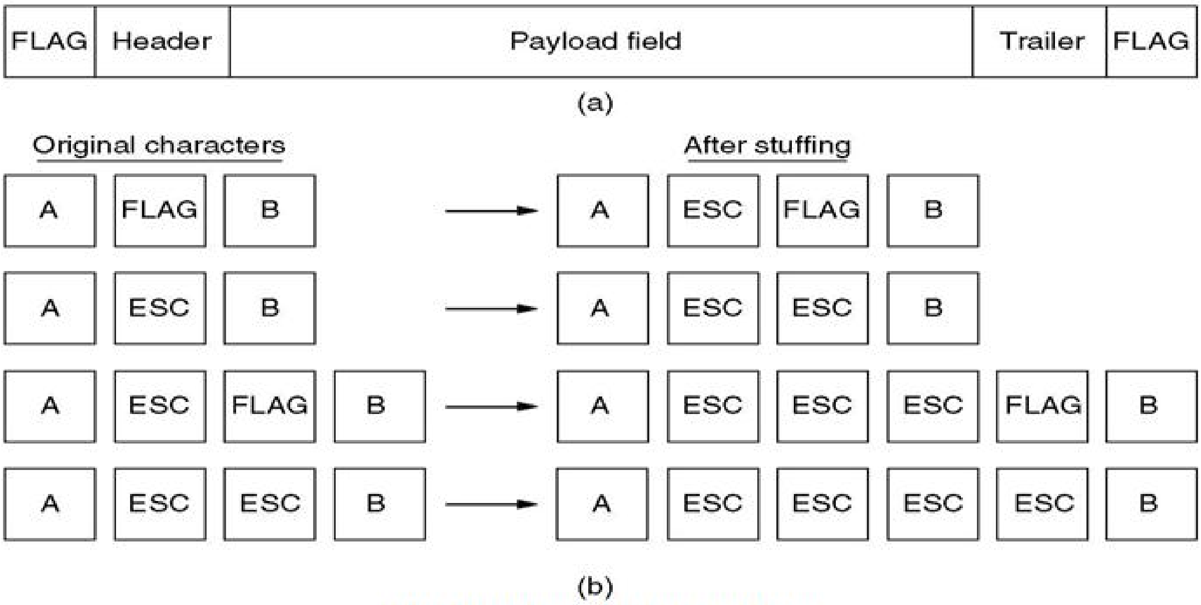

Fig: Framing with byte stuffing

Problem: fixed character size: assumes character size to be 8 bits: can't handle heterogeneous environment.

Bit-Oriented framing approach

? Bit stuffing is the process of adding one extra 0 whenever five consecutive 1's follow a 0 in the data, so that the receiver does not mistake the pattern 0111110 for a flag.

? Most protocols use a special 8-bit pattern flag 01111110 as the delimiter to define the beginning and the end of the frame, as shown in Figure below.

? This flag can create the same type of problem. That is, if the flag pattern appears in the data, we need to somehow inform the receiver that this is not the end of the frame.

? We do this by stuffing 1 single bit (instead of I byte) to prevent the pattern from looking like a flag. The strategy is called bit stuffing.

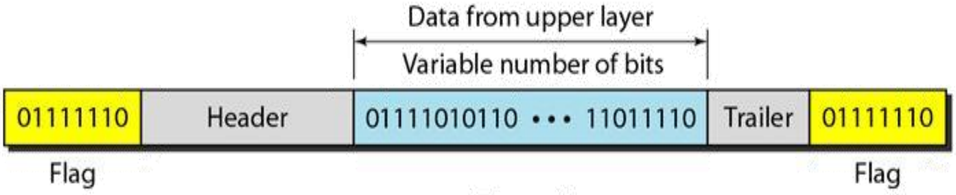

Figure (a)

Bit stuffing is the process of adding one extra 0 whenever five consecutive 1s follow a 0 in the data, so that the receiver does not mistake the pattern 0111110 for a flag.

