UNIT-IV

MEDIUM ACCESS CONTROL SUBLAYER (MAC)

Networks can be categories in to two ways

a) Point to point b) Broad cast channel

- In broadcast network, the key issue is how to share the channel among several users.

- Ex *a conference call with five people*

-Broadcast channels are also called as multi-access channels or random access channels.

-Multi-access channel belong to a sublayer at the DL layer called the MAC sublayer.

The Channel Allocation problem:

a) **Static channel allocation** in LANs & MANs

i) **FDM** ii) **TDM**

Drawbacks: -1) Channel is wasted if one or more stations do not send data.

2) If users increases this will not support.

b) Dynamic channel allocation

i) Pure **ALOHA** & Slotted **ALOHA**

CSMA/CD

ii) CSMA

CSMA/CA

Pure ALOHA

-1970's Norman Abramson end his colleagues devised this method, used ground –based radio broad costing. This is called the **ALOHA** system.

-The basic idea, many users are competing for the use of a single shared channel.

-There are two versions of ALOHA: **Pure and Slotted**.

-Pure ALOHA does not require global time synchronization, where as in slotted ALOHA the time is divided into discrete slots into which all frames must fit.

-Let users transmit whenever they have data to be sent.

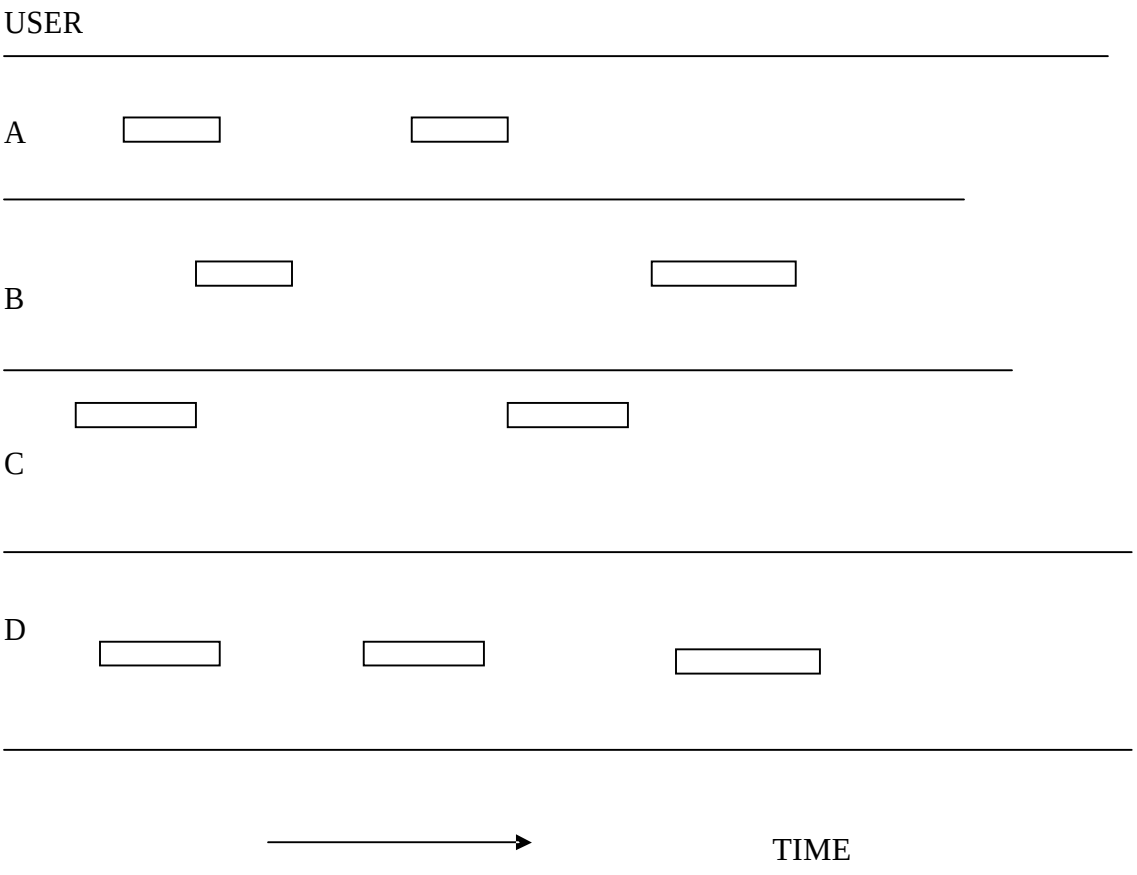
-There will be collisions and all collided frames will be damaged.

-Senders will know through feedback property whether the frame is destroyed or not by listening channel.

[-With a LAN it is immediate, with a satellite, it will take 270m sec.]

-If the frame was destroyed, the sender waits random amount of time and again sends the frame.

-The waiting time must be random otherwise the same frame will collide over and over.



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 destroyed.

-We have to find out what is the efficiency of an ALOHA channel?

-Let us consider an infinite collection of interactive users sitting at their systems (stations).

-A user will always in two states **typing or waiting**.

-Let the 'Frame time' denotes the time required to transmit one fixed length frame.

-Assume that infinite populations of users are generating new frames according to possion distribution with mean N frames per frame time.

-If N>1 users are generating frames at a higher rate than the channel can handle.

-For reasonable throughput 0<N<1.

-In addition to new frames, the station also generates retransmission of frames.

-Old and new frames are G per frame time.

-G> N

-At low load there will be few collisions, so G ~ N

-Under all loads, the throughput $S = GP_0$, 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 with one frame time of its start.

-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 frame.

-Similarly, any other frame started b/w t_0+t and t_0+2t will bump into the end of the shaded frame.

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

$$P_r[k] = \frac{G^k e^{-G}}{k!}$$

-The probability of zero frames is just e^{-G}

-In an interval two frame times long, the mean number at frames generated is 2G.

