

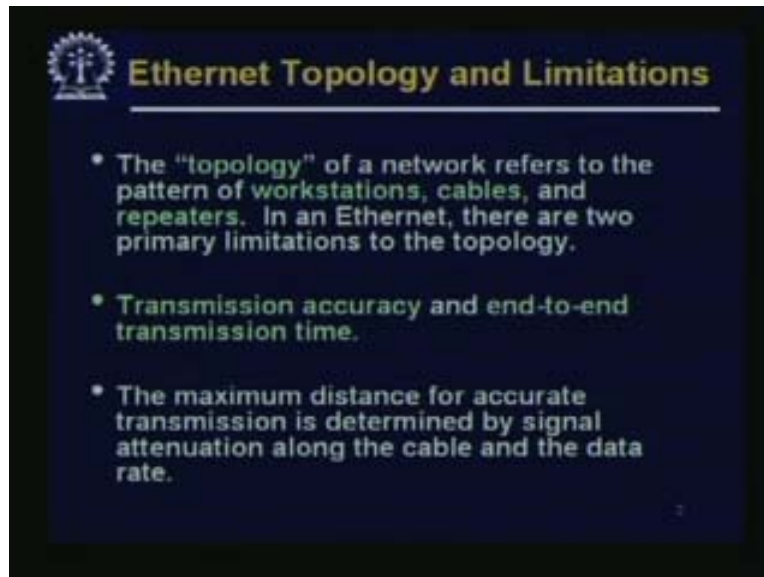
Computer Networks
Prof: Sujoy Ghosh
Department of Computer Science and Engineering
Indian Institute of Technology, Kharagpur
Lecture name - 20
Modern Internet

Good day! In this lecture we will continue our discussion about Ethernet. In the previous lectures, we have seen the basic data link protocol used in Ethernet. Ethernet is a very popular and nowadays it is almost ubiquitous in the LAN area. It is a very widely used system and it has evolved from whatever was the approach to Ethernet in the earlier days and the basic Ethernet protocol that we had discussed. People are sort of shifting away from it because as the technology develops and as things like switches etc., become cheaper as the speed of network goes up, there is a slight shift in emphasis in modern Ethernet. In this lecture, what we are going to do is that first, we are going to have a look at the physical layer of the Ethernet. There is not much to discuss about that, and then we will see how we are shifting from a shared medium to a switched Ethernet kind of concept, and how speed is increasing. Then, we will discuss a little bit about Ethernet, LAN as a concept is. So, we discuss the modern Ethernet that is how the Ethernet is evolved.

(Refer Slide Time: 02:15)

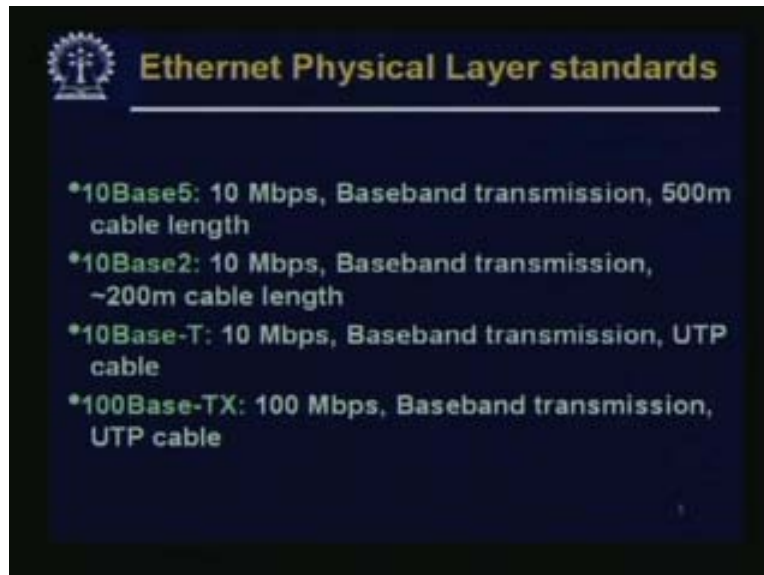


(Refer Slide Time: 02:21)



Before we go into that, we will look at the topology and some of the limitations of the earlier versions of Ethernet. The topology of a network refers to the pattern of workstations, cables and repeaters. In an Ethernet, there are two primary limitations to the topology: transmission accuracy and end-to-end transmission time. Transmission accuracy means that Ethernet really under a data link layer is an unreliable protocol in sense that it does not do a lot of acknowledgement or negative acknowledgement and things like that. Therefore, essentially the transmission accuracy should be fairly tolerable. What happens is that the maximum distance for accurate transmission is determined by signal attenuation along the cable, and it would depend on the on the quality of the cable, the kind of cable we used and the data rate. The other important point is the end-to-end transmission time as we have seen, because of this collision and other things this end-to-end transmission time is important. It's not just for the data rate or things to move faster but for other reasons, also this end-to-end transmission time is a factor.

(Refer Slide Time: 03:48)



These are the Ethernet physical layer standards that we have – some of them, they are actually more. So 10Base5, that is, 10 Mbps base band transmission and 500 m; this 5 here stands for 500 m cable length, this is called thick Ethernet or has thick coaxial cables and this has become obsolete. Then 10Base2 is 10 Mbps base band transmission; once again all of these are base band transmissions; that means we do not modulate them to a higher channel or frequency channel. So, 10 Mbps and 2 for above 200 m cable line, or lines of 180 to 200 m length. Then, 10BaseT; this again is for 10 Mbps base band transmission, which uses UTP cable. There are various versions of UTP cable like categories 3, 4, 5, etc. 100Base-TX is 100 Mbps base band transmission, also uses UTP cable.

(Refer Slide Time: 04:58)

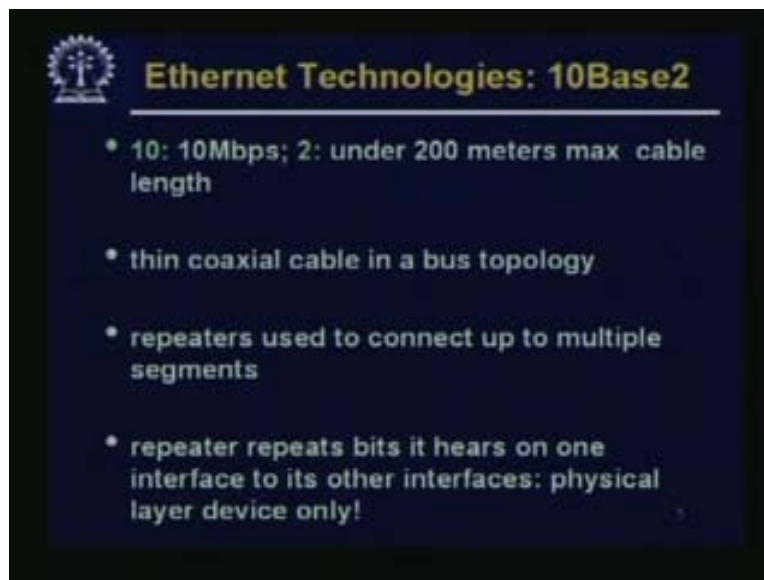
Ethernet Cabling

The most common kinds of Ethernet cabling

Name	Cable	Max. seg.	Nodes/seg.	Advantages
10Base5	Thick coax	500 m	100	Original cable; now obsolete
10Base2	Thin coax	185 m	30	No hub needed
10Base-T	Twisted pair	100 m	1024	Cheapest system
10Base-F	Fiber optics	2000 m	1024	Best between buildings

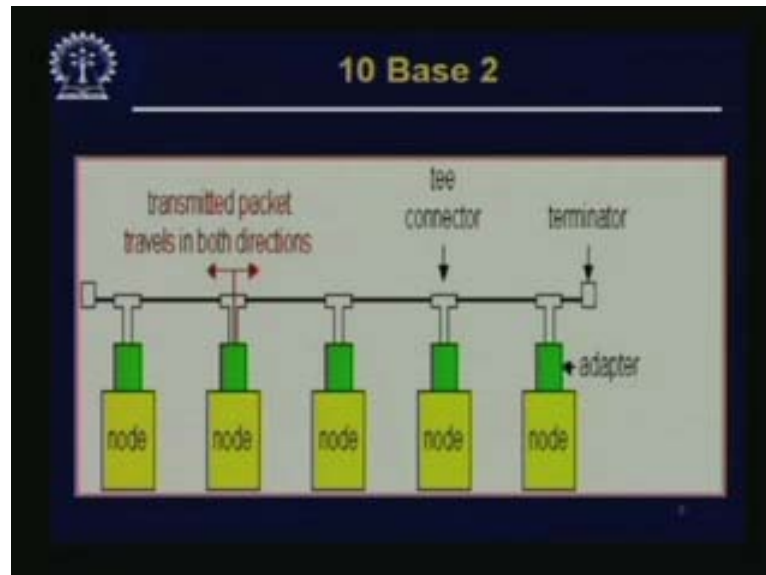
So, these are the of maximum segment length. As we can see, 10Base5 is 500 m; 10Base2 is 185 m; 10Base-T is 100 m. Fibre optic cable 10Base-FI goes up to 2 km. actually, 10Base-F is not in much use these days, because fibre optic components became cheap and available, and people have moved from this 10 Mbps rate to the 100 Mbps rate. The point is that, the fibre optic cable goes much longer than copper cables and it is best between buildings; that means when it goes through some open space, etc., we avoid copper cable most times because of electric interference and other problems, not to mention the distance. The nodes per segment in 10Base-F could also be more; we will come back to this topic of nodes per segment later. As we have seen previously, 10Base5 and 10Base2 have many limitations; for example, in 10Base2, you see the figure of 30 and a network segment having only 30 nodes is a constraint these days because all kinds of things are being networked. For completeness sake, I will make a mention of 10Base2 and other kinds of

(Refer slide Time: 06:51)



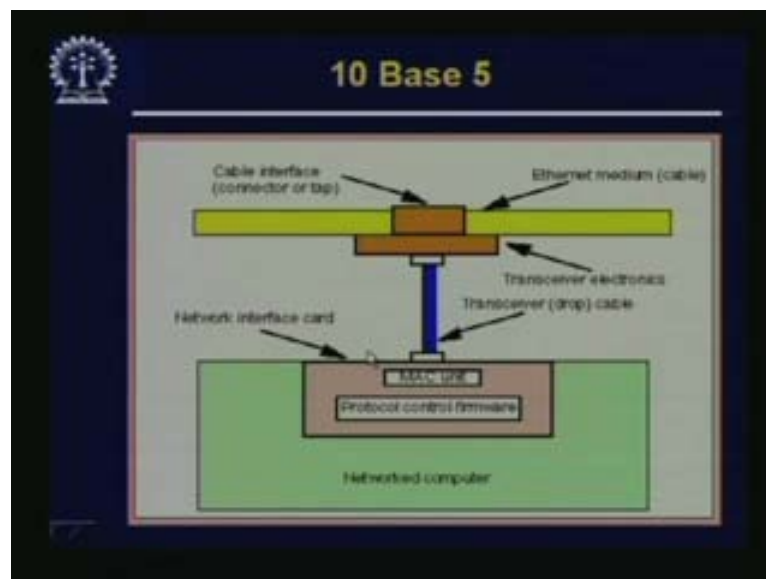
Technologies, although they have become obsolete now – 10Base2 has 10 Mbps, 200 m maximum cable length, is a thin coaxial cable in a bus topology; this is a classical bus topology. Repeaters are used to connect up to multiple segments. Repeater repeats bits; it hears on one interface to the other interface, so this is a physical layer device just for amplification, strengthening the signal.

(Refer slide Time: 07:19)



This is a figure showing how they are connected. We have this thin cable, which has terminated on both the ends and these nodes would be connected using some t connectors over here, and this is the network adapter. When a network adapter pushes some signal on to the cable, the transmitted packet travels in both the directions.

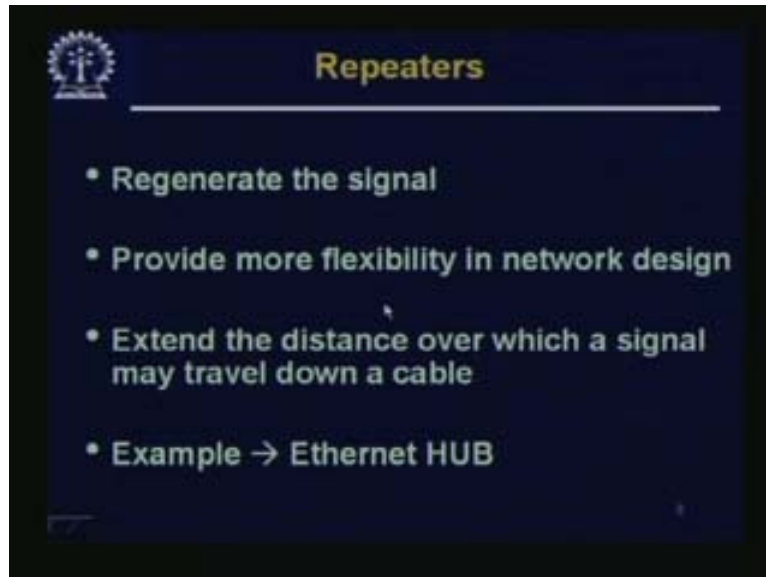
(Refer Slide Time: 07:55)



This is just a mention about 10Base5, because connecting a 10Base5 is more cumbersome and in 10Base2, connection would get loose quite often in the earlier days.

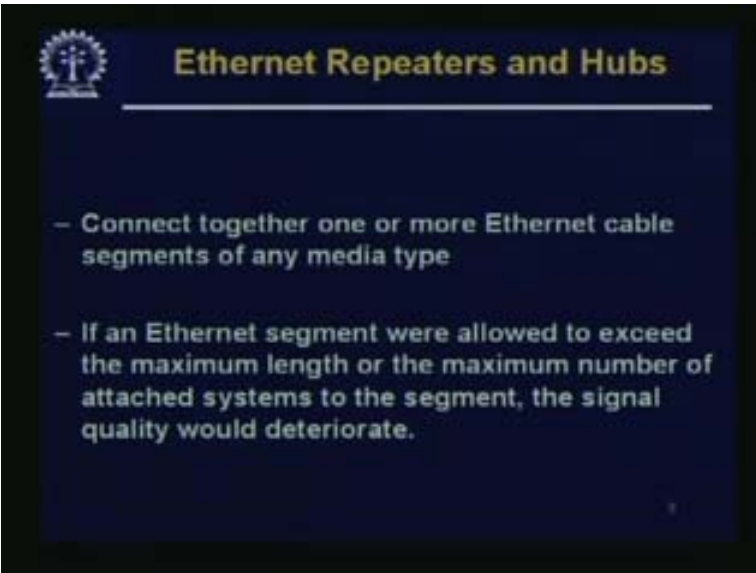
Anyway, this is a 10Base5; it has a thick coaxial cable, the transceiver, and transceiver, drop cable and there would be a network interface card here with this MAC unit and protocol control firmware etc. As I said, this has become obsolete now.