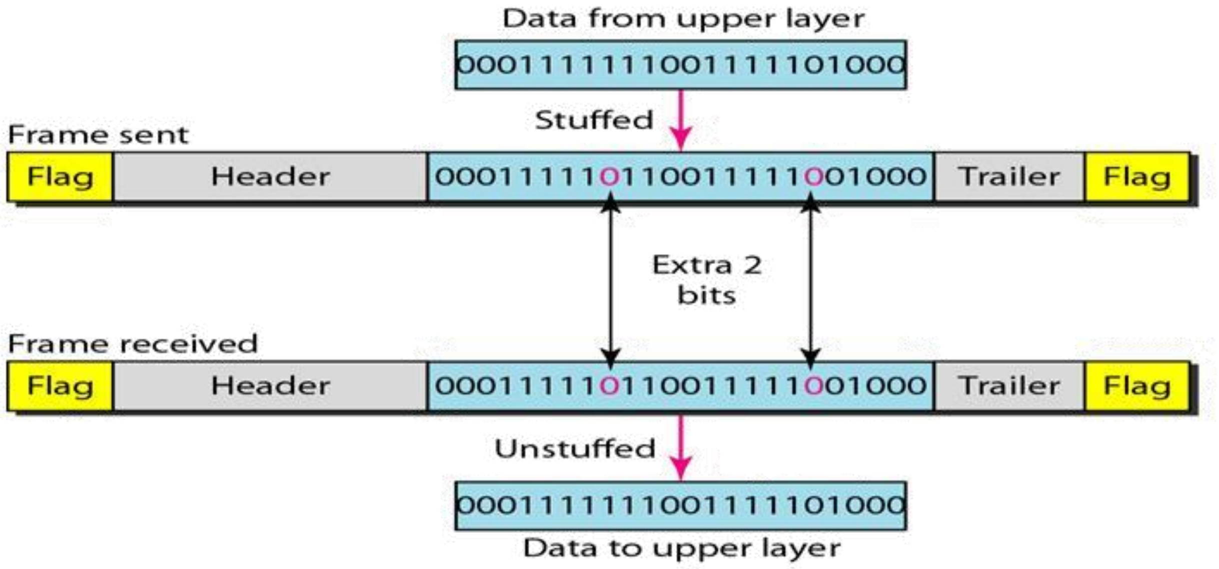

Figure (b)

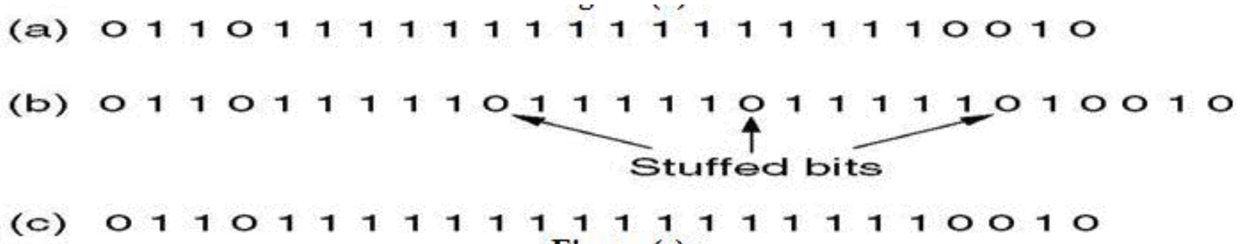
(a) 0 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 1 0
(b) 0 1 1 0 1 1 1 1 1 0 1 1 1 1 1 0 1 1 1 1 1 0 1 0 0 1 0
Stuffed bits
(c) 0 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 1 0

Figure (c)

(a) The original data.
(b) The data as they appear on the line.
(c) The data as they are stored in receiver's memory after destuffing.

Physical layer coding violation:

The last method of framing is only applicable to networks in which the encoding on the physical medium contains some redundancy

. For example, some LANs encode 1 bit of data by using 2 physical bits. Normally, a 1 bit is a high- low pair and a 0 bit is a low-high pair. The scheme means that every data bit has a transition in the middle, making it easy for the receiver to locate the bit boundaries. The combinations high-high and low- low are not used for data but are used for delimiting frames in some protocols.

As a final note on framing, many data link protocols use a combination of a character count with one of the other methods for extra safety. When a frame arrives, the count field is used to locate the end of the frame. Only if the appropriate delimiter is present at that position and the checksum is correct is the frame accepted as valid. Otherwise, the input stream is scanned for the next delimiter.

2.2.3 Error Control

How do we make sure that all frames are eventually delivered to the network layer at the destination and in the proper order?

? Provide sender with some acknowledgement about what is happening with the receiver

? Sender could wait for acknowledgement

Disadvantages

? If a frame vanishes, the receiver will not send an acknowledgement thus, sender will wait forever

? Dealt with by timers and sequence numbers – important part of DLL

? Sender transmits a frame, starts a timer.

? Timer set to expire after interval long enough for frame to reach destination, be processed, and have acknowledgement sent to sender.

? Is a danger of frame being transmitted several times, however dealt with by assigning sequence numbers to outgoing frames, so that receiver can distinguish retransmissions from originals.

2.2.4 Flow Control

What do we do when a sender transmits frames faster than the receiver can accept them?

? **Feedback-based flow control** – receiver sends back information to the sender, giving it permission to send more data or at least telling the sender how the receiver is doing

? **Rate-based flow control** – the protocol has a built-in mechanism that limits the rate at which the sender may transmit data, using feedback from the receiver.

2.3 The Channel Allocation Problem

The channel allocation problem is how to allocate a single broadcast channel among competing users.

2.3.1 Static Channel Allocation in LAN and MAN

? The traditional way of allocating a single channel, such as a telephone trunk, among multiple competing users is Frequency Division Multiplexing (FDM).

? If there are N users, the bandwidth is divided into N equal-sized portions, each user being assigned one portion.Since each user has a private frequency band, there is no interference between users.

? When there is only a small and constant number of a user, each of which has a heavy (buffered) load of traffic (e.g., carriers' switching offices), FDM is a simple and efficient allocation mechanism.

? However, when the number of senders is large and continuously varying or the traffic is bursty, FDM presents some problems.

? If the spectrum is cut up into N regions and fewer than N users are currently interested in communicating, a large piece of valuable spectrum will be wasted.

? If more than N users want to communicate, some of them will be denied permission for lack of bandwidth, even if some of the users who have been assigned a frequency band hardly ever transmit or receive anything.