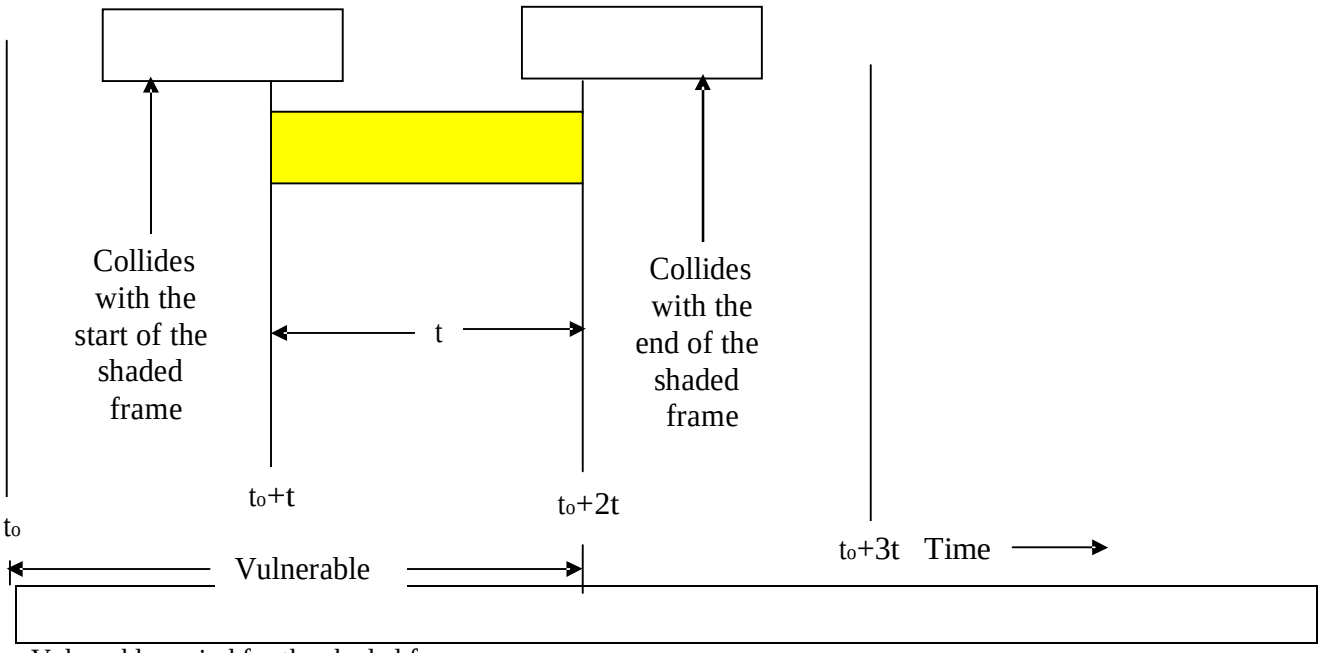
-The probability at no other traffic being initiated during the entire vulnerable period is given by

$$P_0 = e^{-2G}$$

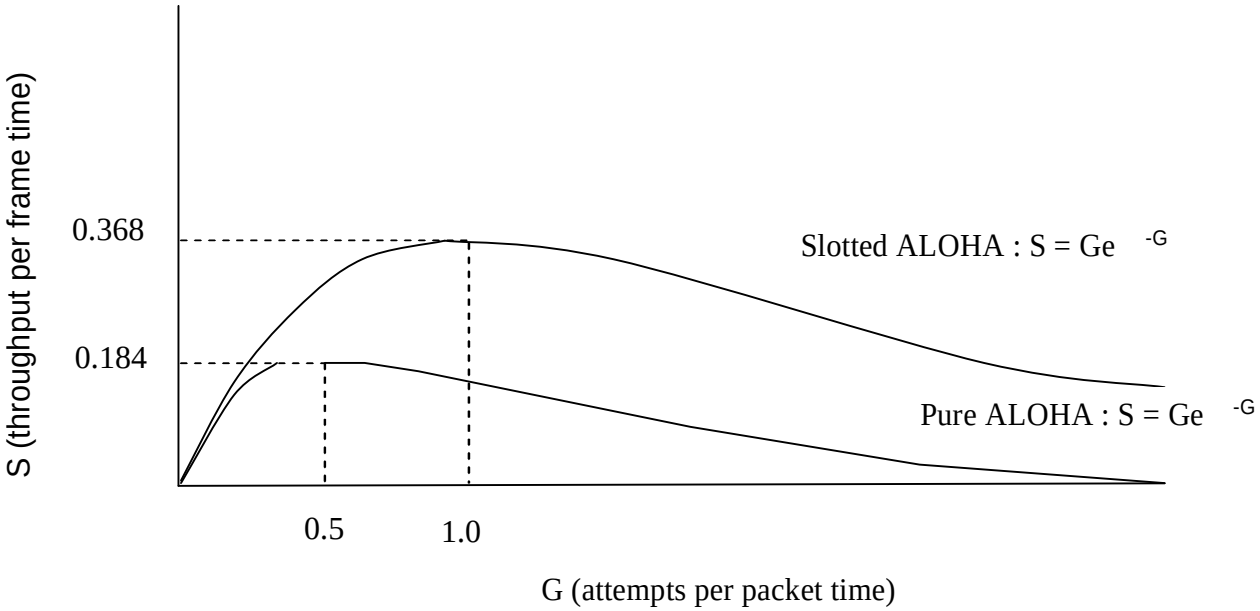
$$S = Ge^{-2G} \qquad [S=GP_0]$$

The Maximum through put occurs at G=0.5 with S=1/2e = 0.184

The channel utilization at pure ALOHA =18%.



Vulnerable period for the shaded frame



Throughput versus offered traffic for ALOHA systems

Slotted ALOHA

-In 1972, Roberts' devised a method for doubling the capacity of ALOHA system.

-In this system the time is divided into discrete intervals, each interval corresponding to one frame.

-One way to achieve synchronization would be to have one special station emit a pip at the start of each interval, like a clock.

-In Roberts' method, which has come to be known as slotted ALOHA, in contrast to Abramson's 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.

-Thus the continuous pure ALOHA is turned into a discrete one.

-Since the vulnerable period is now halved, the of no other traffic during the same slot as our test frame is e^{-G} which leads to

$S = Ge^{-G}$

- At G=1, slotted ALOHA will have maximum throughput.

- So S=1/e or about 0.368, twice that of pure ALOHA.

- The channel utilization is 37% in slotted ALOHA.