(Refer Slide Time: 08:33)



Let us talk a little about repeaters. Repeater is a physical layer kind of device, which regenerates the signal, and provides more flexibility in network design because it depends on where we want to extend your LAN. So if you want to extend your LAN in some direction, make running into this distance limitation 200 m for 10Base2 etc. You may sort of increase that by just putting another segment on the other side of repeater so that you can go another 200 m. Therefore you can extend the distance over which a signal may travel down a cable; an example of a repeater is Ethernet hub. A hub has two purposes: one is to provide a collision domain and replace the cable. The other function of the hub is the repeater function; that means it regenerates the signal. This connects together or a repeater or a hub. It connects together one or more Ethernet cable segments of any media type. If an internet segment were allowed to exceed the maximum length or the maximum number of attached systems to the segment, the signal quality would deteriorate.

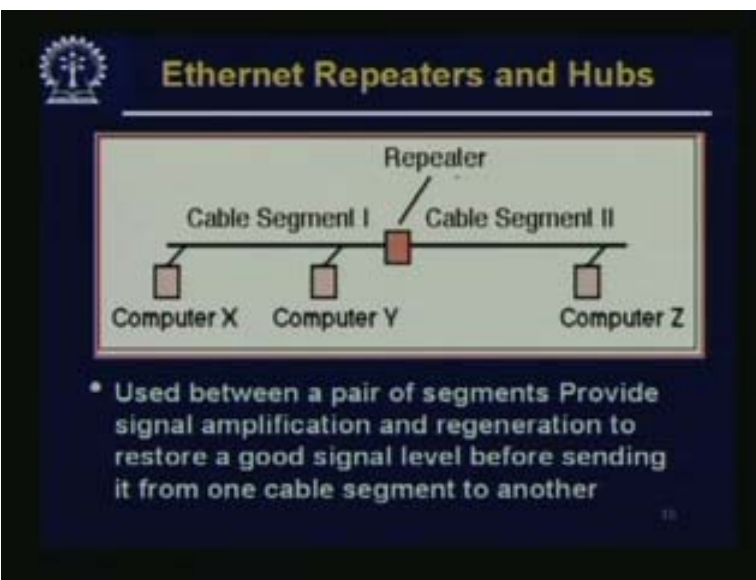
(Refer Slide Time: 09:48)



Ethernet Repeaters and Hubs

- Connect together one or more Ethernet cable segments of any media type
- If an Ethernet segment were allowed to exceed the maximum length or the maximum number of attached systems to the segment, the signal quality would deteriorate.

(Refer Slide Time: 10:03)



Ethernet Repeaters and Hubs

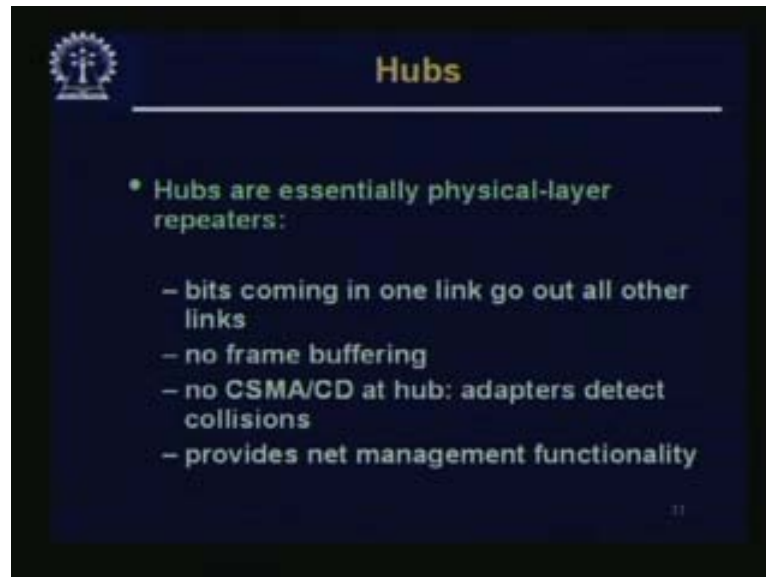
Diagram illustrating the use of a Repeater:

The diagram shows a horizontal line representing a cable. A red square labeled "Repeater" is positioned in the middle. To the left of the repeater is "Cable Segment I" and to the right is "Cable Segment II". Below Cable Segment I are two small squares labeled "Computer X" and "Computer Y". Below Cable Segment II is one small square labeled "Computer Z".

- Used between a pair of segments Provide signal amplification and regeneration to restore a good signal level before sending it from one cable segment to another

You can see that you have one cable segment here. This goes up to its maximum distance there is a cable segment too again going up to the maximum distance. But we put this repeater in between so that, we can increase the total distance that may be covered. These are the different computers connected to the cable. It is used between a pair of segments; this is a simple repeater to provide signal amplification and regeneration to restore a good signal level before sending it from one cable segment to another.


(Refer Slide Time: 10:37)



Hubs are essentially physical layer repeaters; bits come in one link and go out all other links. There is no frame buffering or CSMA/CD at the hub. Adapters detect the collisions and provide net management functionality, which means that this hub, although it is an active device, in the sense that this has some electronics in it, in its action it is more passive. Its active component is restricted to the fact that it regenerates the signal; it amplifies and regenerates it. So far as the other intelligence is concerned, like detecting collisions, it does nothing like that. If two signals coming at two different ports of the hub come to the hub at the same point of time, they will collide. So, a hub is a whole co-axial cable, which has been collapsed into one collision domain inside the hub; we could look at that way.

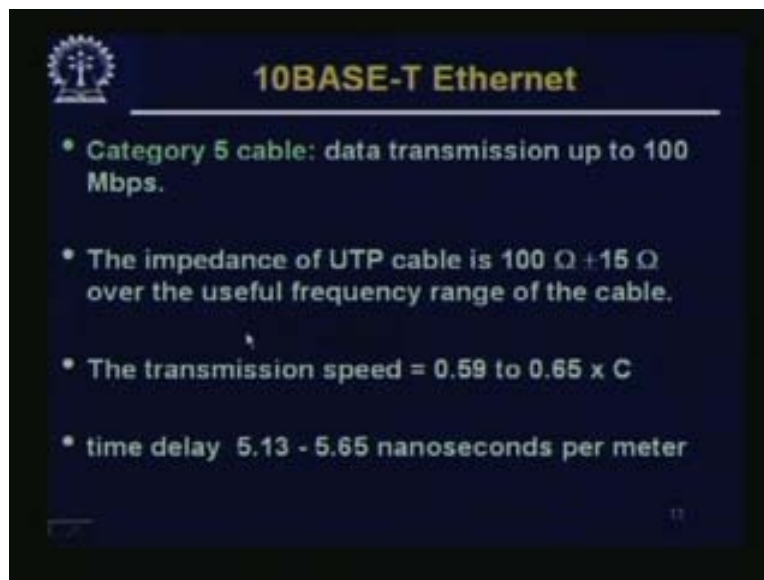
But it does provide some net management functionality. That means, if you have a managed hub, there are two kinds of devices. The point is that a network is usually a geographically distributed, dispersed kind of entity; but this has to be managed. Therefore if there is a problem somewhere, may be some user has put in a complaint, what we will do is that there are so many things which connect from the user to the central network. There are cables, may be switches, hubs, and other things, so it is important that you should be able to remotely tell whether some active device is functioning. We cannot do this on passive things like a cable but there are managed switches and managed hubs etc., whom a central network station may interrogate and find out whether its health is alright or not. So if you have a managed hub, you have net management functionality; of course, the managed hubs are costlier than unmanaged hubs.

(Refer Slide Time: 13:03)



So this is the picture of hub. This is a co-axial cable coming in here and this is a t joint; the t joint is feeding into the hub. There may be a number of ports and a number of segments may be connected together.

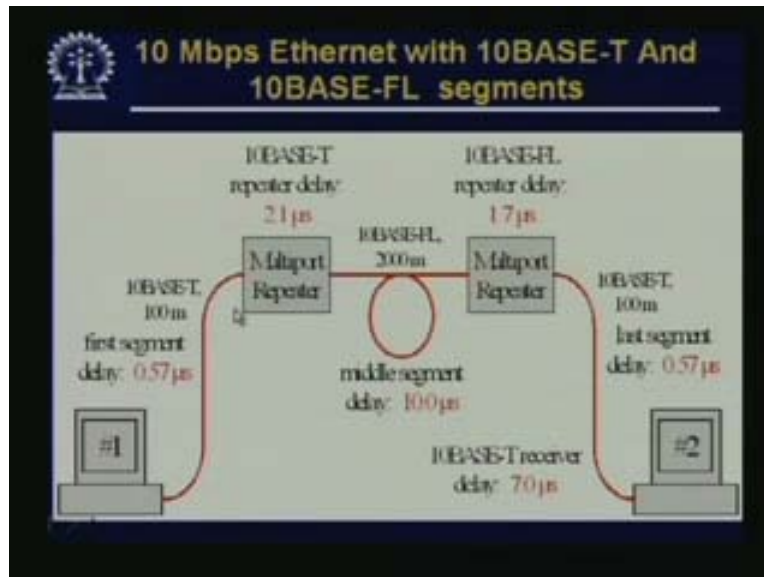
(Refer Slide Time: 13:25)



Now, we come to 10BaseT; so far as the physical part is concerned, this 10BaseT can go on to become 100BaseT; that means, from 10 Mbps it can go to 100Mbps. But for this UTP, unshielded twisted pair cable of various categories like 3, 4, 5, etc., are the dominant physical medium and the varying technology today for local area networks.

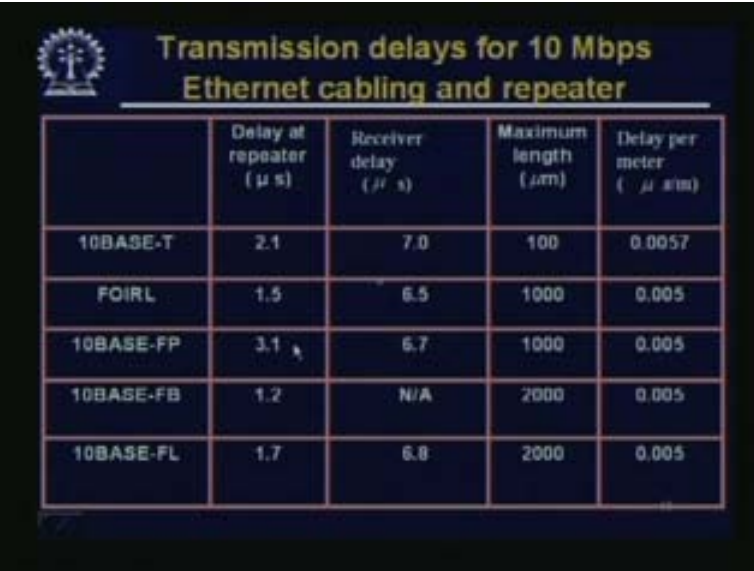
We can have category 5 cable data transmission up to 100 Mbps and we can go a little bit more on this depending on the distance. The impedance of a UTP cable is about 100 ohm, give or take another 15 ohm over the useful frequency range of the cable. The transmission speed is 0.6 to 0.65 into c , which is about 210^8 m/s. Time delay is about 5 nanoseconds /m.

(Refer Slide Time: 14:42)



Here, you can see there are some 10BaseT repeaters and then 10Base-FL segments. This is just to give you an idea about the various delays which are involved. The first segment delay is about 0.57 μ s for this 100-m segment, middle segment delay is about 100 μ s and last segment delay of about 0.57 μ s. The transceiver delay is about 7 μ s, FL repeater delay is about 1.7 μ s, and 10BaseT delay is about 21 μ s and so on. All these delays really add up and you have some limit on this in order for your CSMA/CD to work properly.

(Refer Slide Time: 15:41)

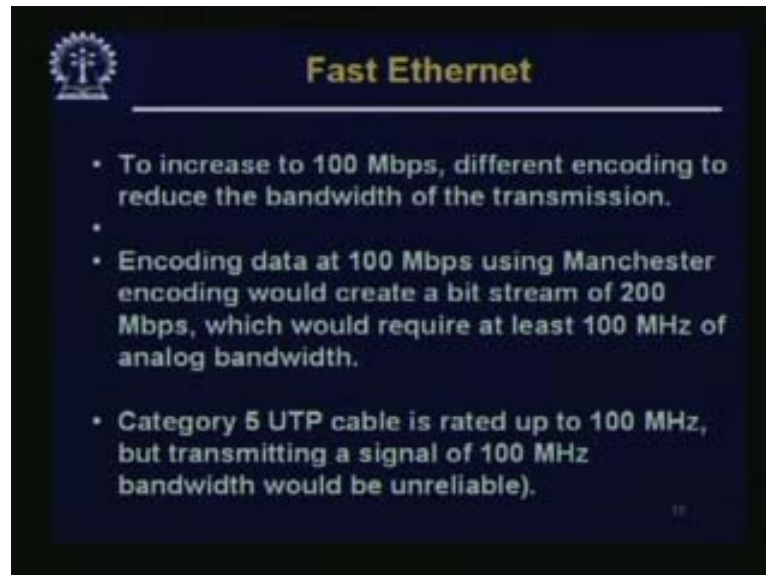


	Delay at repeater (μ s)	Receiver delay (μ s)	Maximum length (μ m)	Delay per meter (μ s/m)
10BASE-T	2.1	7.0	100	0.0057
FOIRL	1.5	6.5	1000	0.005
10BASE-FP	3.1	6.7	1000	0.005
10BASE-FB	1.2	N/A	2000	0.005
10BASE-FL	1.7	6.8	2000	0.005

So this is the same data in tabular form. It has various types of cables and these are all fibre optic kind of cables, which go up to 2 km or 1 km. These are obsolete now – these are the maximum lengths and these are the delays, the receiver delay and the delay at the repeater and so on. Now, what happened was that the people moved as the demand for bandwidth grew; so at least in the LAN it moved much faster when LAN bandwidth became quite cheap and everybody wanted to move from 10 Mbps to the next speed, which is 100 Mbps, and this was called fast Ethernet? One thing was that, when we come to fast Ethernet, this co-axial cable really was out. So the thing preferred was the unshielded twisted pair or UTP. Category 5 is preferred for 100 Mbps operation, although category 4 would also work.