? However, even assuming that the number of users could somehow be held constant at N, dividing the single available channel into static sub channels is inherently inefficient.

? 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.

? Furthermore, in most computer systems, data traffic is extremely bursty (peak traffic to mean traffic ratios of 1000:1 are common). Consequently, most of the channels will be idle most of the time.

? The poor performance of static FDM can easily be seen from a simple queuing theory calculation. Let us start with the mean time delay, T, for a channel of capacity C bps, with an arrival rate of λ frames/sec, each frame having a length drawn from an exponential probability density function with mean 1/μ bits/frame. With these parameters the arrival rate is λ frames/sec and the service rate is μ C frames/sec. From queuing theory it can be shown that for Poisson arrival and service times,

$$T = \frac{1}{\mu C - \lambda}$$

? For example, if C is 100 Mbps, the mean frame length, 1/μ, is 10,000 bits, and the frame arrival rate, λ, is 5000 frames/sec, then T = 200 μ sec. Note that if we ignored the queuing delay and just asked how long it takes to send a 10,000 bit frame on a 100-Mbps network, we would get the (incorrect) answer of 100 μ sec. That result only holds when there is no contention for the channel.

? Now let us divide the single channel into N independent sub channels, each with capacity C/N bps. The mean input rate on each of the sub channels will now be λ/N. Recomputing T we get Equation 4.

? The mean delay using FDM is N times worse than if all the frames were somehow magically arranged orderly in a big central queue.

$$T_{\text{FDM}} = \frac{1}{\mu(C/N) - (\lambda/N)} = \frac{N}{\mu C - \lambda} = NT$$

? Precisely the same arguments that apply to FDM also apply to time division multiplexing (TDM). Each user is statically allocated every Nth time slot. If a user does not use the allocated slot, it just lies fallow. The same holds if we split up the networks physically. Using our previous example again, if we were to replace the 100-Mbps network with 10 networks of 10 Mbps each and statically allocate each user to one of them, the mean delay would jump from 200 μ sec to 2 msec.

? Since none of the traditional static channel allocation methods work well with bursty traffic, we will now explore dynamic methods.

2.3.2 Dynamic Channel Allocation in LAN and MAN

FIVE KEY ASSUMPTIONS

1.Station Model.

? The model consists of N independent stations (e.g., computers, telephones, or personal communicators), each with a program or user that generates frames for transmission. Stations are sometimes called terminals.

? The probability of a frame being generated in an interval of length Δt is λΔt, where λ 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 Assumption.

? A single channel is available for all communication. All stations can transmit on it and all can receive from it.

? As far as the hardware is concerned, all stations are equivalent, although protocol software may assign priorities to them.

3.Collision Assumption.

? 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 collisions. A collided frame must be transmitted again later. There are no errors other than those generated by collisions.

4a. Continuous Time.

Frame transmission can begin at any instant. There is no master clock dividing time into discrete intervals.

4b. Slotted Time.

? Time is divided into discrete intervals (slots). Frame transmissions always 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.

5a. Carrier Sense.

Stations can tell if the channel is in use before trying to use it. If the channel is sensed as busy, no station will attempt to use it until it goes idle.

5b. 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.

2.4 Multiple Access Protocol

? The media access control (MAC) data communication protocol sub-layer, also known as the **medium access control** , is a **sublayer of the data link layer** specified in the seven-layer OSI model

? It provides addressing and channel access control mechanisms that make it possible for several terminals or network nodes to communicate within a multiple access network that incorporates a shared medium, e.g, Ethernet. The hardware that implements the MAC is referred to as a medium access controller.

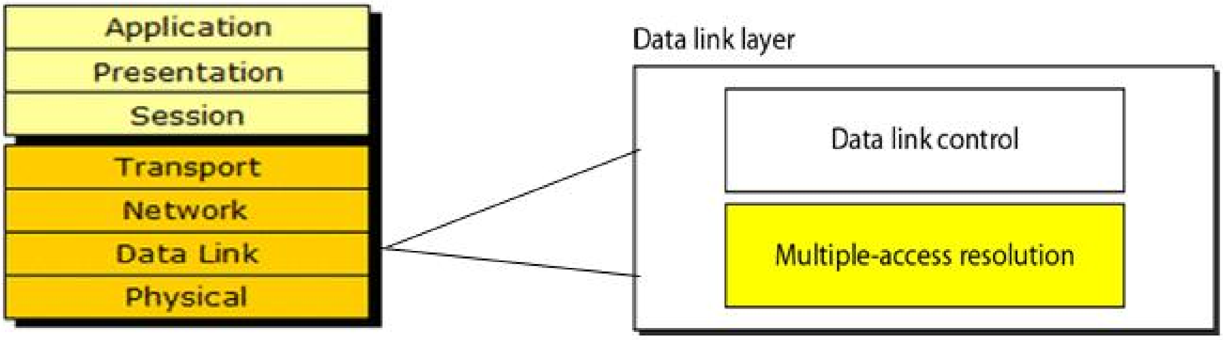

? The MAC sub-layer acts as an interface between the logical link control (LLC) sublayer and the network's physical layer. The MAC layer emulates a full-duplex logical communication channel in a multi-point network. This channel may provide unicast, multicast or broadcast communication service.