Carrier Sense Multiple Access Protocols

Protocols in which stations listen for a carrier (transmission) and act accordingly are called carries sense protocols.

Persistent CSMA

When a station has data to send, it first listens to the channel to see if any one else is transmitting at that moment. If the channel is busy, the station waits until it become 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 also 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. The longer the propagation delay the worse the performance of the protocol.

Even if the propagation delay is zero, there will be collisions. If two stations listen the channel, that is idle at the same, both will send frame and there will be collision.

Non persistent CSMA

In this, before sending, a station sense the channel. If no one else is sending, the station begins doing so it self. However, if the channel is busy, the station does not continually sense it but it waits a random amount of time and repeats the process.

This algorithms leads to better channel utilization but longer delays then 1-persistent CSMA.

With persistent CSMA, what happens if two stations become active when a third station is busy? Both wait for the active station to finish, then simultaneously launch a packet, resulting a collision. There are two ways to handle this problem.

a) P-persistent CSMA b) exponential backoff.

P-persistent CSMA

The first technique is for a waiting station not to launch a packet immediately when the channel becomes idle, but first toss a coin, and send a packet only if the coin comes up heads. If the coin comes up tails, the station waits for some time (one slot for slotted CSMA), then repeats the process. The idea is that if two stations are both waiting for the medium, this reduces the chance of a collision from 100% to 25%. A simple generalization of the scheme is to use a biased coin, so that the probability of sending a packet when the medium becomes idle is not 0.5, but p, where $0 < p < 1$. We call such a scheme **P-persistent CSMA**. The original scheme, where p=1, is thus called 1-persitent CSMA.

Exponential backoff

The key idea is that each station, after transmitting a packet, checks whether the packet transmission was successful. Successful transmission is indicated either by an explicit acknowledgement from the receiver or the absence of a signal from a collision detection circuit. If the transmission is successful, the station is done. Otherwise, the station retransmits the packet, simultaneously realizing that at least one other station is also contending for the medium. To prevent its retransmission from colliding with the other station's retransmission, each station backs off (that is, idles) for a random time chosen from the interval

[0,2*max-propagation_delay] before retransmitting its packet. If the retransmission also fails, then the station backs off for a random time in the interval [0,4* max_propagation_delay], and tries again. Each subsequent collision doubles the backoff interval length, until the retransmission finally succeeds. On a successful transmission, the backoff interval is reset to the initial value. We call this type of backoff exponential backoff.

CSMA/CA

In many wireless LANS, unlike wired LANS, the station has no idea whether the packet collided with another packet or not until it receives an acknowledgement from receiver. In this situation, collisions have a greater effect on performance than with CSMA/CD, where colliding packets can be quickly detected and aborted. Thus, it makes sense to try to avoid collisions, if possible. CSMA/CA is basically p-persistence, with the twist that when the medium becomes idle, a station must wait for a time called the interframe spacing or IFS before contending for a slot. A station gets a higher priority if it is allocated smaller inter frame spacing.

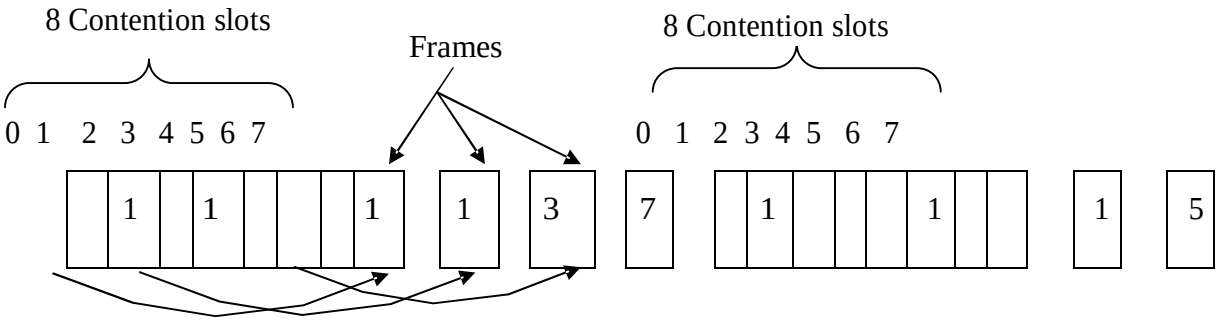
When a station wants to transmit data, it first checks if the medium is busy. If it is, it continuously senses the medium, waiting for it to become idle. When the medium becomes idle, the station first waits for an interframe spacing corresponding to its priority level, then sets a contention timer to a time interval randomly selected in the range [0,CW], where CW is a predefined contention window length. When this timer expires, it transmits a packet and waits for the receiver to send an ack. If no ack is received, the packet is assumed lost to collision, and the source tries again, choosing a contention timer at random from an interval twice as long as the one before(binary exponential backoff). If the station senses that another station has begun transmission while it was waiting for the expiration of the contention timer, it does not reset its timer, but merely freezer it, and restarts the countdown when the packet completes transmission. In this way, stations that happen to choose a longer timer value get higher priority in the next round of contention.

Collision-Free Protocols

A Bit-Map Protocol

In 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 zeroth slot. No other station is allowed to transmit during this slot. Regardless of what station 0 does, station 1 gets the

opportunity to transmit a 1during slot 1, but only if it has a frame queued. In general, station j may announce the fact 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 with to transmit.



The basic bit-map protocol

Since everyone agrees on who goes next, there will never be any collisions. After the last ready station has transmitted its frame, an event all stations can easily monitor, another N bit contention period is begun. If a station becomes ready just after its bit slot has passed by, it is out of luck and must remain silent until every station has had a chance and the bit map has come around again. Protocols like this in which the desire to transmit is broadcast before the actual transmission are called reservation protocols.