It is possible to transmit it over category 4 cables also; but category 5 cables are preferred. Now one thing about this UTP cables is that because they are used for connection from end to end unlike the coaxial cables – in coaxial cables, we have one cable and many nodes may be connected at intermediate points – you put a connector, usually an RJ45 kind of connector on the two ends and then go into two nodes. It is as simple as that. So if there are no nodes in between, just two end point connections, where this part of the medium is not shared, at one end may be the computer, and the other end might go into another machine straight away. These two machines network together. Usually the other end would go to a network device. Let us say, to start with, it might go into a 100 Mbps hub. So, all the nodes which are connecting, that means which are part of the network, all of them are connected by this UTP cable to the central hub. So the central hub is a shared medium at the collision domain. But so far as the UTP part is concerned, there is no collision.

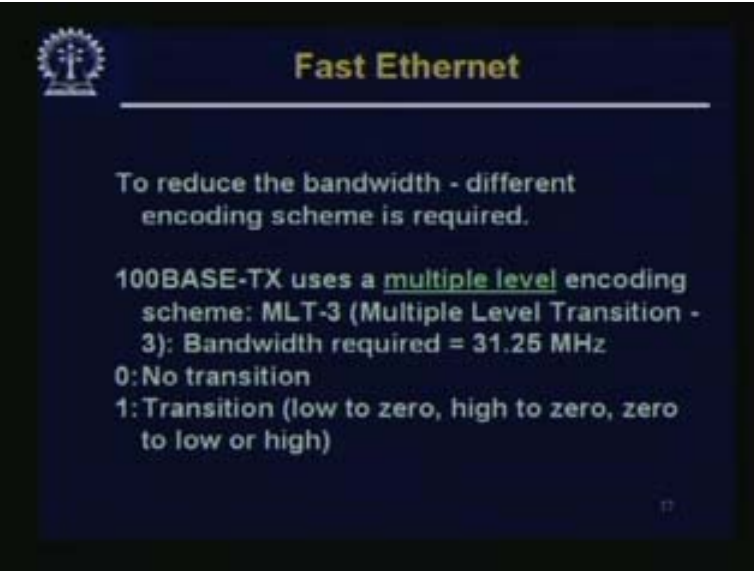
(Refer Slide Time: 19:16)



Just as we had the 10Base-T, which supported 10 Mbps, may be with the same kind of cable with this UTP cable, we could support fast Ethernet, that is, 100Base-T. There was 100Base-TX, But there are some problems with this: one is to increase to 100 Mbps, we have got one order of magnitude jump in the speed. That is not a mean achievement on the same medium. Depending on the quality of the medium, as the speed goes up; its performance tends to decrease. So depending on what kind of medium it is, there is only a certain amount of bandwidth in a certain range that you can support on that medium. So to increase to 100 Mbps, different encoding is used to reduce the bandwidth of the transmission.

Encoding data at 100 Mbps using Manchester encoding would create a bit stream of 200 Mbps. So, we cannot straight away use Manchester encoding here, because this is for supporting 200 Mbps, we requires at least 100 MHz of analog bandwidth, which is too much. Category 5 UTP cable is rated up to 100 MHz but for keeping some margin etc., transmitting a signal of 100 MHz bandwidth would be unreliable. So what we do is that instead of Manchester encoding, we use a different kind encoding, which is called multiple or multilevel encoding. To reduce the bandwidth, a different encoding scheme is required. 100Base-TX uses a multiple level encoding scheme, which is MLT3, Multiple Level Transition 3. The bandwidth required in this case has come down quite sharply to 31.25 MHz. The encoding is somewhat like this – 0: there is no transition and 1: there is a transition. And this transition could be of various types from low to 0, high to 0, 0 to low or 0 to high, depending on the context.

(Refer Slide Time: 20:57)



Fast Ethernet

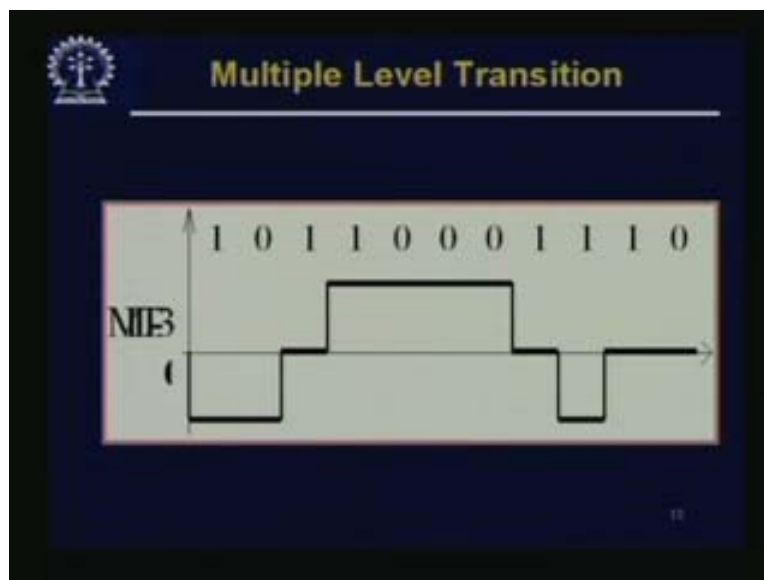
To reduce the bandwidth - different encoding scheme is required.

100BASE-TX uses a multiple level encoding scheme: MLT-3 (Multiple Level Transition - 3): Bandwidth required = 31.25 MHz

0: No transition
1: Transition (low to zero, high to zero, zero to low or high)

17

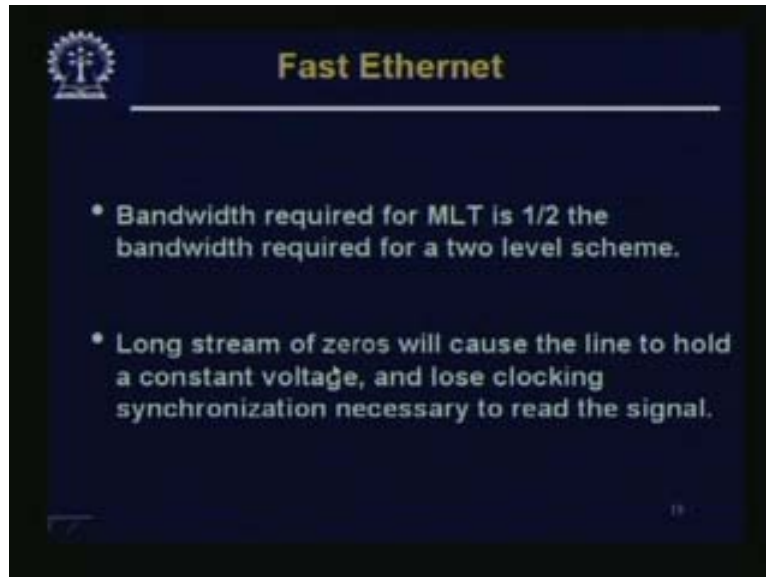
(Refer Slide Time: 21:33)



For example, this is a 10110001110 that is being encoded. Suppose with the 1 there is a transition from 0 to low, that is, 0 to low level, then for 0 there is no transition, after 0 there is 1. Now there is a transition; since we go to low now, and the only place to go to is high, we go to a high and then there is another 1 and since we are going high, we go high some more, that is, another 1. Should this go from 0 to high, then three 0s, that is, no transition at all. Then from 0 to 1, now from high since 1 has come, you will have to come to low and again another one has come; we are going in the downward direction. So we continue from 0 to low and then, another 1 from low to high and then 0; there is no transition. So, this is the way an MLT encodes using different levels or multiple levels of signals.

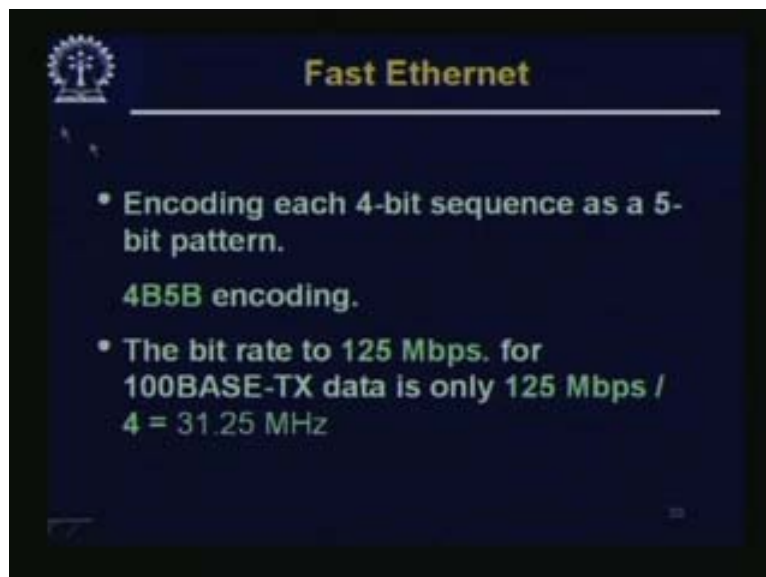
This reduces the net analog bandwidth requirement for the cable. So this is how, the ten-fold increase in the speed was achieved on the same physical infrastructure. So bandwidth required for MLT is half the bandwidth required for a two-level scheme.

(Refer Slide Time: 22:48)



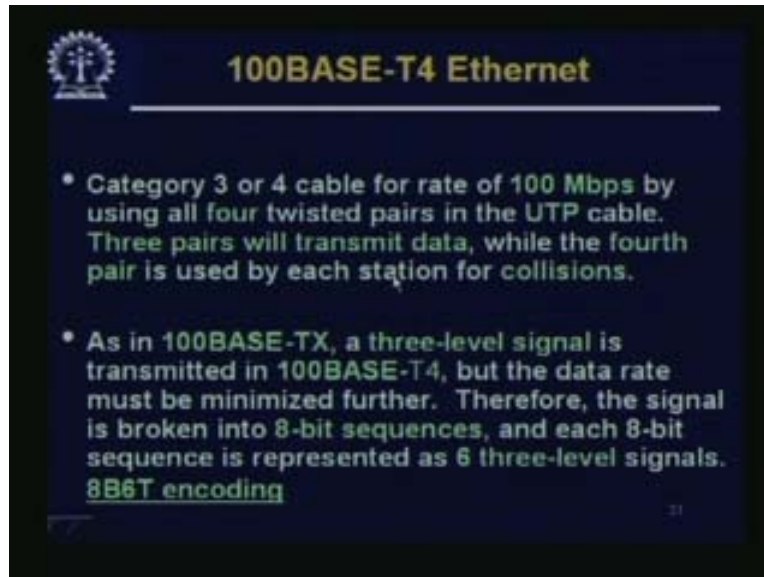
Half is a rough figure. Long stream of zeros will cause the line to hold a constant voltage and lose clocking, synchronization is necessary to read the signal. This is one problem, which is handled in various ways.

(Refer Slide Time: 23:27)




One way is to encode each 4-bit sequence as a 5-bit pattern 4B5B encoding. In this 4B5B encoding what is done is that instead of a 4-bit sequence we introduce a 5-bit sequence so we get some transition and we can hold on to synchronization.

(Refer Slide Time: 24:03)



This pushes up the bit rate to 125Mbps for 100BASE-TX, where data is only 125Mbps/4 is equal to 31.25 MHz. Category 3 or 4 cable is used for a rate of 100 Mbps by using all four twisted pairs in the UTP cable; so this is also possible. Three pairs will transmit data while the fourth pair is used by each station for collisions. As in 10BASE-TX, a three-level signal is transmitted in 100BASE-T4. T4 is when we are using category 4 cable but the data rate must be minimised further. Therefore the signal is broken into 8-bit sequence and each 8-bit sequence is represented as 6 three level signals, or 8/6t. If we are trying to put in a new network today we are never going to use category 4 UTP, we are always use category 5 or better. You know about the increasing requirement of speed in the LAN, but these techniques were developed so that Ethernet could move into some existing physical infrastructure like wiring, where category 4 cables is used for moving from 10 to 100 Mbps; that is why these were developed.

(Refer Slide Time: 25:19)




Fast Ethernet

The original fast Ethernet cabling.

Name	Cable	Max. segment	Advantages
100Base-T4	Twisted pair	100 m	Uses category 3 UTP
100Base-TX	Twisted pair	100 m	Full duplex at 100 Mbps
100Base-FX	Fiber optics	2000 m	Full duplex at 100 Mbps; long runs

Fast Ethernet is done with cabling of the following types – 100BASE-T4 with category 5 100BASE-TX; and 100BASE-FX is quite common. As I said, 10BASE-FI has become outdated but 100BASE-FX is still there, which uses fibre optics and goes up to 2 km. Category 5 UTP goes only up to 100 m or may be some people are conservative and limit it to 75 m. At full duplex at 100Mbps, fibre is again full duplex at 100 Mbps; these are the advantages. The advantage of T4 is that, it uses existing category 3 UTP.

(Refer Slide Time: 25:58)

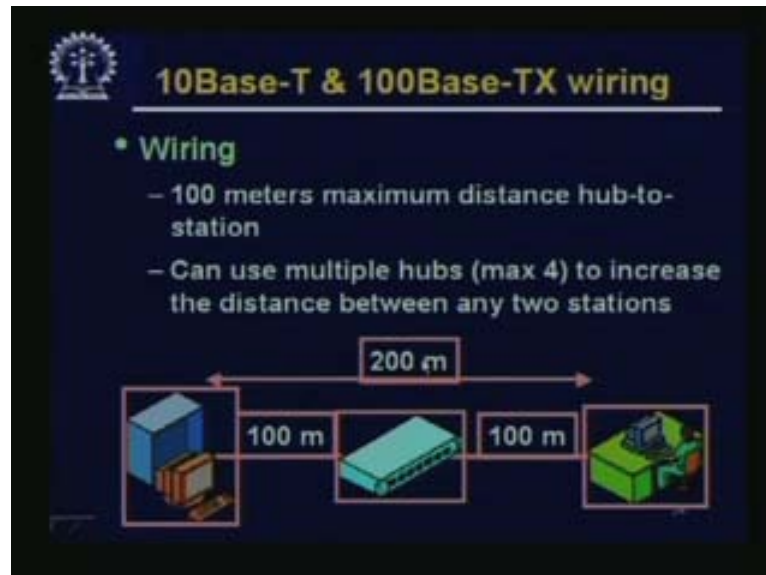


Characteristics of a 100BASE-T Ethernet system

100BASE-TX	delay per meter	0.0057	/m
	maximum length	100	m
100BASE-FX	delay per meter	0.0050	/m
	maximum length	412	m
Class I Repeater Delay		1.40	s
Class II Repeater Delay		0.92	s
Receiver Delay		0.50	s

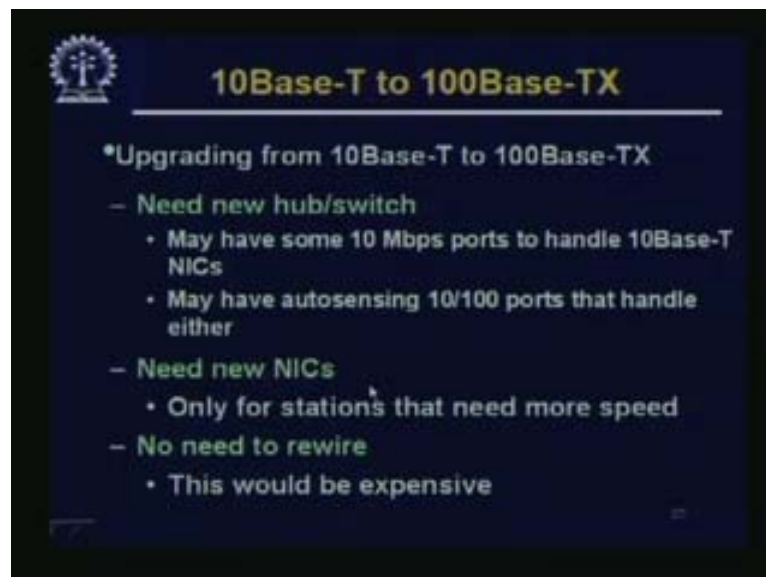
23

(Refer Slide Time: 26:09)



These are the repeater delays. 10BASE-T and 100BASE-TX wiring goes like this – from the end station a maximum distance of 100 m to the hub. We can use multiple hubs, a maximum of four, to increase the distance between any two stations; a hub may go to another hub; that may go to a station again. Now, there are three segments. We can use a maximum of four hubs and five segments with 100 m from the node to the hub.