? The channel access control mechanisms provided by the MAC layer is known as a multiple access protocol. This makes it possible for several stations connected to the same physical medium to share it.

? Examples of shared physical media are bus networks, ring networks, hub networks, wireless networks and half- duplex point-to-point links.

? The multiple access protocol may detect or avoid data packet collisions if a packet mode contention based channel access method is used, or reserve resources to establish a logical channel if a circuit switched or channelization based channel access method is used. The channel access control mechanism relies on a physical layer multiplex scheme.

RANDOM ACCESS PROTOCOLS

? In a random access protocol, a transmitting node always transmits at the full rate of the channel, namely, R bps.

? When there is a collision, each node involved in the collision repeatedly retransmits its frame(that is ,packet) until the frame gets through without a collision. But when a node experiences a collision, it doesn’t necessarily retransmitting the frame right away. Instead it waits a random delay before retransmitting the frame.

? Each node involved in a collision chooses independent random delays .Because the random delays are independently chosen, it is possible that one of the nodes will pick a delay that is sufficiently less than the delays of the other colliding nodes and will therefore be able to sneak its frame into the channel without a collision.

The most commonly used random access protocols

1. The ALOHA protocol ,
2. CSMA (carrier sense multiple access) protocol ,
3. CSMA/CD (carrier sense multiple access /collision detection) protocol and
4. Collision-Free Protocols
5. Limited-Contention Protocols

2.4.1 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**.

? Although Abramson's work, called the ALOHA system, used ground-based radio broadcasting, the basic idea is applicable to any system in which uncoordinated users are competing for the use of a single shared channel.

? The two versions of ALOHA here: **pure and slotted** .

? They differ with respect to whether time is divided into discrete slots into which all frames must fit.

? Pure ALOHA does not require global time synchronization; slotted ALOHA does.

PURE ALOHA

? If you have data to send, send the data

? If the message collides with another transmission, try resending “later”

Note that the first step implies that Pure ALOHA does not check whether the channel is busy before transmitting. The critical aspect is the "later" concept: the quality of the backoff scheme chosen significantly influences the efficiency of the protocol, the ultimate channel capacity, and the predictability of its behavior.

A sketch of frame generation in an ALOHA system is given in Fig. 4-1. 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.

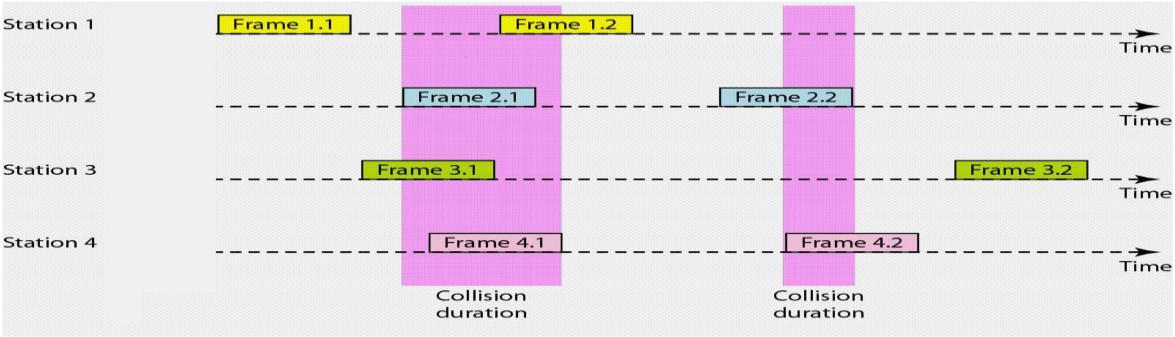

Figure 4-1. In pure ALOHA, frames are transmitted at completely arbitrary times.

To assess Pure ALOHA, we need to predict its throughput, the rate of (successful) transmission of frames. First, let's make a few simplifying assumptions:

? All frames have the same length.
? Stations cannot generate a frame while transmitting or trying to transmit.
? The population of stations attempts to transmit (both new frames and old frames that collided) according to a Poisson distribution.

Let "T" refer to the time needed to transmit one frame on the channel, and let's define "frame-time" as a unit of time equal to T. Let "G" refer to the mean used in the Poisson distribution over transmission-attempt amounts: that is, on average, there are G transmission-attempts per frame-time.

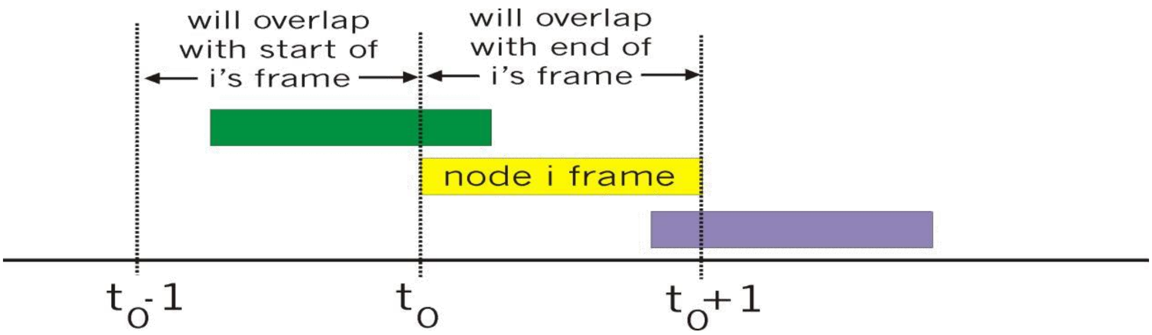

will overlap with start of i's frame
will overlap with end of i's frame
node i frame
t_o-1
t_o
t_o+1

Overlapping frames in the pure ALOHA protocol. Frame-time is equal to 1 for all frames.