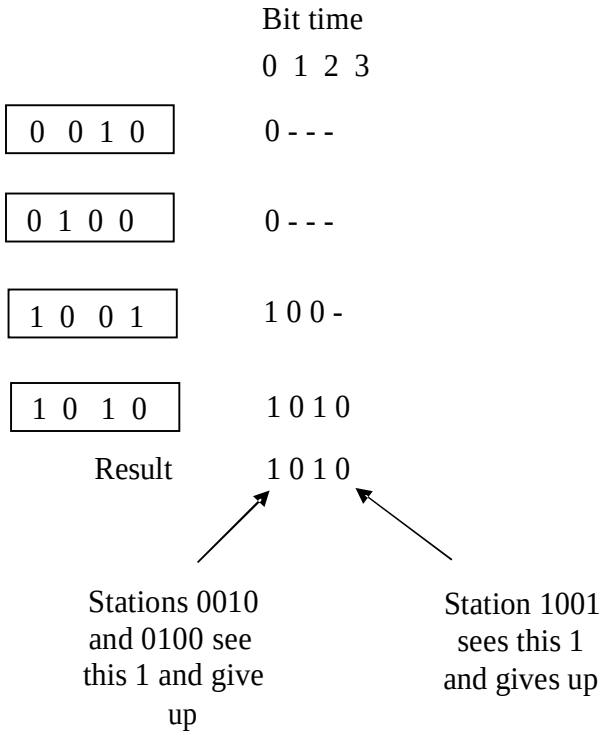
Binary Countdown

A problem with the basic bit-map protocol is that the overhead is 1 bit per station. A station wanting to use the channel now broadcasts its address as a binary bit string, starting with the high-order bit. All addresses are assumed to be the same length. The bits in each address position from different stations are BOOLEAN ORed together. We will call this protocol binary countdown. It is used in Datakit.

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 station 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. 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.

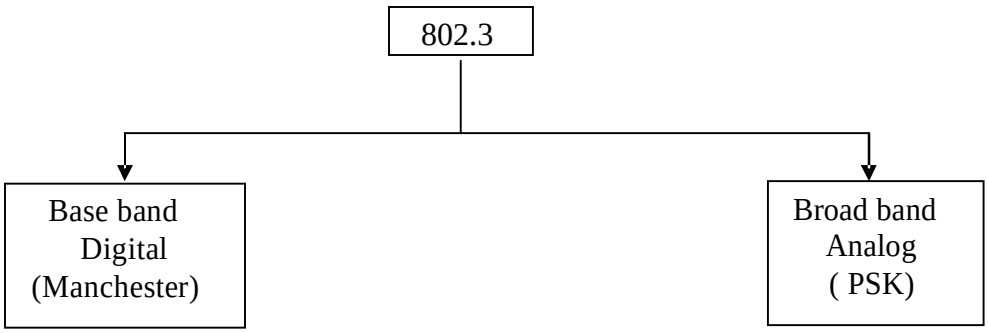
The next bit is 0, and both stations continue. The next bit is 1, so station 1001 gives up. The winner is station 1010, because it has the highest address. After winning the bidding, it may now transmit a frame, after which another bidding cycle starts.

The binary countdown protocol. A dash indicates silence



IEEE Standard 802 for LANS and MANS

The IEEE 802.3 is for a 1-persistent CSMA/CD LAN. Xerox built a 2.94 Mbps CSMA/CD system to connect over 100 personal workstations on 1-Km cable. This system was called Ethernet through which electromagnetic radiation was once thought to propagate. Xerox DEC and Intel came with another standard for 100 Mbps Ethernet. This differs from old one that it runs at speeds from 1 to 10 Mbps on various media. The second difference between these two is in one header (802.3 length field is used for packet type in Ethernet).



10Base5, 10Base2

10Base-T, 1Base5

100 Base-T

10 Broad 36

802.3 Cabling

Five types of cabling are commonly used, 10Base5 cabling called thick Ethernet, came first. It resembles a yellow garden hose, with markings every 2.5 m to show where the taps go. Connections to it are generally made using **vampire taps**, in which a pin is carefully forced halfway into the coaxial cable's core. The notation 10Base5 means that it operates at 10 Mbps, uses baseband signaling, and can support segments of up to 500m.

Name	**Cable**	**Max. segment**	**Nodes/seg.**	**Advantages**
10Base5	Thick coax	500 m	100	Good for backbones
10Base2	Thin coax	200 m	30	Cheapest system
10Base-T	Twisted pair	100 m	1024	Easy maintenance
10Base-F	Fiber optics	2000 m	1024	Best between buildings

The second cable type was **10Base2** or thin Ethernet, which, in contrast to the garden-hose-like thick Ethernet, bends easily. Connections to it are made using industry standard BNC connectors to form T-junctions, rather than using vampire taps. These are easier to use and more reliable. Thin Ethernet is much cheaper and easier to install, but it can run for only 200m and can handle only 30 machines per cable segment.

Cable breaks, bad taps, or loose connectors can be detected by a devise called time domain reflectometry.

For 10Base5, a transceiver is clamped securely around the cable so that its tap makes contact with the inner core. The transceiver contains the electronics that handle carrier detection and collision detection. When a collision is detected, the transceiver also puts a

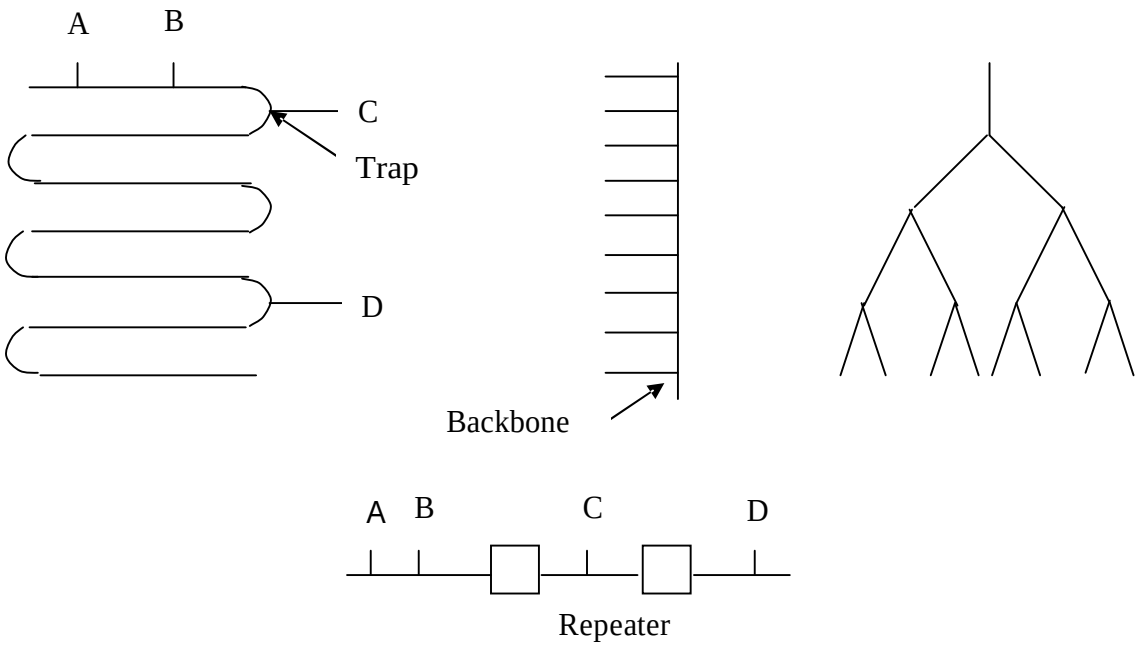
special invalid signal on the cable to ensure that all other transceivers also realize that a collision has occurred.