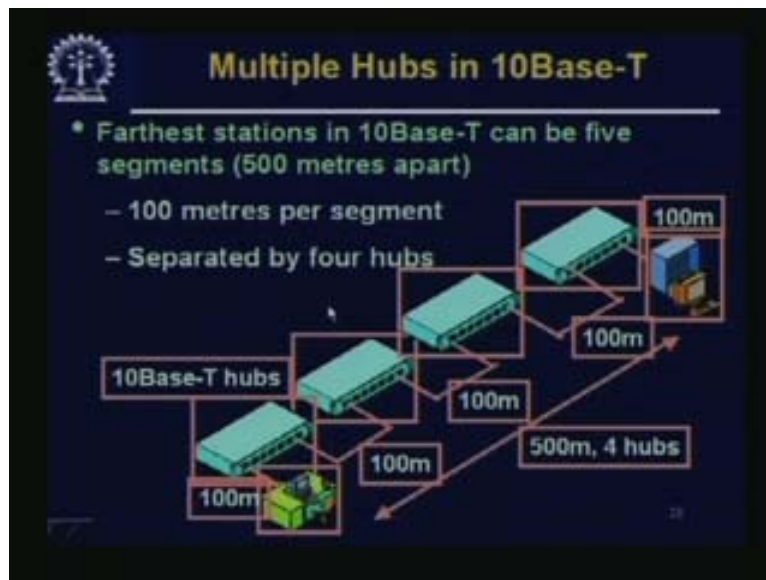
(Refer Slide Time: 26:49)



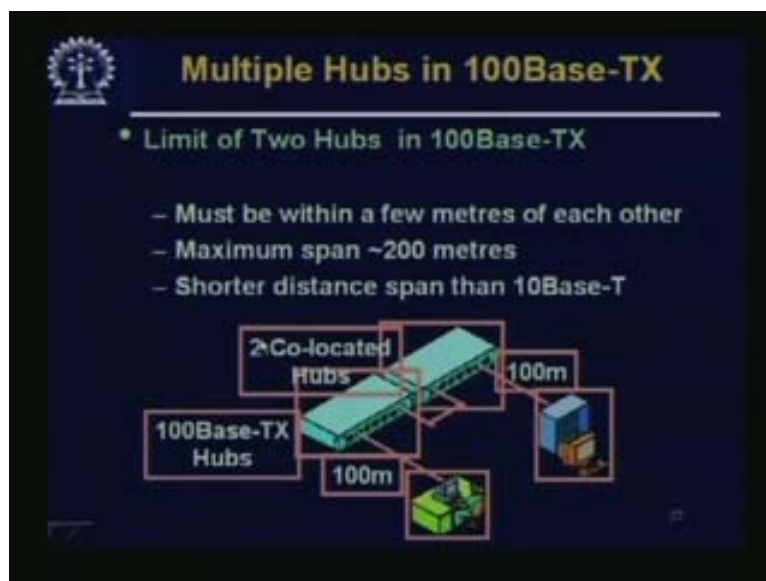
For upgrading from 10BASE-T to 100BASE-TX, we need new hubs or switches. You may have some 10 Mbps ports to handle 10BASE-T; NICs may have auto sensing 10/100 ports that handle either. Auto sensing ports are quite common.

As a matter of fact, most of the network interface cards you get today are 10/100; that means they are auto sensing ports. An NIC senses the network device at the other end. If the other end supports a rate of 100 Mbps, this NIC will operate at 100 Mbps. If it senses that the opposite end can handle only 10 Mbps, it will use only 10 Mbps. By the way, if there are problems with the varying connection, etc., it may send 100 Mbps port as 10 Mbps. So, if you want to upgrade from 10BASE-T to 100BASE-TX, you may need new NICs, and only for stations that need more speed and no need to rewire – that may be a big advantage in some cases. We may have 100 multiple hubs, one connecting to the other; you can have multiple hubs for connecting etc.

(Refer Slide Time: 27:50)

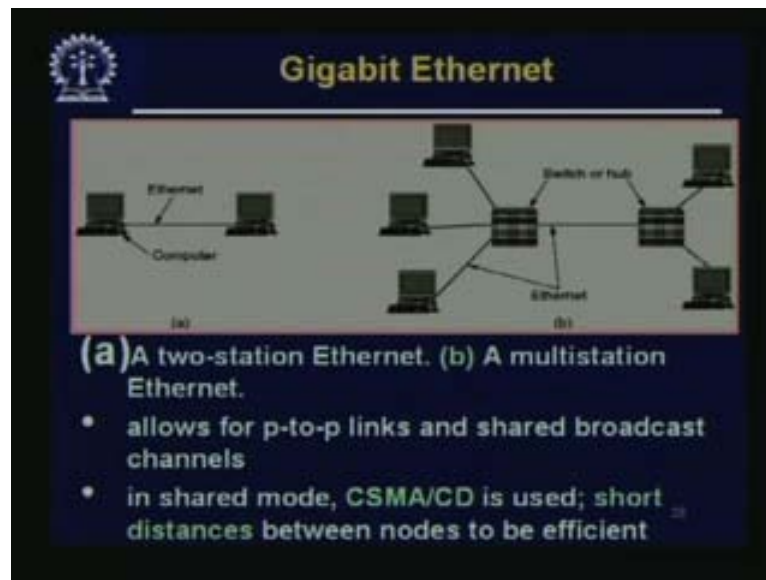


(Refer Slide Time: 27:50)



In 100BASE-TX there is a limitation that you can have only maximum two hubs and they must be within a few meters of each other. So, the maximum span using these hubs is only 200 m, which is shorter than 10BASE-T. It could be a big problem and that is why 100 Mbps hubs are also becoming obsolete. People have migrated from hubs to switches, and that is what we will be talking about. So if you had used hubs, this would have been a limitation; but nowadays people do not usually use hubs.

(Refer Slide Time: 29:19)



We will come back to the topic about hubs and switches later on. What happened is that this development went on and now we are moving into the era of Gigabit Ethernet. The Gigabit Ethernet in the desktop is still not very common, but it is coming and for servers etc., nowadays we routinely get Gigabit ports and, may be, soon your desktop computer may also be connected with Gigabit ports. So, this is a depiction of a two-station Ethernet – two computers connected straight away, this is a switch or a hub which operates at Gigabit speed. So, these allow point-to-point links and shared broadcast channels through some hubs if they are there. In shared mode, CSMA/CD is used for the short distance between the nodes to be efficient; but this is not very common. It uses standard Ethernet frame format 802.3z; the normal mode is full duplex mode with a central switch. This is the most common one connected to computers. In this configuration all lines are buffered.

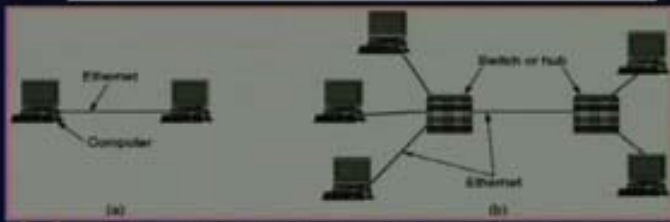
(Refer Slide Time: 30:06)

Gbit Ethernet

- Use standard Ethernet frame format (802.3z standard)
- The normal mode is full-duplex mode, with a central switch connected to computers. In this configuration, all lines are buffered.
- No sense the channel to see if it is idle because contention is impossible. CSMA/CD protocol is not used

(Refer Slide Time: 30:31)

Gigabit Ethernet

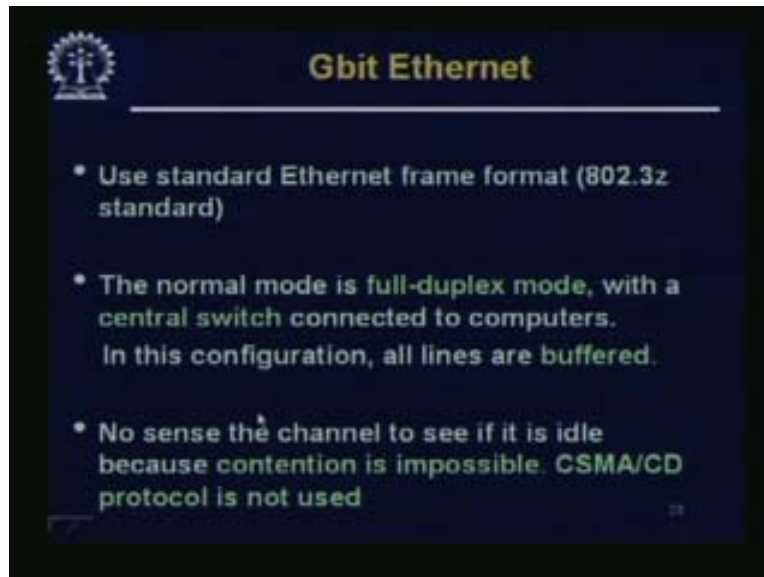


(a) A two-station Ethernet. (b) A multistation Ethernet.

- allows for p-to-p links and shared broadcast channels
- in shared mode, CSMA/CD is used; short distances between nodes to be efficient

On the right-hand side, you see one switch and the nodes are connected straight to the switch. So this is a switch rather than a hub, where each of the ports has a buffering.

(Refer Slide Time: 30:56)



Since each of the port is buffered, contention is impossible; because data can always and gets it on the buffer. So, it makes no sense to the channel to see if it is idle or not because if there is no contention, CSMA/CD protocol is not used at all. And this is the mode, in which Gigabit is usually used, although there are buffered distributors also.

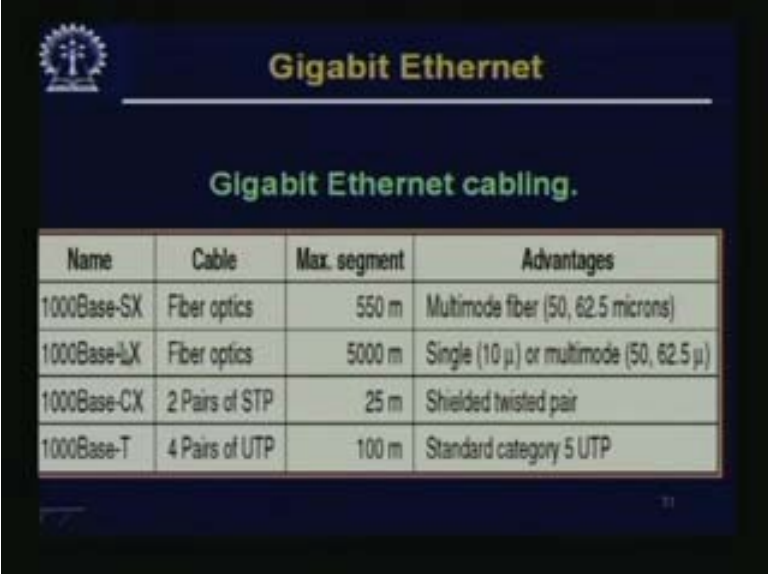
(Refer Slide Time: 31:11)



Now people are talking about moving to 10 gbps, at least in the metropolitan area network. They use copper and fibre with 850 nm or 1300 nm lasers and some kind of encoding.

The Ethernet is moving really fast, because it is upgrading quite fast – it went from 10 mbps to 100 mbps; from 100 mbps we are moving to 1gbps; and from 1gbps, we can see the possibility of moving to 10 gbps, and 10 gbps is already in operation today. From 10 mbps to 10 gbps is a 1000-fold increase in a very short span of time, which is really amazing.

(Refer Slide Time: 32:16)

A presentation slide titled "Gigabit Ethernet" with a subtitle "Gigabit Ethernet cabling." It features a table with four columns: Name, Cable, Max. segment, and Advantages. The table lists four types of Gigabit Ethernet: 1000Base-SX, 1000Base-LX, 1000Base-CX, and 1000Base-T, each with its corresponding cable type, maximum segment length, and advantages.

Name	Cable	Max. segment	Advantages
1000Base-SX	Fiber optics	550 m	Multimode fiber (50, 62.5 microns)
1000Base-LX	Fiber optics	5000 m	Single (10 μ) or multimode (50, 62.5 μ)
1000Base-CX	2 Pairs of STP	25 m	Shielded twisted pair
1000Base-T	4 Pairs of UTP	100 m	Standard category 5 UTP

For Gigabit Ethernet, we have 1000Base-SX; this SX shows the multimode fibre for 1000Base, so the maximum segment size would be about 550 m. 1000Base-LX uses single or multimode, so this goes up to 5000 m. Shielded twisted pair would only take to 25 m and 1000BaseT can go up to 100 m using category 5, category 5b or, for gigabit Ethernet, category 6 is preferred. We will talk more in detail about the switching, which is very important. So, Ethernet is usually considered as a shared media LAN.

(Refer Slide Time: 33:17)

Latency and Congestion with hubs

- **Ethernet is a shared media LAN**
 - Only one station can transmit at a time
 - Even in multi-hub LANs
 - Others must wait
 - This causes delay

The diagram illustrates a multi-hub LAN. A central hub is connected to several stations. One station is shown sending data, with a label 'One Station Sends'. A callout box points to the hub and other stations, stating 'All Other Stations Must Wait', highlighting the shared media nature of the LAN.