Pure Aloha efficiency

P(success by given node) = P(node transmits) .

P(no other node transmits in [t 0-1,t0] .

P(no other node transmits in [t 0,t0+1]

= p . (1-p)N-1 . (1-p)N-1

= p . (1-p)2(N-1)

… choosing optimum p and then letting n -> ∞ ...

Efficiency = 1/(2e) = .18

SLOTTED ALOHA

? An improvement to the original ALOHA protocol was "Slotted ALOHA", which introduced discrete timeslots and increased the maximum throughput.

? A station can send only at the beginning of a timeslot, and thus collisions are reduced. In this case, we only need to worry about the transmission-attempts within 1 frame-time and not 2 consecutive frame-times, since collisions can only occur during each timeslot. Thus, the probability of there being zero transmission-attempts in a single timeslot is:

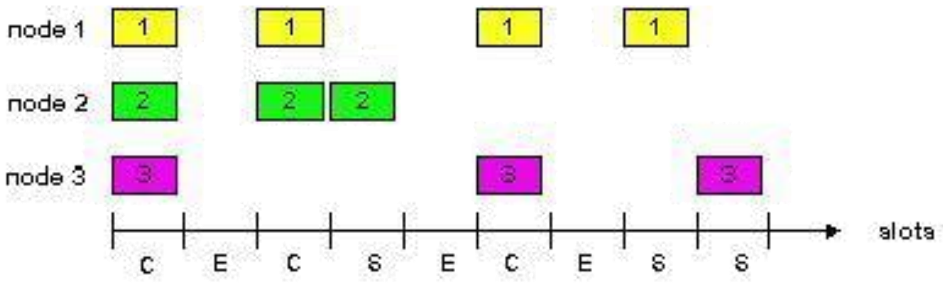

? Slotted ALOHA is used in low-data-rate tactical satellite communications networks by military forces, in subscriber-based satellite communications networks, mobile telephony call setup, and in the contactless RFID technologies.

Pros

? single active node can continuously transmit at full rate of channel
? highly decentralized: only slots in nodes need to be in sync
? simple

Cons

? collisions, wasting slots.
? idle slots

? nodes may be able to detect collision in less than time to transmit packet

? clock synchronization

Efficiency is the long-run fraction of successful slots when there are many nodes, each with many frames to send

- Suppose N nodes with many frames to send, each transmits in slot with probability p
- prob that node 1 has success in a slot = p(1-p)N-1
- prob that any node has a success = Np(1-p)N-1
- For max efficiency with N nodes, find p* that maximizes Np(1-p)N-1
- For many nodes, take limit of Np*(1-p*)N-1 as N goes to infinity, gives 1/e = .37
- Efficiency is 37%, even with optimal p

2.4.2 Carrier Sense Multiple Access

Carrier Sense Multiple Access (**CSMA**) is a probabilistic Media Access Control (MAC) protocol in which a node verifies the absence of other traffic before transmitting on a shared transmission medium, such as an electrical bus, or a band of the electromagnetic spectrum.

" **Carrier Sense**" describes the fact that a transmitter uses feedback from a receiver that detects a carrier wave before trying to send. That is, it tries to detect the presence of an encoded signal from another station before attempting to transmit. If a carrier is sensed, the station waits for the transmission in progress to finish before initiating its own transmission.

" **Multiple Access**" describes the fact that multiple stations send and receive on the medium. Transmissions by one node are generally received by all other stations using the medium.

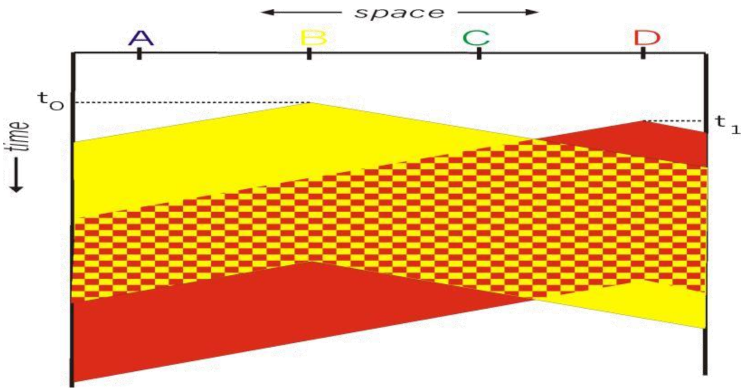

ADVANTAGES

? Fairly simple to implement

? Functional scheme that works

DISADVANTAGES

? Cannot recover from a collision (inefficient waste of medium time)

2.4.3 Collision Free Protocols

Although collisions do not occur with CSMA/CD once a station has unambigously seized the channel, they can still occur during the contention period. These collisions adversely affect the efficiency of transmission. Hence some protocols have been developed which are contention free.