The transceiver cable terminates on an interface board inside the computer. The interface board contains a controller chip that transmits frames to, and receives frames from, the transceiver. The controller is responsible for assembling the data into the proper frame format, as well as computing checksums on outgoing frames and verifying them on incoming frames.

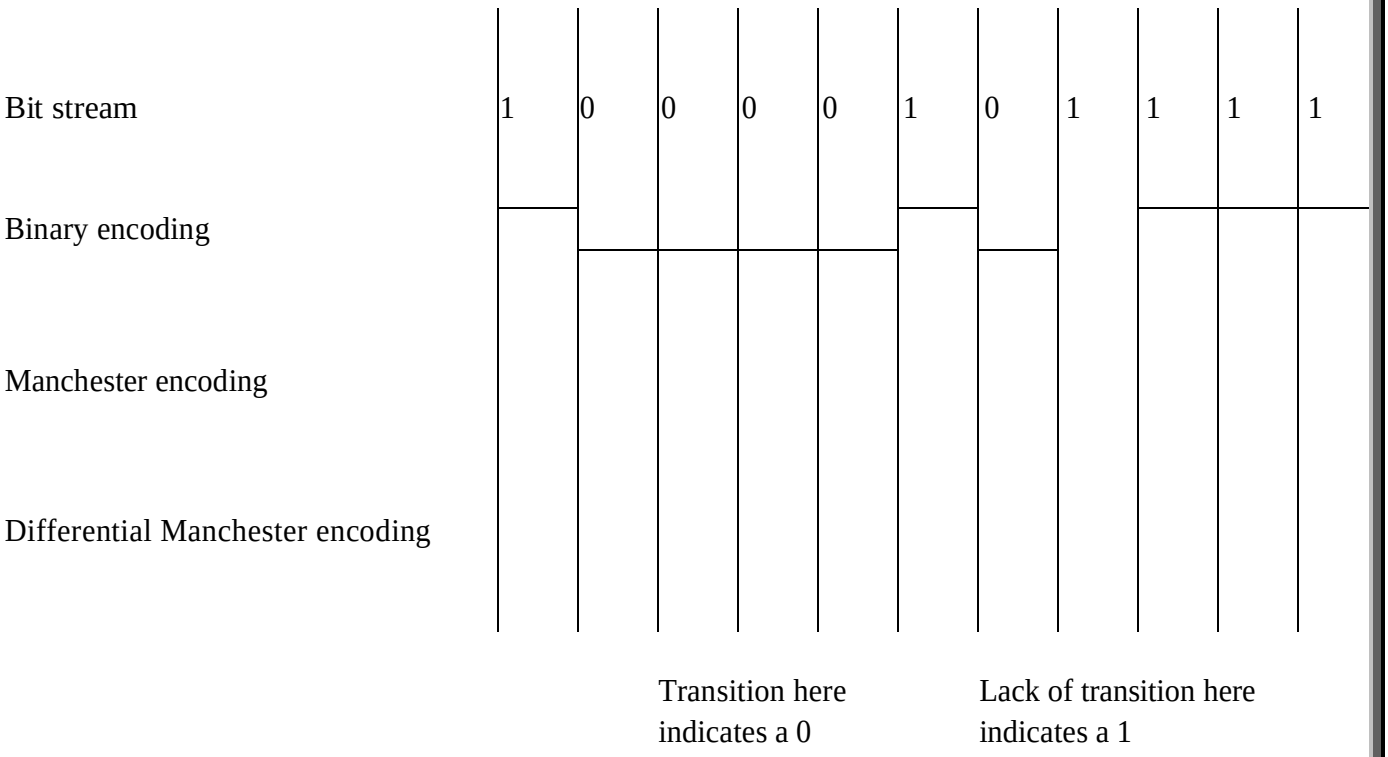
With 10Base2, the connection to the cable is just a passive BNC T-junction connector. The transceiver electronics are on the controller board, and each station always has its own transceiver.

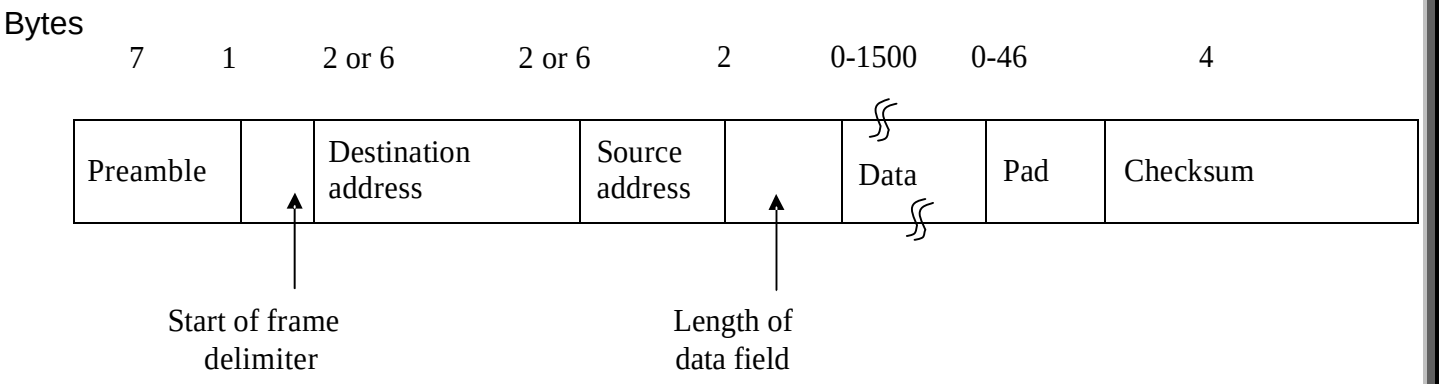
With 10Base-T, there is no cable at all, just the hub (a box full of electronics). Adding or removing a station is simple in this configuration, and cable breaks can be detected easily. The disadvantage of 10Base-T is that the maximum cable run from the hub is only 100m, may be 150m if high-quality (category 5) twisted pairs are used. 10Base-Tis becoming steadily more popular due to the ease of maintenance. 10Base-F, which uses fiber optics. This alternative is expensive due to the cost of the connectors and terminators, but it has excellent noise immunity and is the method of choice when running between buildings or widely separated hubs.