Only one station can be transmitted at a time even when you have multiple hubs, these hubs are all shared domain, which means that only one can transmit. All other stations must wait while one station sends. So this is the big limitation of latency and congestion with hubs

(Refer Slide Time: 33:39)

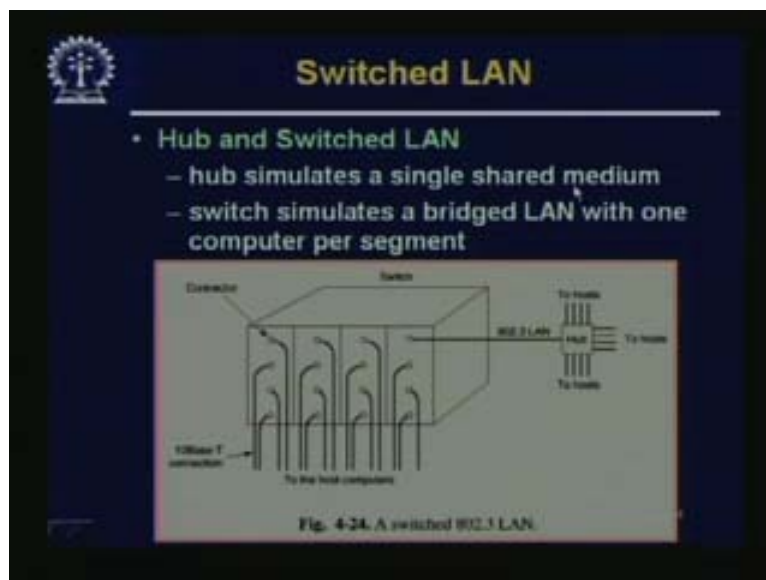
Ethernet Switch → Benefits

- **Improved security**
 - users are less able to tap-in into other user's data
- **Better management**
 - control who receives what information (i.e. Virtual LANs)
 - limit the impact of network problems
- **Full duplex**
 - rather than half duplex required for shared access

To improve this, people moved over from hubs to switches. Packet switches are very useful. They give you improved security. Users are less able to tap into other users' data. In a shared medium, we know everything is being broadcast so all packets are going to every other node. That is poor security; this gives you much better security; better management – we can control who receives what information.

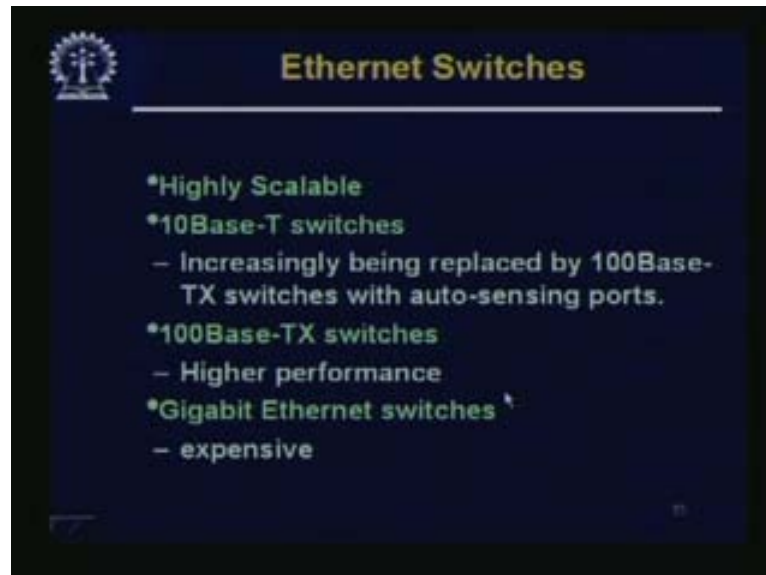
That means, even using a switch, even within a LAN, we can create a work group and make users members of different virtual LANs. These virtual LANs may be packet filtering between the virtual LANs, etc., to give a better management and limit the impact of network problems. This switch will also give you full duplex communication, rather than half duplex. One side effect of having a shared medium is that only one sends, and everybody else is in the receiving mode. Obviously that channel cannot work as a full duplex mode. Because two of them cannot send simultaneously at the same time, only one may be sending and the other may send. At most, this works as a half duplex channel, whereas over here this can work as full duplex channel. A hub simulates a single shared medium, whereas switch simulates a bridged LAN with one computer per segment.

(Refer Slide Time: 35:20)



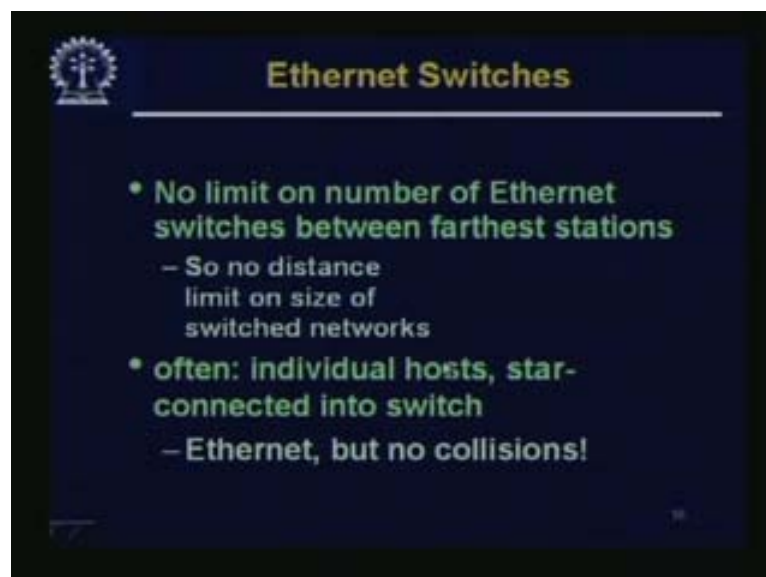
For example, we have only one computer and many segments, and there is only one computer for each segment. So it does not have any collision and from two different computers you can send data at the same time to the switch. The switch will buffer it and there will be no collision. So collision detection may be eliminated. Some of these may connect to a hub instead of going to a computer. So a switch may be connected to another switch, that switch may connect to hub, and that hub may connect to so many computers. So there are many possibilities.

(Refer Slide Time: 36:08)



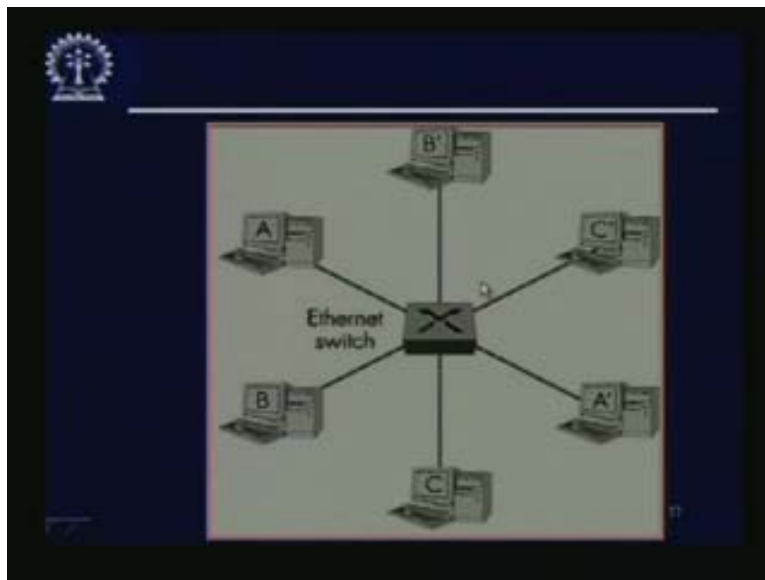
Ethernet switches are highly scalable. We had 10Base originally; we went from 10BaseT hubs to 10BaseT switches. 10BaseT switches are increasingly being replaced by 100Base-TX switches with auto sensing codes. 100Base-TX switches are very common these days, they have higher performance and their costs have come down. These days we get 100Base-TX switch at the cost of what a 10BaseT hub would cost a few years ago. We can see how the technology and the economy have improved. Gigabit Ethernet switches are still expensive; it will take some more time in order for them to move to the desktop.

(Refer Slide Time: 37:02)



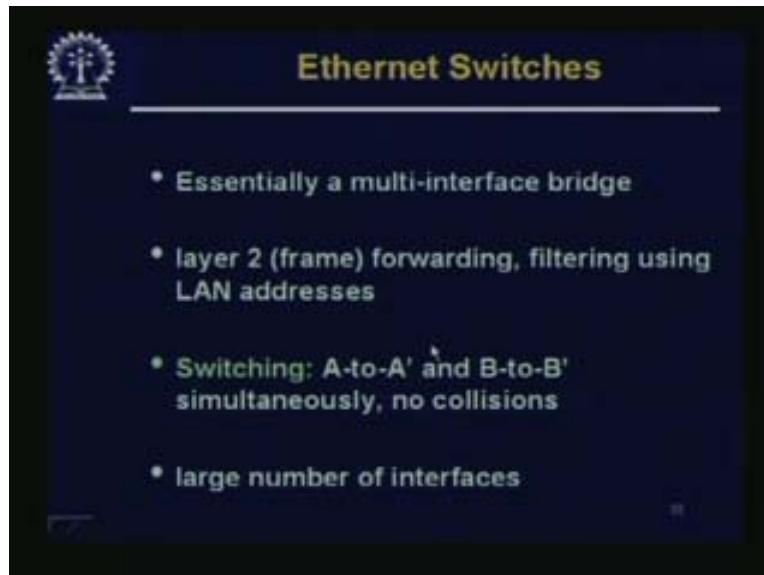
Other advantages of switches are that there are no limits on the number of Ethernet switches between the farthest stations and no distance limit on size of switched networks. You remember when we were talking about 100Base-TX and when we were using 100 base TX hubs. You can only have a maximum of two hubs and these two hubs have to be close together; they cannot go far long. So, the maximum span of a local area network could almost be 200 m, which was really a limitation, but then people moved to 100Base-TX. They also moved from hubs to switches. With these switches, now there is no problem because the distance limitation had to do with the round-trip propagation delay and things like that. Since there is no contention and the switches have buffers, there is no distance limitation at all. You can have any number of switches connected to each other; there is no distance limit on size of switched networks. So often individual hosts, are connected to switches and may be different switches are also connected. This is Ethernet without collisions or CSMA/CD.

(Refer Slide Time: 38:17)



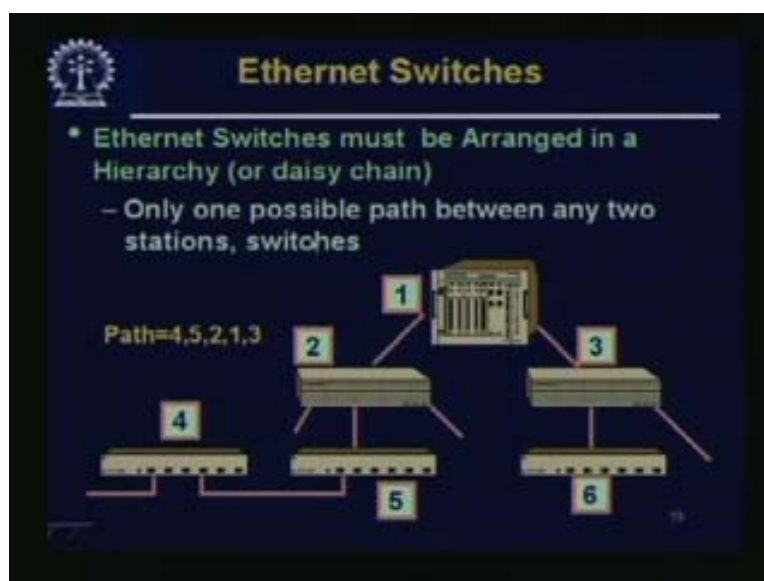
Here, we have A, B, C, and so on. Suppose A wants to communicate with A' and B wants to communicate with B', both the communications may take place simultaneously. This is a switch. If this had been a hub that would not be possible but since this is a switch, B can communicate to B' and A can communicate to A', just as in a telephone switch, two pairs or distinct pairs of people can communicate at the same time. Similarly, two distinct pairs of machines, through this packet switch, can communicate at the same time.

(Refer Slide Time: 38:56)



This is essentially a multi-interface bridge; it has a layer 2 frame forwarding, filtering using LAN addresses; LAN addresses mean MAC addresses. We will talk more about MAC addresses presently. It allows switching A to A' and B to B' simultaneously, but there are no collisions. This can have a large number of interfaces. A modern switch may have hundreds of ports, so hundreds of machines can be connected to the same switch. Similarly, a number of such switches can connect to each other.

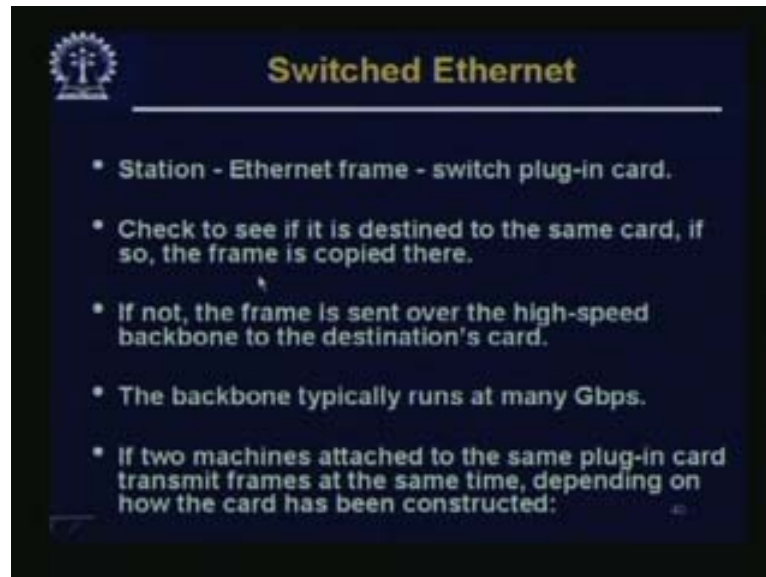
(Refer Slide Time: 39:34)



Ethernet switches must be arranged in a hierarchy or daisy chain with only one possible path between any two stations.

Suppose, from 4 to 3 there is only a single path that is 4, 2, 1, 3; you do not make a loop connecting two ports of a switch because that would really create a lot of problem. Therefore, Ethernet switches must be arranged in a hierarchy or daisy chain. There is only one possible path between any two stations or switches.