Bit-Map Method

In this method, there N slots. If node 0 has a frame to send, it transmit a 1 bit during the first slot. No other node is allowed to transmit during this period. Next node 1 gets a chance to transmit 1 bit if it has something to send, regardless of what node 0 had transmitted. This is done for all the nodes. In general node j may declare the fact that it has a frsme to send by inserting a 1 into slot j. Hence after all nodes have passed, each node has complete knowledge of who wants to send a frame. Now they begin transmitting in numerical order. Since everyone knows who is transmitting and when, there could never be any collision. The basic problem with this protocol is its inefficiency during low load. If a node has to transmit and no other node needs to do so, even then it has to wait for the bitmap to finish. Hence the bitmap will be repeated over and over again if very few nodes want to send wasting valuable bandwidth.

Binary Countdown

In this protocol, a node which wants to signal that it has a frame to send does so by writing its address into the header as a binary number. The arbitration is such that as soon as a node sees that a higher bit position that is 0 in its address has been overwritten with a 1, it gives up. The final result is the address of the node which is allowed to send. After the node has transmitted the whole process is repeated all over again. Given below is an example situation.

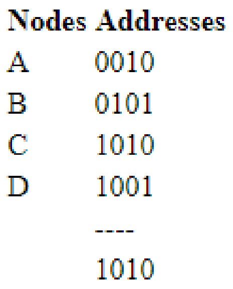

Nodes	Addresses
A	0010
B	0101
C	1010
D	1001

	1010

Node C having higher priority gets to transmit. The problem with this protocol is that the nodes with higher address always wins. Hence this creates a priority which is highly unfair and hence undesirable.

2.4.4 Limited Contention Protocols

Both the type of protocols described above - Contention based and Contention - free has their own problems. Under conditions of light load, contention is preferable due to its low delay. As the load increases, contention becomes increasingly less attractive, because the overload associated with channel arbitration becomes greater. Just the reverse is true for contention - free protocols. At low load, they have high delay, but as the load increases , the channel efficiency improves rather than getting worse as it does for contention protocols.

Obviously it would be better if one could combine the best properties of the contention and contention - free protocols, that is, protocol which used contention at low loads to provide low delay, but used a cotention-free technique at high load to provide good channel efficiency. Such protocols do exist and are called Limited contention protocols.

It is obvious that the probablity of some station aquiring the channel could only be increased by decreasing the amount of competition. The limited contention protocols do exactly that. They first divide the stations up into (not necessarily disjoint) groups. Only the members of group 0 are permitted to compete for slot 0. The competition for aquiring the slot within a group is contention based. If one of the members of that group succeeds, it aquires the channel and transmits a frame. If there is collision or no node of a particular group wants to send then the members of the next group compete for the next slot. The probablity of a particular node is set to a particular value (optimum).

Adaptive Tree Walk Protocol

The following is the method of adaptive tree protocol. Initially all the nodes are allowed to try to aquire the channel. If it is able to aquire the channel, it sends its frame. If there is collision then the nodes are divided into two equal groups and only one of these groups compete for slot 1. If one of its member aquires the channel then the next slot is reserved for the other group. On the other hand, if there is a collision then that group is again subdivided and the same process is followed. This can be better understood if the nodes are thought of as being organised in a binary tree as shown in the following figure.

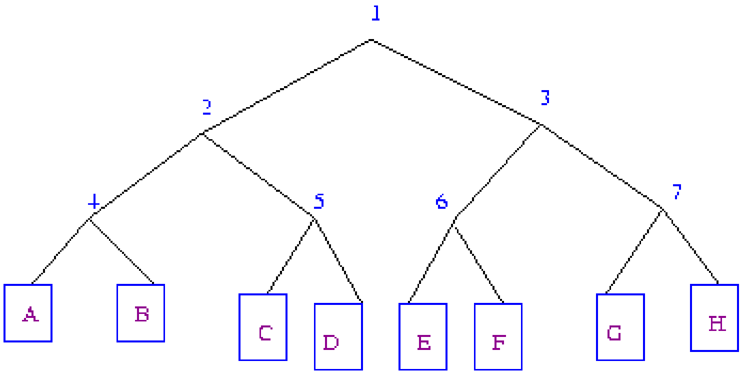

Fig. Adaptive Tree Walk abstraction of nodes in binary tree.

Many improvements could be made to the algorithm. For example, consider the case of nodes G and H being the only ones wanting to transmit. At slot 1 a collision will be detected and so 2 will be tried and it will be found to be idle. Hence it is pointless to probe 3 and one should directly go to 6,7.

2.5 Ethernet

IEEE 802.3 Local Area Network (LAN) Protocols : Ethernet protocols refer to the family of local- area network (LAN)covered by the IEEE 802.3. In the Ethernet standard, there are twomodes of operation: half-duplex and full-duplex modes. In the halfduplex mode, data are transmitted using the popular Carrier-SenseMultiple Access/Collision Detection (CSMA/CD) protocol on ashared medium.