Each version of 802.3 has a maximum cable length per segment. To allow larger networks, multiple cables can be connected by repeaters. A repeater is a physical layer device. It receives, amplifies, and retransmits signals in both directions. As far as the software is concerned, a series of cable segments connected by repeaters is no different than a single cable (except for some delay introduced by the repeater). A system may contain multiple cable segments and multiple repeaters, but no two transceivers may be more than 2.5km apart and no path between any two transceivers any traverse more than four repeaters.



802.3 uses Manchester Encoding and differential Manchester Encoding





The 802.3 MAC sub layer protocol:

I) Preamble:

Each frame start with a preamble of 7 bytes each containing a bit pattern 10101010.

II) Start of frame byte:

It denotes the start of the frame itself. It contains 10101011.

III) Destination address:

This gives the destination address. The higher order bit is zero for ordinary address and 1for group address (Multi casting). All bits are 1s in the destination field frame will be delivered to all stations (Broad casting).

The 46 thbi (adjacent to the high-order bit) is used to distinguish local from global addresses.

IV) Length field:

This tells how many bytes are present in the data field from 0 to 1500.

V) Data field:

This contains the actual data that the frame contains.

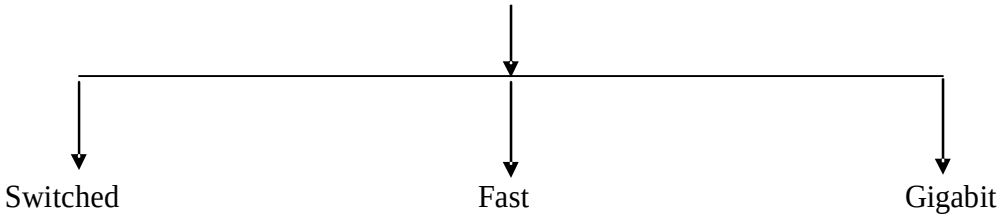
VI) Pad:

Valid frame must have 64 bytes long from destination to checksum. If the frame size less than 64 bytes pad field is used to fill out the frame to the minimum size.

VII) Checksum:

It is used to find out the receiver frame is correct or not. CRC will be used here.