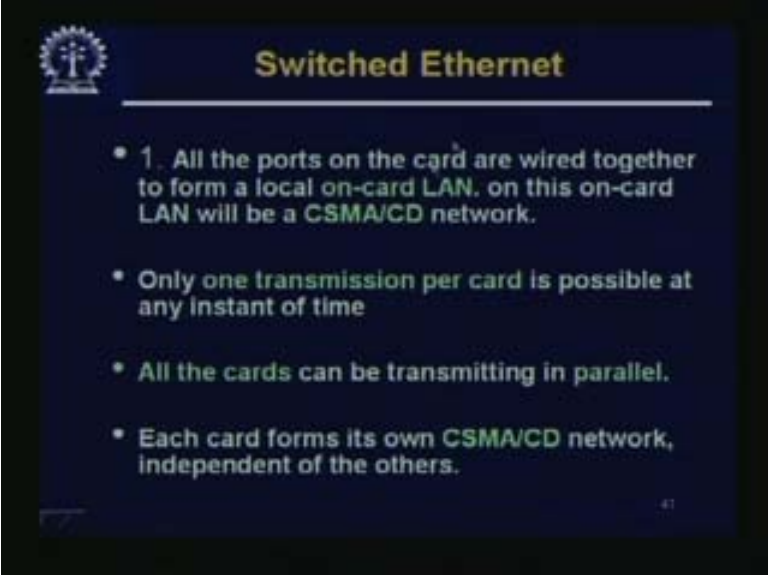
(Refer Slide Time: 40:05)



A station is an Ethernet frame switch plug-in card. It checks to see if it is destined to the same card. If so, the frame is copied there. We mentioned that in a packet switch, usually the ports would be grouped into small groups and each of these groups will go into a particular line card. This line card is similar to the Ethernet. A switch is nothing but a packet switch; there will be line cards over there and may be 4 lines coming to one particular line card and maybe there are 8 line cards for a 32-port switch. Now, for those that are coming into the line card, there are variations of how they are handled by the line card. It could be that a frame which has arrived in a switch plug in card is destined to the same card; that means it is destined to some other line in the same card. If so, the frame is simply copied there; if not, the frame is sent. Now these line cards should get connected to each other through the backplane.

If the frame is sent over from the high-speed backbone to the destinations card, it will travel to the other distinct destination through the switching fabric. This backplane may be active or passive; they are of various types. Anyway, the backbone typically runs at many gigabits per second. As a matter of fact, nowadays it is possible to have a high-end switch with something of the order of 100 gbps backplane switch. It is a very high-speed switch that can support a lot of machines at the same time. If two machines attached to the same plug in card transmit frames at the same time, what would happen? There are two machines attached to the same plug in a card and they transmit frames at the same time – what would happen would really depend on the way the card handles it; here there are some variations. For example, let us say, there are 4 line cards – 1, 2, 3, 4 – coming into this one particular line card. Line 1 and line 2 frames arrived simultaneously and both of them are destined to line 3; what is going to happen? It depends on how it is handled.

(Refer Slide Time: 42:54)




The slide is titled "Switched Ethernet" in yellow text on a dark blue background. In the top left corner, there is a small circular logo featuring a cross and a gear. The slide contains four bullet points, each starting with a green number and followed by green text. The first bullet point describes the on-card LAN configuration. The second and third bullet points discuss transmission constraints and parallelism. The fourth bullet point states that each card forms its own CSMA/CD network. A small number "41" is visible in the bottom right corner of the slide area.

- 1. All the ports on the card are wired together to form a local on-card LAN, on this on-card LAN will be a CSMA/CD network.
- Only one transmission per card is possible at any instant of time
- All the cards can be transmitting in parallel.
- Each card forms its own CSMA/CD network, independent of the others.

One possibility could be that all the ports on the card are wired together to form a local, on-card LAN. This LAN card itself is a local on-card LAN, and on this on-card LAN, there will be a CSMA/CD network. CSMA/CD may operate on that particular card itself; in that case, there will be a collision that will be detected. Only one transmission per card is possible at any instant of time, so if you are doing that, you have made that line card to be a shared medium. All the cards can be transmitting in parallel; the different cards may be transmitting in parallel. Each card forms its own CSMA/CD network independent of the others; that is one possibility.

(Refer Slide Time: 43:42)

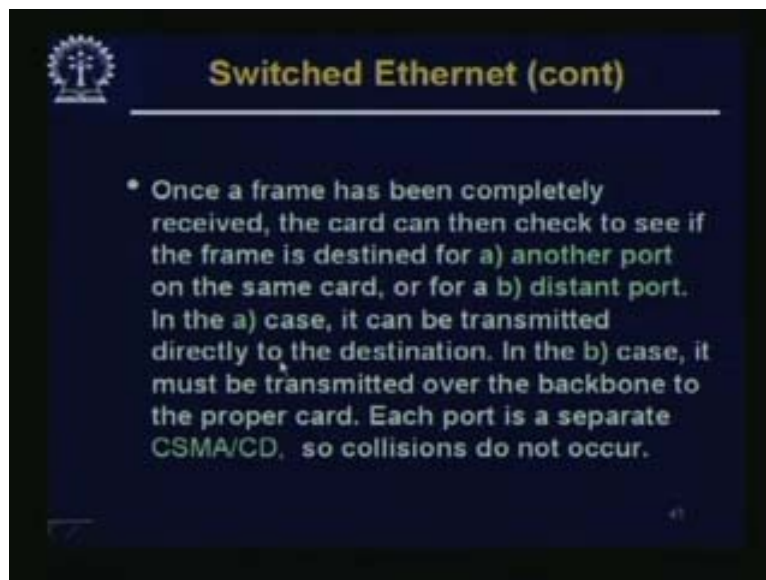


The slide is titled "Switched Ethernet (cont)" in yellow text on a dark blue background. In the top left corner, there is a small circular logo featuring a cross and a gear. The slide contains one bullet point, starting with a green number and followed by green text. The bullet point describes how a plug-in card buffers incoming frames to allow for parallel, full-duplex operation, which is not possible with CSMA/CD on a single channel. A small number "42" is visible in the bottom right corner of the slide area.

- 2. plug-in card of each input port is buffered, so incoming frames are stored in the card's on-board. This allows all input ports to receive/transmit frames at the same time, for parallel, full-duplex operation, this is not possible with CSMA/CD on a single channel.

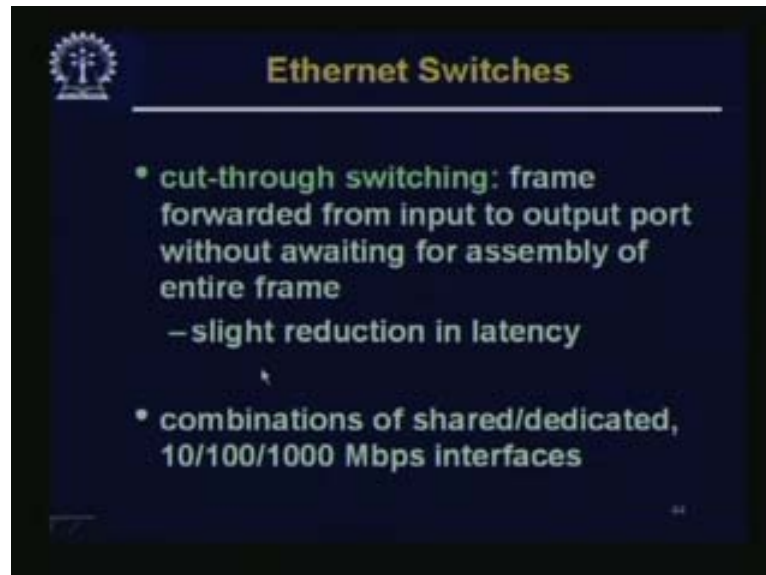
The other possibility is that the plug-in card of each input port is buffered; so incoming frames are stored in the cards on board. This allows all input ports to receive or transmit frames at the same time for a parallel, full duplex operation. This is not possible with CSMA/CD on a single channel. Previously we had talked about buffers in packet switches. We mentioned how buffers could be on input side, how buffers could be on output side and distributed throughout the switch fabric. That was from the point of view of handling head of line, blocking, switching, speed etc. Here, it is slightly different in the sense that if there is some small buffer at each of the ports and on input side, full duplex transmission at the same time is possible on all the lines. Instead of making a local CSMA/CD and allowing some kind of collision, this naturally is a better kind of switch. Once a frame has been completely received, the card will then check to see the frame is destined for another port on the same card, or for a distant port. In the A case, it can be directed to the destination; in the B case it must be transmitted over the backbone to the proper card.

(Refer Slide Time: 45:12)



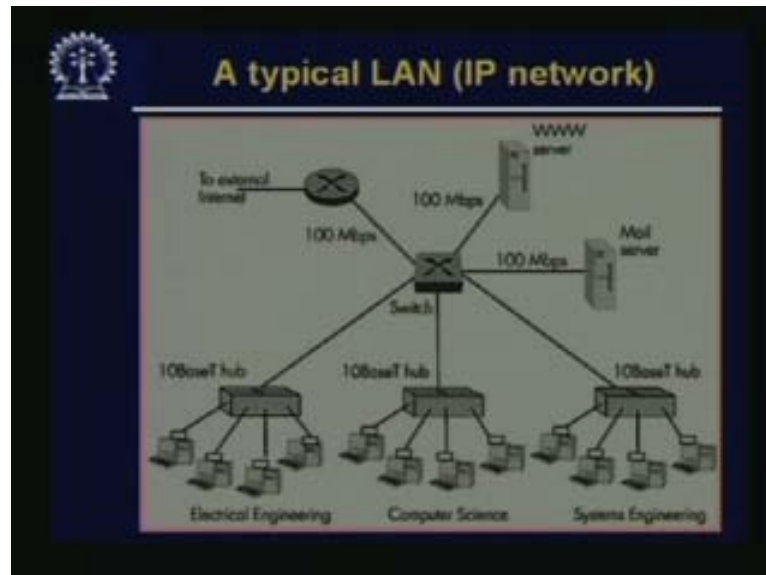
Each port has a separate CSMA/CD so collisions do not occur. There is now no question of any CSMA/CD because each port has only one machine. Then there is another term called cut through switching: frame forwarded from input to output port without waiting for assembly of the entire frame. That means even when the frame has not fully come in, the bits are transmitted and there is a slight reduction in latency. So, combinations of shared dedicated 10/100/1000 mbps interfaces are possible.

(Refer Slide Time: 45:23)



On the same switch you might get some 10 mbps ports, some 100 mbps ports, also some 1000 mbps or Giga bit ports for up linking to the main backbone.


(Refer Slide Time: 46:03)



This is a typical LAN of an IP network. Let us look at this figure – some hubs are remaining, but these hubs will give way to switches from 10BaseT to 100BaseT very soon, but let us say, there is a legacy system, through which some 10BaseT hubs and some computers are connected to this. These hubs may be connected to a switch. This switch could be connected to another switch and that switch might connect to some other nodes, etc.

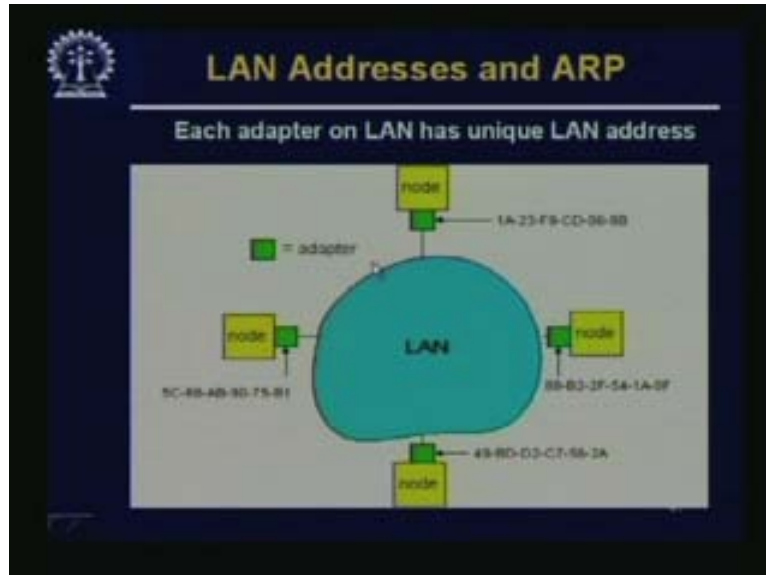
This constitutes the LAN – maybe there is a 100 mbps connection from this switch to the mail server, because mail server may be something which everybody is using. This is a server, and you might want to have a higher speed of this one also – these are 10 mbps ports. You might want to make it 100 mbps. There is a www server, that is, a web server, which is also connected through a 100 mbps port. So from the same switch, you may connect to some hubs, you may connect to some servers, you may connect to some individual desktops, PCs or machines. So this is what the LAN part looks like. And of course, nowadays nobody would want a LAN which is standing by itself. People want to connect their LANs to the internet, which is the great network of networks. There are probably lot of networks, may be millions of networks, in other places; we want to connect to other networks also. So we have to go through what is known as a router. We go through a router to the external internet. We will be talking about this part when we talk about the network layer in more detail; so this is a typical picture of a LAN.

(Refer Slide Time: 48:05)




If we abstract this a little bit, we have a LAN, and different nodes or the different machines are connected to this LAN. And then there is a router, which connects you to the internet. This internet may be a connection between a lot of different networks, may be a lot of different LANs to their own routers etc. This is a slightly more abstract view. And then looking into it in detail in the LAN part of it, if we go to the previous picture, this is the LAN part of it and these nodes are connected.