The main disadvantages of the half-duplex are the efficiency and distance limitation, in which the link distance is limited by the minimum MAC frame size. This restriction reducesthe efficiency drastically for high-rate transmission. Therefore, the carrier extension technique is used to ensure the minimum framesize of 512 bytes in Gigabit Ethernet to achieve a reasonable link distance. Four data rates are currently defined for operation over optical fiber and twisted-pair cables :

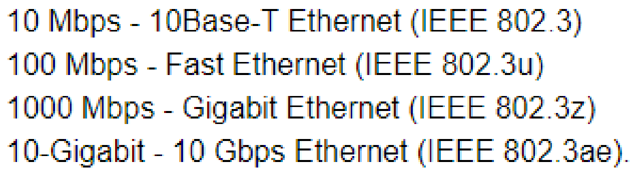

10 Mbps - 10Base-T Ethernet (IEEE 802.3)
100 Mbps - Fast Ethernet (IEEE 802.3u)
1000 Mbps - Gigabit Ethernet (IEEE 802.3z)
10-Gigabit - 10 Gbps Ethernet (IEEE 802.3ae).

2.5.1 Classic Ethernet 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.

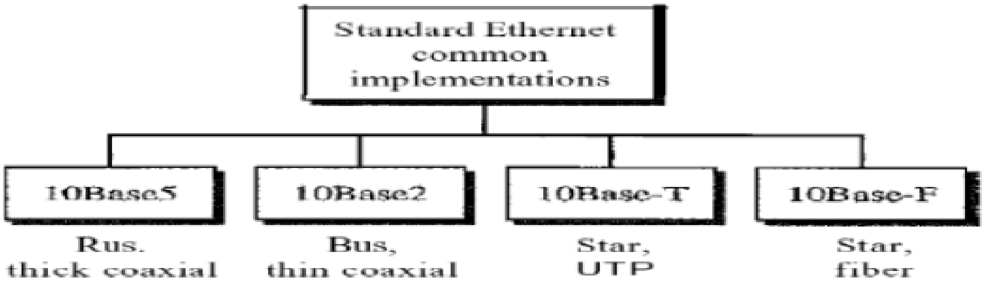

2.5.2 Classic Ethernet MAC Sublayer Protocol

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 CRe. 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 ofthe MAC frame is shown in Figure 4.

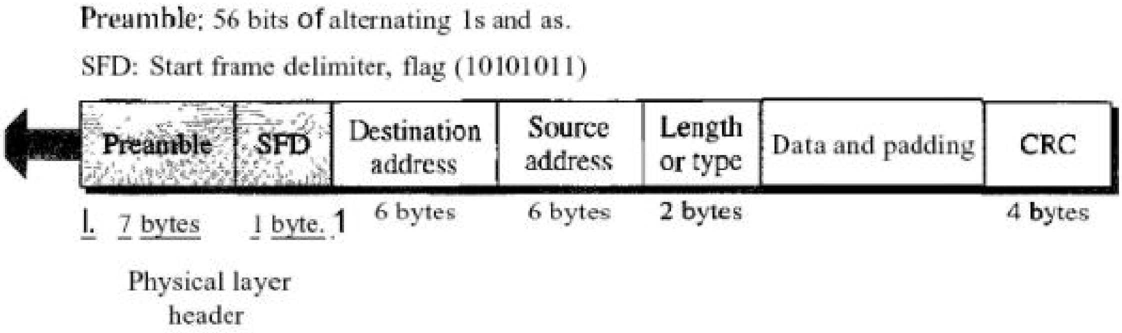

Figure 4 *802.3 MACframe*

D Preamble. The first field of the 802.3 frame contains 7 bytes (56 bits) of alternating Os and Is 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.

D Start frame delimiter (SFD). The second field (l byte: 10101011) signals the beginning of the frame.

The SFD warns the station or stations that this is the last chance for synchronization. The last 2 bits is 11 and alerts the receiver that the next field is the destination address.

Destination address (DA). The DA field is 6 bytes and contains the physical address of the destination station or stations to receive the packet.

Source address (SA). The SA field is also 6 bytes and contains the physical address of the sender of the packet. We will discuss addressing shortly.

Length or type . This field is defined as a type field or length field. The original Ethernet used this field

as the type field to define the upper-layer protocol using the MAC frame. The IEEE standard used it as the length field to define the number of bytes in the data field. Both uses are common today.

Data. This field carries data encapsulated from the upper-layer protocols. It is a minimum of 46 and a maximum of 1500 bytes.

CRC. The last field contains error detection information, in this case a CRC-32.

2.5.3 Fast Ethernet

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:

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

2.5.4 Gigabit Ethernet

Gigabit Ethernet works much the same way as 10 Mb/s and 100 Mb/s Ethernet, only faster. It uses the same IEEE 802.3 frame format, full duplex, and flow control methods. Additionally, it takes advantage of CSMA/CD when in half-duplex mode, and it supports simple network management protocol (SNMP) tools.

Gigabit Ethernet takes advantage of jumbo frames to reduce the frame rate to the end host. Standard Ethernet frame sizes are between 64 and 1518 bytes. Jumbo frames are between 64 and 9215 bytes. Because larger frames translate to lower frame rates, using jumbo frames on Gigabit Ethernet links greatly reduces the number of packets (from more than 80,000 to less than 15,000 per second) that are received and processed by the end host.