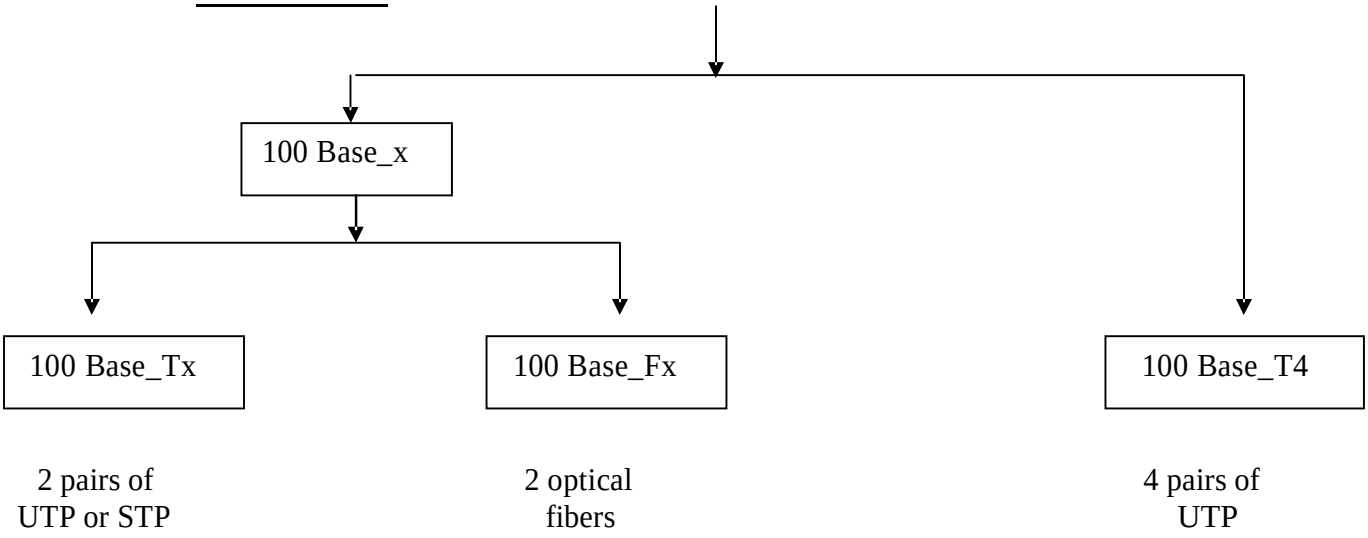
Other Ethernet Networks



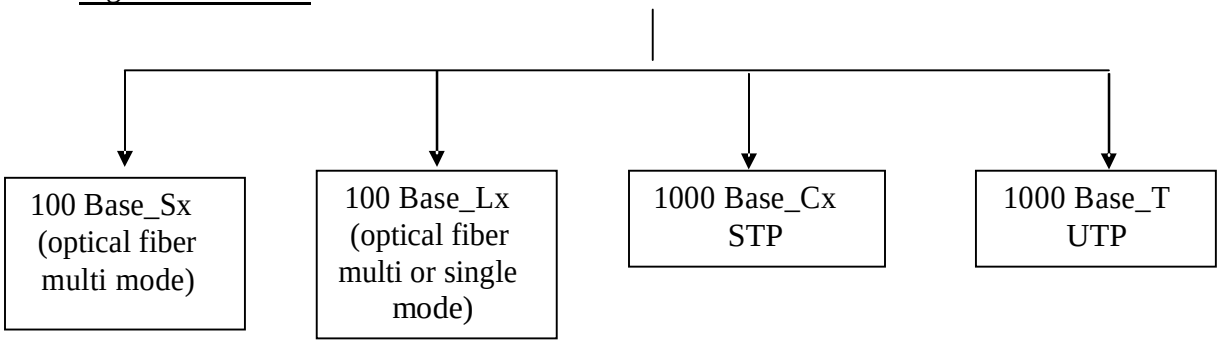
Switched Ethernet:

- 10 Base-T Ethernet is a shared media network.
- The entire media is involved in each transmission.
- The HUB used in this network is a passive device. (not intelligent).
- In switched Ethernet the HUB is replaced with switch. Which is a active device (intelligent)

Fast Ethernet

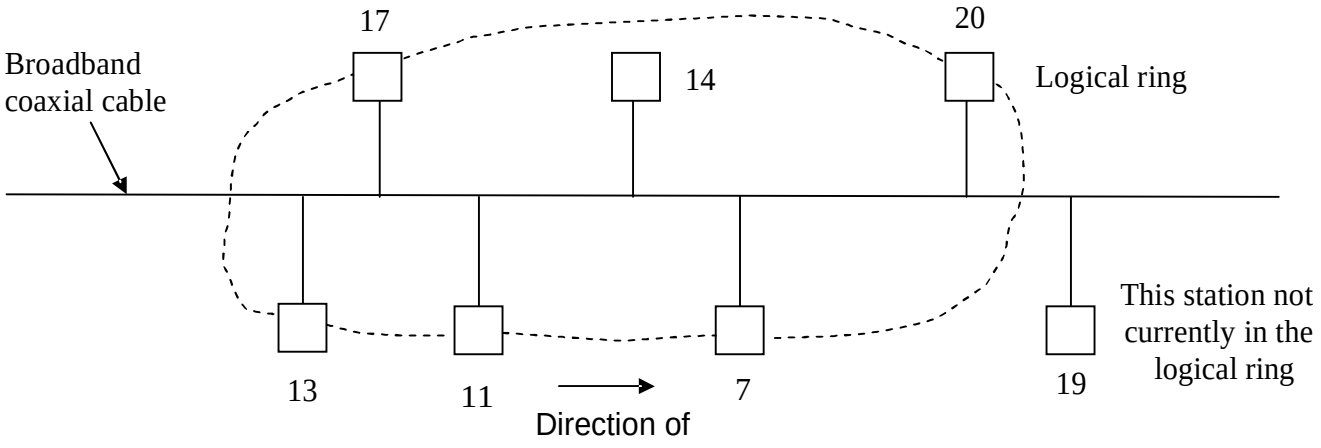


Gigabit Ethernet



IEEE 802.4 (Token Bus)

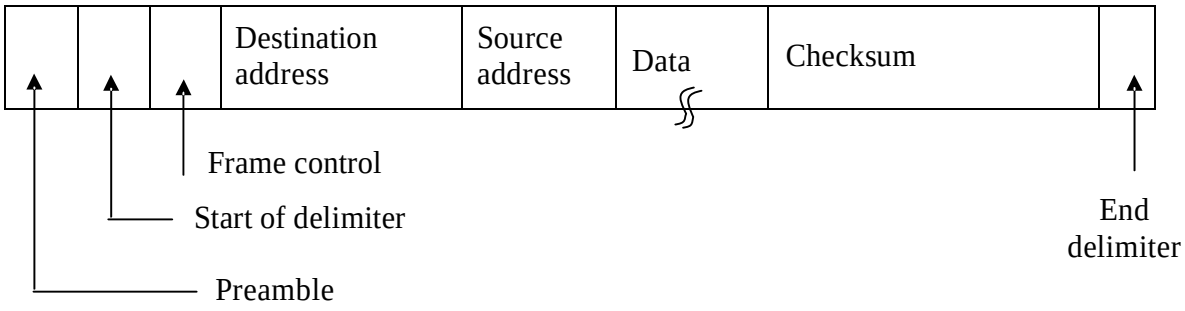
802.3 frames do not have priorities, making them unsuited for real-time systems in which important frames should not be held up waiting for unimportant frames. A simple system with a known worst case is a ring in which the stations take turns sending frames. If there are n stations and it takes T sec to send a frame, no frame will ever have to wait more than nT sec to be sent.



This standard, 802.4, describes a LAN called a token bus. Physically, the token bus is a linear or tree-shaped cable onto which the stations are attached. Logically, the stations are organized into a ring, with each station knowing the address of the station to its "left" and "right." When the logical ring is initialized, the highest numbered station may send the first frame. After it is done, it passes permission to its immediate neighbor by sending the neighbor a special control frame called a token. The token propagates around the logical ring, with only the token holder being permitted to transmit frames. Since only one station at a time holds the token, collisions do not occur.

Since the cable is inherently a broadcast medium, each station receives each frame, discarding those not addressed to it. When a station passes the token, it sends a token frame specifically addressed to its logical neighbor in the ring, irrespective of where that station is physically located on the cable. It is also worth noting that when stations are first powered on, they will not be in the ring, so the MAC protocol has provisions for adding stations to, and deleting stations from, the ring. For the physical layer, the token bus uses the 75-ohm broadband coaxial cable used for cable television. Both single and dual-cable systems are allowed, with or without head-ends.