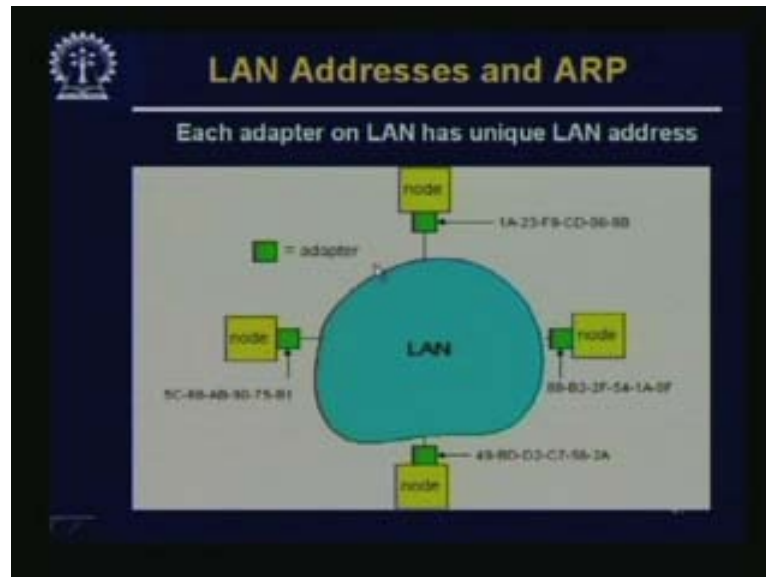
(Refer Slide Time: 46:35)



(Refer Slide Time: 48:40)



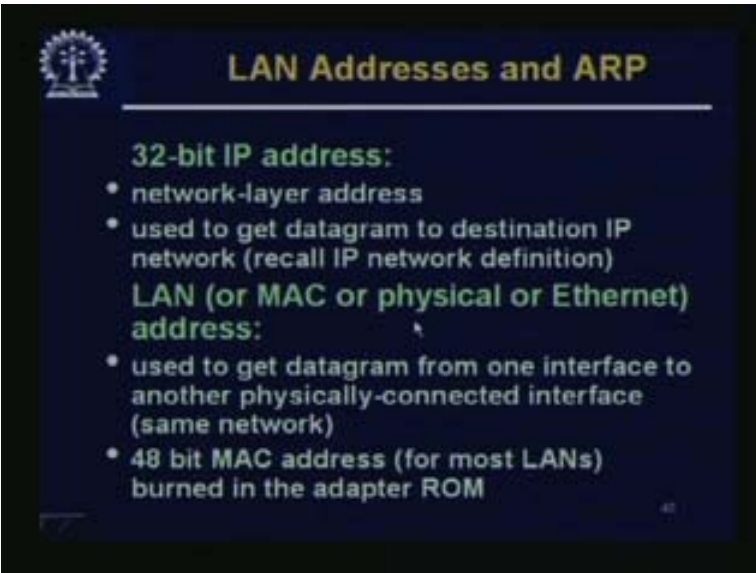
(Refer Slide Time: 48:45)



These nodes are usually connected through adapters. That means the network interface card, NIC, or the adapter. These nodes may communicate with each other in the LAN if it is connected to the internet. Naturally it may communicate with the other machines in the outside world also; but let us for the time being just focus our attention on the nodes which are in the LAN, and they are sort of trying to talk to each other. If they want to talk to each other using either a hub or a switch, or in this case, the Ethernet frame format, what is going to happen is that the frame format has certain fields and those fields have to be filled up. That means the destination address has to be there, and the destination address is the LAN address or the MAC address or the hardware address. Sometimes this is also called an Ethernet address.

And this Ethernet address is distributed by IEEE; the manufacturers of these adapters buy blocks of addresses from IEEE. And these are all 6-byte addresses, something like this: 1a23f9, each of them is 1 byte. And if you look at this figure, say 1a23f9cd069b and 88b22a do not have any relationship with each other. They are just 6 bytes, may be 6 arbitrary bytes. This card may have come from some manufacturer and got entirely different blocks from IEEE, the other may have an entirely different address, which has no relationship with the other; but they are all sharing this LAN. There is a small problem – how does the source node get to know the destination address? One thing is sure – since one central agency, namely IEEE, gives this Ethernet addresses, each adapter on LAN has a unique LAN address; no two adapters will ever have the same address. Sometimes, there may be some nodes with a number of LAN addresses, with a number of adapters etc., but one particular adapter will have one unique address, two addresses will never be the same. The point is that, how do get to know this is the source? This is a destination – how does the source know the destination, LAN address or the MAC address?

(Refer Slide Time: 51:35)



LAN Addresses and ARP

32-bit IP address:

- network-layer address
- used to get datagram to destination IP network (recall IP network definition)

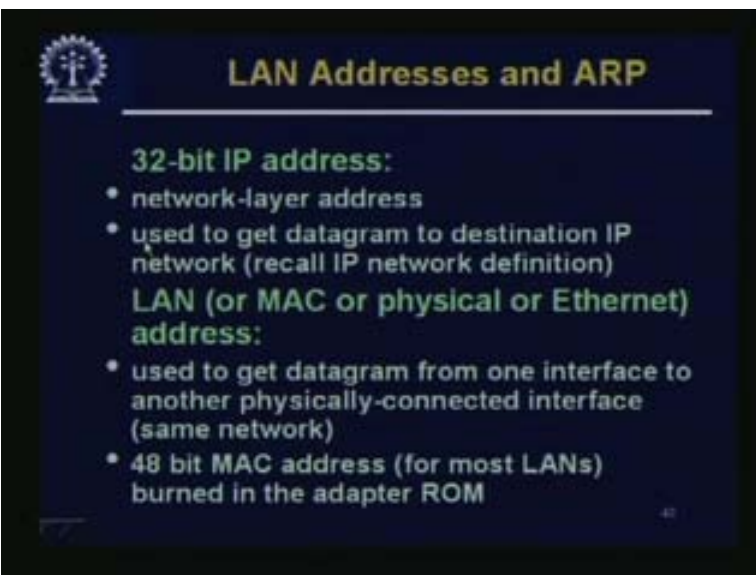
LAN (or MAC or physical or Ethernet) address:

- used to get datagram from one interface to another physically-connected interface (same network)
- 48 bit MAC address (for most LANs) burned in the adapter ROM

40

There are at least two kinds of addresses, which are used in networks. One is the IP address, which is used by the network layer to decide on the total route from the source to the destination. What will happen is that it will know where it will have to go; in order to reach its destination address; it must have some kind of relationship with geography. And then, there are these hardware or MAC addresses, which are at a lower level and are just as unique; they do not have any relationship with each other. So we have this 32-bit IP address, network layer addresses used to get datagram to the destination IP.

(Refer Slide Time: 52:32)



LAN Addresses and ARP

32-bit IP address:

- network-layer address
- used to get datagram to destination IP network (recall IP network definition)

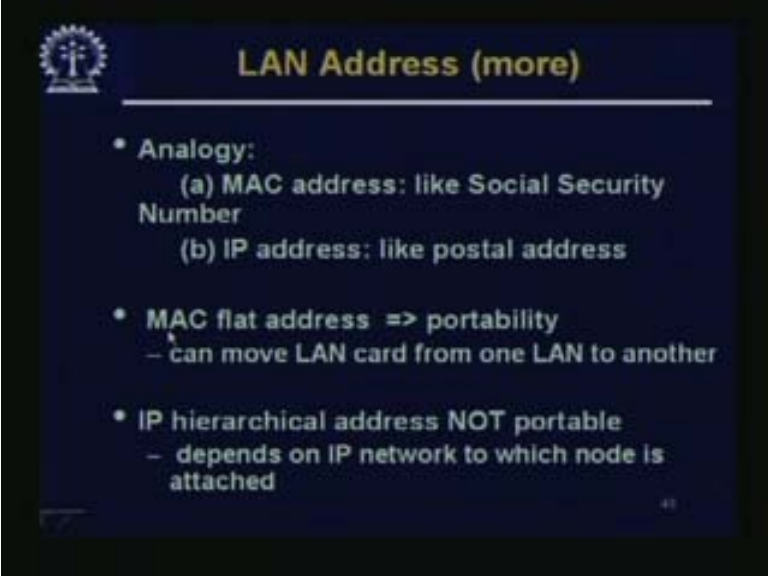
LAN (or MAC or physical or Ethernet) address:

- used to get datagram from one interface to another physically-connected interface (same network)
- 48 bit MAC address (for most LANs) burned in the adapter ROM

40

We have LAN, MAC, physical, Ethernet or hardware address used to get datagram from one interface to another physically connected interface in the same network. So this is a 48-bit MAC address for most LANs burned in the adapter ROM.

(Refer Slide Time: 52:51)

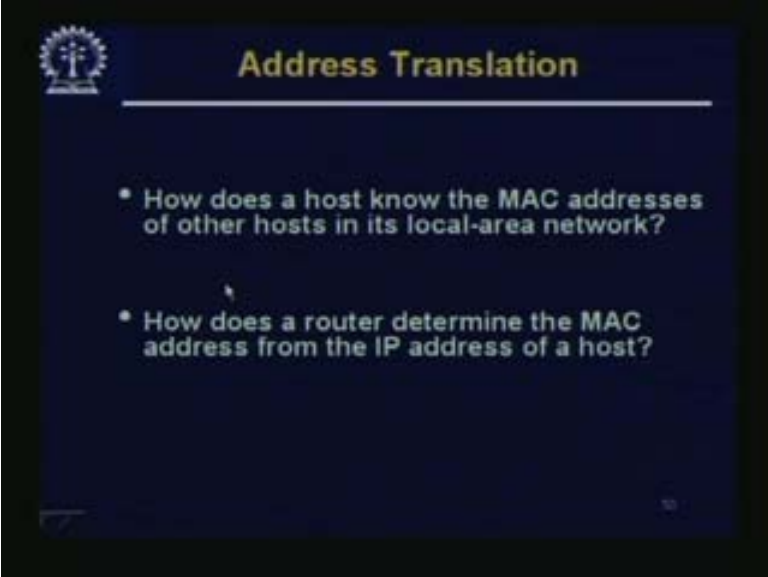


LAN Address (more)

- Analogy:
 - (a) MAC address: like Social Security Number
 - (b) IP address: like postal address
- MAC flat address => portability
 - can move LAN card from one LAN to another
- IP hierarchical address NOT portable
 - depends on IP network to which node is attached

40

(Refer Slide Time: 53:02)



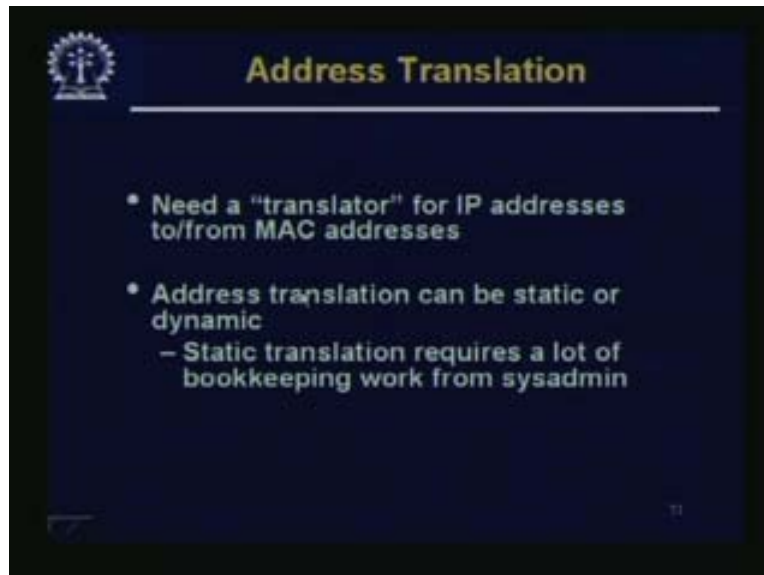
Address Translation

- How does a host know the MAC addresses of other hosts in its local-area network?
- How does a router determine the MAC address from the IP address of a host?

41

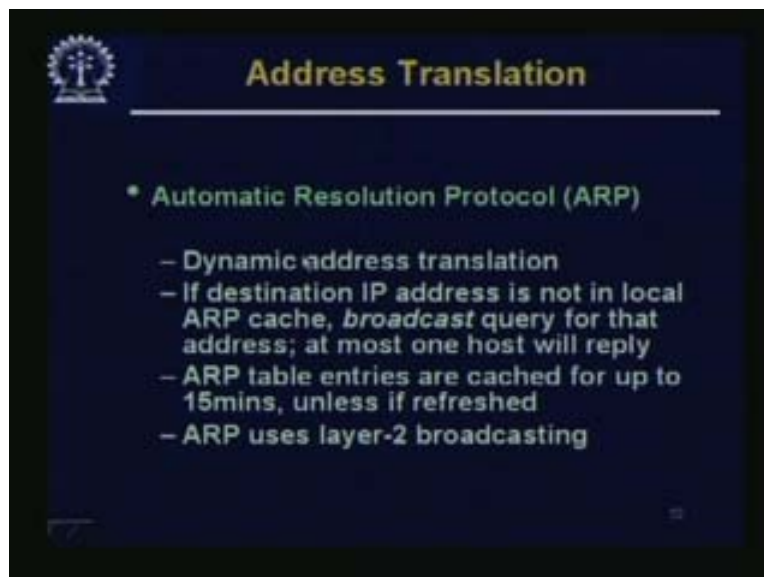
Analogy of MAC address is like a social security number; IP address is like postal address. The question is how do you get to know the other person's address? Address translation: how does a host know the MAC address of the other host in its local area network? And secondly, how does a router determine the MAC address from the IP address of a host? These are two interrelated questions.

(Refer Slide Time: 53:17)



We need some kind of a translator for IP addresses from MAC addresses. Address translation can be static or dynamic; static translation means you keep some static table. Static translation requires a lot of book-keeping work from system administration and this is not feasible.

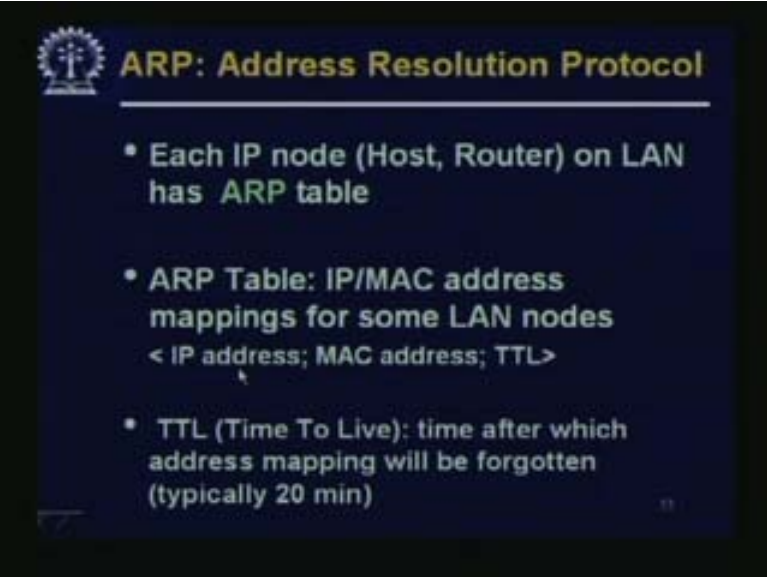
(Refer Slide Time: 53:37)



We have this address resolution protocol or automatic resolution protocol, the so-called ARP protocol, which is a dynamic address translation scheme. If the destination IP address is not in the local ARP cache – ARP cache is a table, which is stored in the local machine and which gives this IP address to MAC address mapping (the broadcast query for that address) – at the most one host will reply because the IP address will match with only one of the host addresses.