Gigabit Ethernet can be transmitted over CAT 5 cable and optical fiber such as the following:

- 1000Base-CX—Short distance transport (copper)
- 1000Base-SX—850 nm wavelength (fiber optics)
- 1000Base-LX—1300 nm wavelength (fiber optics)

2.5.5 10 Gigabit Ethernet

The operation of 10 Gigabit Ethernet is similar to that of lower speed Ethernets. It maintains the IEEE 802.3 Ethernet frame size and format that preserves layer 3 and greater protocols. However, 10 Gigabit Ethernet only operates over point-to-point links in full-duplex mode. Additionally, it uses only multimode and single mode optical fiber for transporting Ethernet frames.

Note: Operation in full-duplex mode eliminates the need for CSMA/CD.

The 10 Gigabit Ethernet standard (IEEE 802.3ae) defines two broad physical layer network applications:

? Local area network (LAN) PHY
? Wide area network (WAN) PHY

LAN PHY The LAN PHY operates at close to the 10 Gigabit Ethernet rate to maximize throughput over short distances.

WAN PHY The WAN PHY supports connections to circuit-switched SONET networks.

2.6 Wireless LAN

? Wireless LANs are increasingly popular, and homes, offices, cafes, libraries, airports, zoos, and other public places are being outfitted with them to connect computers, PDAs, and smart phones to the Internet.
? Wireless LANs can also be used to let two or more nearby computers communicate without using the Internet.
? The main wireless LAN standard is 802.11.

2.6.1 The 802.11 Architecture and Protocol Stack

The standard defines two kinds of services: the basic service set (BSS) and the extended service set (ESS).

Basic Service Set

IEEE 802.11 defines the basic service set (BSS) as the building block of a wireless LAN. A basic service set is made of stationary or mobile wireless stations and an optional central base station, known as the access point (AP). Figure 9 shows two sets in this standard.

The BSS without an AP is a stand-alone network and cannot send data to other BSSs. It is called an ad hoc architecture. In this 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.

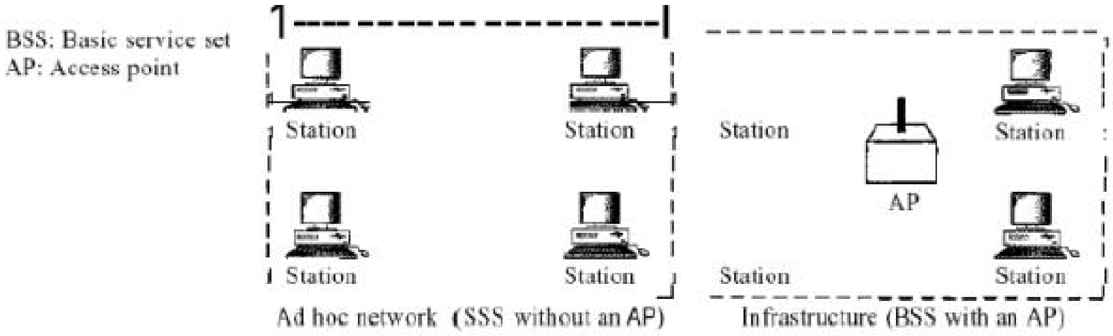

Figure. 9 *Basic service sets (BSSs)*

Extended Service Set

An extended service set (ESS) is made up of two or more BSSs with APs. In this case, the BSSs are connected through a distribution system, which is usually a wired LAN. The distribution system connects the APs in the BSSs. IEEE 802.11 does not restrict the distribution system; it can be any IEEE LAN such as an Ethernet. Note that 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. Figure 10 shows an ESS.

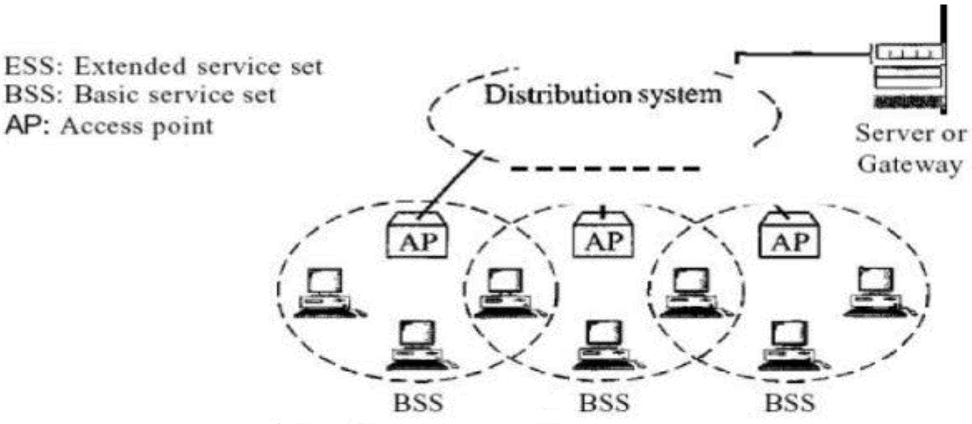

Figure 10 *Extended service sets (ESSs)*