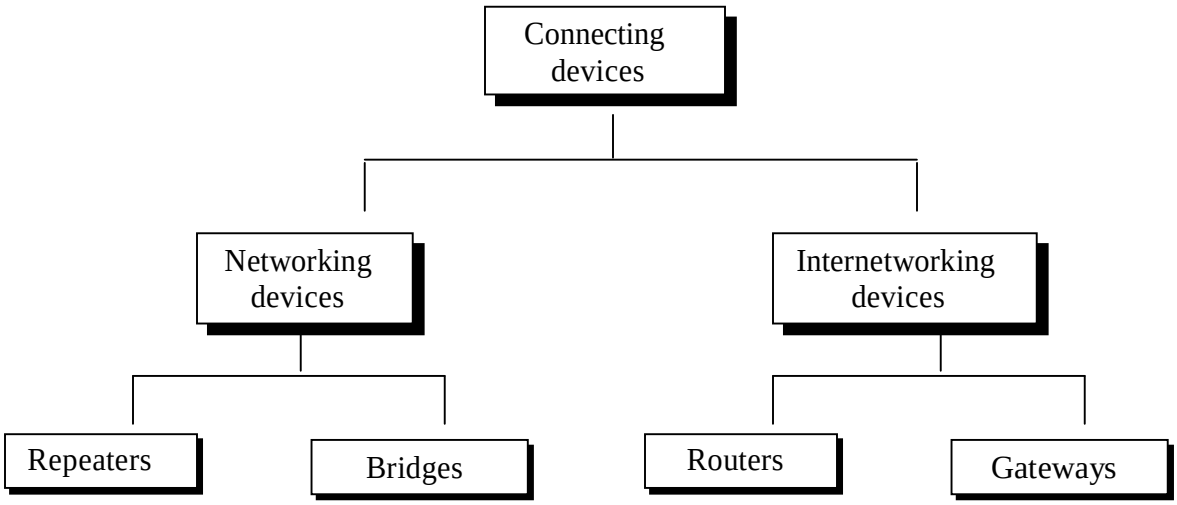
Bytes ≥	1	1	1	2 or 6	2 or 6	0-8182	4	1



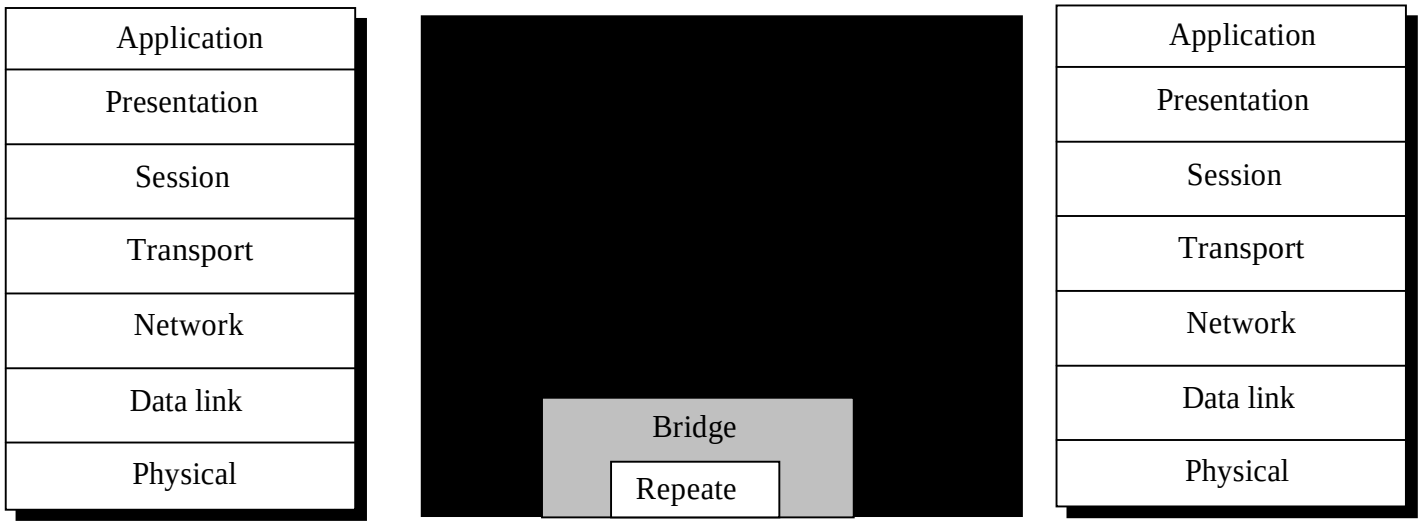
The frame control field is used to distinguish data frames from control frames. Fro data frames, it carries the frame's priority. It can also carry an indicator requiring the destination station to acknowledge correct or incorrect receipt of the frame.

For control frames, the frame control field is used to specify the frame type. The allowed types include token passing and various ring maintenance frames, including the mechanism for letting new stations enter the ring, the mechanism for allowing stations to leave the ring, and so on.

Connecting devices



Connecting devices and the OSI model



Bridges

LANS can be connected by devices called bridges, which operate in the data link layer. Bridges do not examine the network layer header and can thus copy IP, IPX, and OSI packets equally well.

The various reasons why the bridges are used.

1) Many university and corporate departments have their own LANS, primarily to connect their own personal computers, workstations, and servers. Since the goals of the various departments differ, different departments choose different LANS, without regard to what other departments are doing. Sooner or later, there is a need for interaction, so bridges are needed.

2) The organization may be geographically spread over several buildings separated by considerable distances. It may be cheaper to have separate LANS in each building and connect them with bridges and infrared links than to run a single coaxial cable over the entire site.

3) It may be necessary to split what is logically a single LAN into separate LANS to accommodate the load. Putting all the workstations on a single LAN- the total bandwidth needed is far too high. Instead multiple LANS connected by bridges are used.

4) In some situations, a single LAN would be adequate in terms of the load, but the physical distance between the most distant machines is too great (e.g., more than 2.5km for 802.3). Even if laying the cable is easy to do, the network would not work due to the

excessively long round-trip delay. Only solution is to partition the LAN and install bridges between the segments.

5) There is the matter of reliability. On a single LAN, a defective node that keeps outputting a continuous stream of garbage will cripple the LAN. Bridges can be inserted at critical places, to prevent a single node which has gone berserk from bringing down the entire system.

6) And last, bridges can contribute to the organization's security. By inserting bridges at various places and being careful not to forward sensitive traffic, it is possible to isolate parts of the network so that its traffic cannot escape and fall into the wrong hands.