So, the ARP table entries are cached for up to 15 minutes, unless if refreshed. It is firmly at 15 minutes so that a machine can go from one LAN and can also be shifted to another LAN. What will happen is that these old entries must automatically wash out. And if the new entries communicate to their communicating machine, their MAC addresses are going to come into the cache. So this is how it is always kept fresh. ARP uses layer 2 broadcasting

(Refer Slide Time: 54:40)

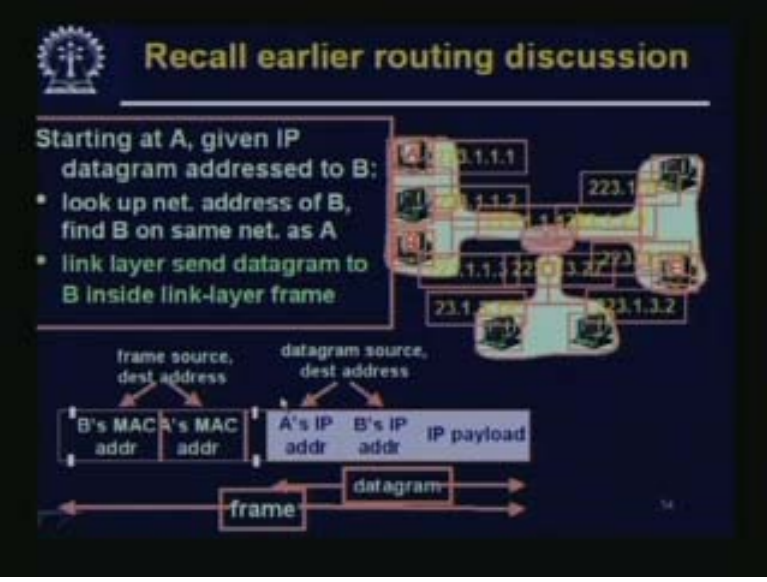


ARP: Address Resolution Protocol

- Each IP node (Host, Router) on LAN has ARP table
- ARP Table: IP/MAC address mappings for some LAN nodes
< IP address; MAC address; TTL >
- TTL (Time To Live): time after which address mapping will be forgotten (typically 20 min)

Each IP node on LAN has an ARP table. ARP table is IP, MAC address mapping for some LAN nodes. And there is a time to leave, something like 15 minutes or so, it depends on the time after which the address mapping will be forgotten – it is typically around 15, 20 or 25 minutes.

(Refer Slide Time: 55:00)



Recall earlier routing discussion

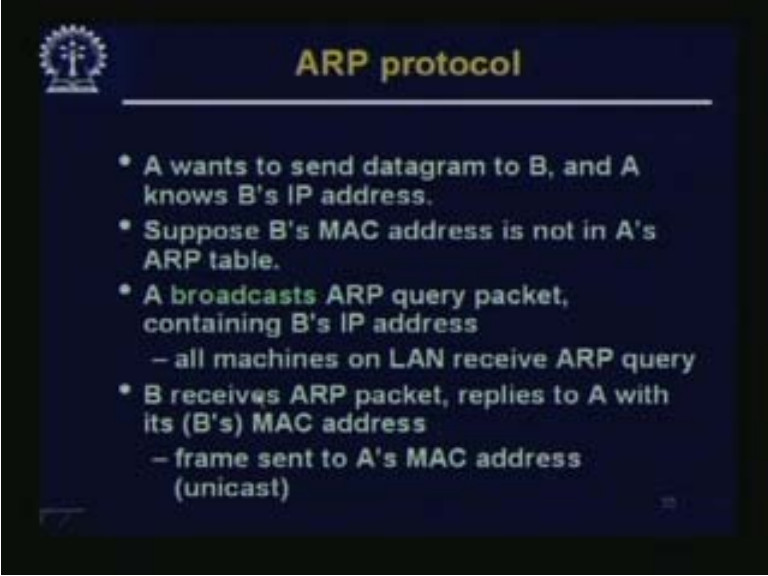
Starting at A, given IP datagram addressed to B:

- look up net. address of B, find B on same net. as A
- link layer send datagram to B inside link-layer frame

The diagram shows a network topology with nodes labeled with IP addresses: 1.1.1.1, 1.1.2, 223.1, 1.1.3, 223.1, 1.1.3.2, 223.1, 1.1.3.2, 223.1, 1.1.3.2. Below the diagram, a frame structure is shown with fields for frame source/dest address (B's MAC, A's MAC) and datagram source/dest address (A's IP, B's IP), followed by the IP payload.

So starting at A, we want to send a datagram addressed to B. We look up the net address of B or link layer. What will happen is that we want we do not need B's MAC address? That means this 2 23.1.1.1 is an IP address – we know it is B's IP address, but we want to know its MAC address. So A's IP address, B's IP address and some IP pay load form a frame like this, which is broadcast.

(Refer Slide Time: 55:37)

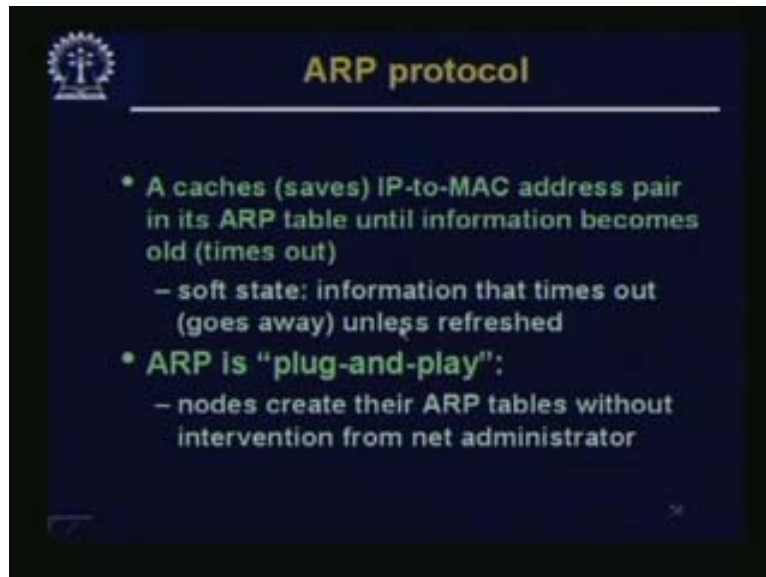


ARP protocol

- A wants to send datagram to B, and A knows B's IP address.
- Suppose B's MAC address is not in A's ARP table.
- A broadcasts ARP query packet, containing B's IP address
 - all machines on LAN receive ARP query
- B receives ARP packet, replies to A with its (B's) MAC address
 - frame sent to A's MAC address (unicast)

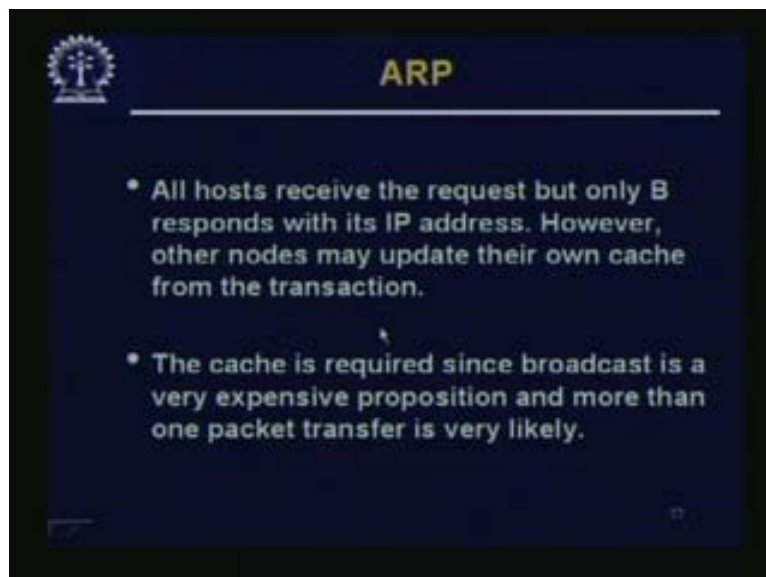
A wants to send datagram to B and A knows B's IP address. Suppose B's MAC address is not in the ARP table, A broadcasts an ARP query packet containing B's IP address. All machines on LAN receive the ARP query. B receives the ARP packet; replies to A with B's MAC address; the frame is sent to A's MAC address, which is unicast. It is now specifically given as MAC address is known, because A sent the query with a source address.

(Refer Slide Time: 56:09)



A cache saves IP to MAC address pair in its ARP table until information becomes old, that is, it times out. But if it is being used, then it will not become older. So in a soft state, information that time out goes away unless refreshed. ARP is plug_and_play, otherwise it will get washed out and new entries will come in automatically. Nodes create their ARP table without intervention from the node administrator.

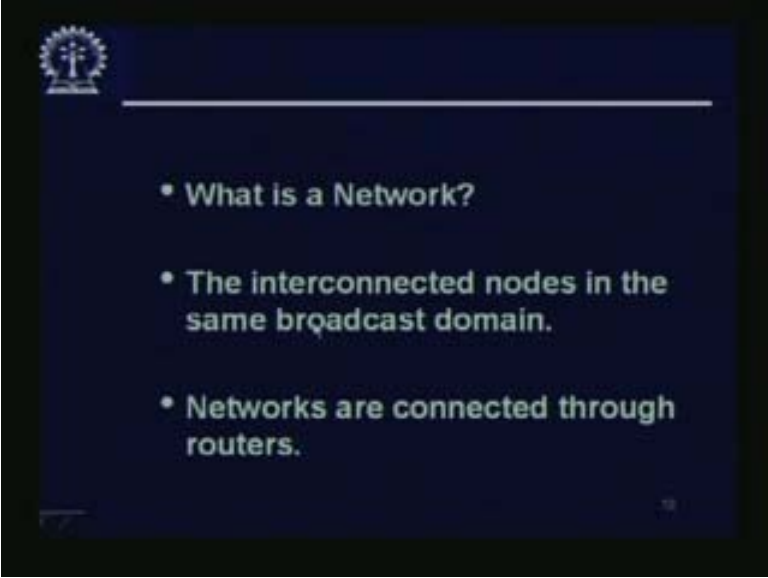
(Refer Slide Time: 56: 35)



All hosts receive the request but only B responds with the IP address; however, other nodes may update their own cache from the transaction.

The cache is required since broadcast is a very expensive proposition and more than one packet transfer is very likely between any two stations.

(Refer Slide Time: 56:54)

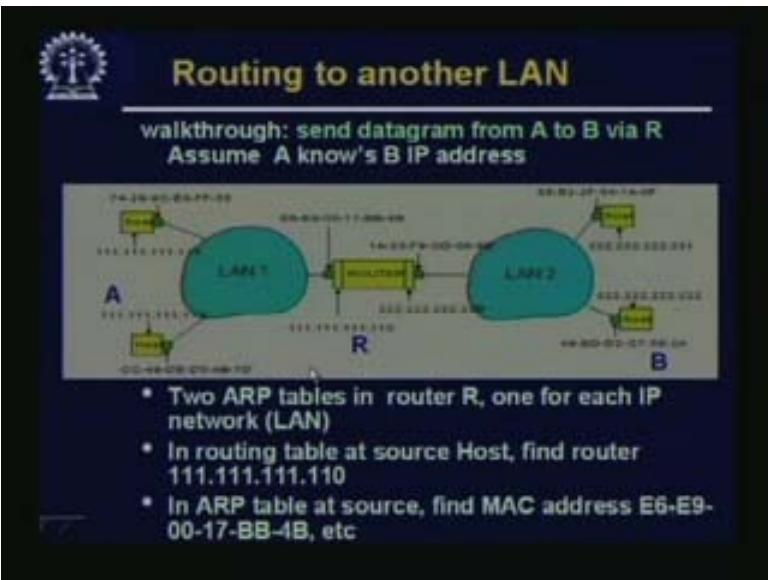


What is a Network?

- The interconnected nodes in the same broadcast domain.
- Networks are connected through routers.

This is an important question, what is a network? We always talk about networks; it is the interconnected nodes in the same broadcast domain. We are doing the broadcast – this broadcast is limited to that domain. Networks are connected through routers.

(Refer Slide Time: 57:13)



Routing to another LAN

walkthrough: send datagram from A to B via R
Assume A know's B IP address

The diagram shows two LANs, LAN 1 and LAN 2, connected by a router R. LAN 1 contains Host A and Host B. LAN 2 contains Host C and Host D. Router R is in the center, connecting the two LANs. Various IP and MAC addresses are shown for the hosts and the router interfaces.

- Two ARP tables in router R, one for each IP network (LAN)
- In routing table at source Host, find router 111.111.111.110
- In ARP table at source, find MAC address E6-E9-00-17-BB-4B, etc

When we want to route to another LAN, it is more involved. We will talk about going through a router and then we will specifically discuss bridges. Thank you.

(Refer Slide Time: 57:45)

Preview of the next lecture

(Refer Slide Time: 57:49)

Lecture Name # 21

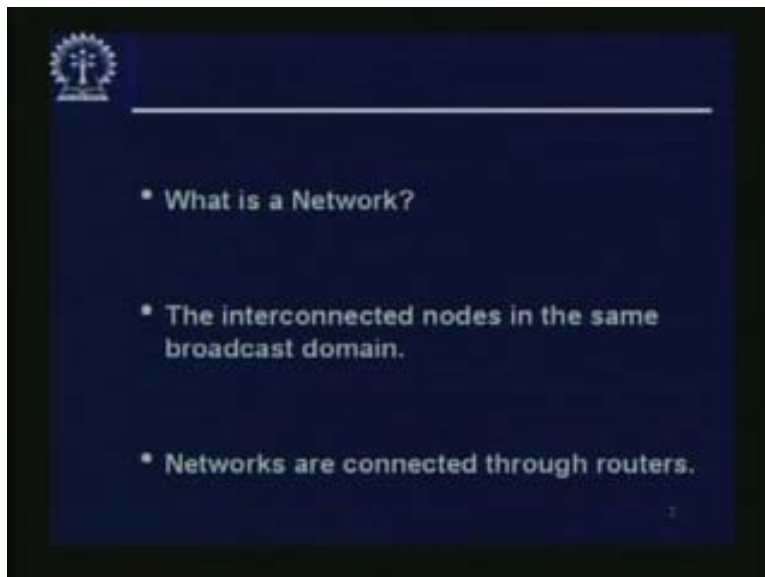
Local Internetworking

(Refer Slide Time: 57:56)



Today we will be talking about local internetworking. Our topic of today is local internetworking, what is internetworking? Internetworking is connection of different networks, everybody is aware of the term internet today why are comes to this term internetworking and by internetworking we mean connecting different networks. Just to remind you, what is a network? We discussed this in the last lecture that the interconnected nodes in the same broadcast domains.

(Refer Slide Time: 58:20)



And networks are usually connected through routers but as we will see for local internetworking we may not need a router we may need something called the bridge we will come to that. This is the set of nodes which are in the same broadcast domain; this broadcast may be a good thing to have for some applications. But for the operation of the network there is one very crucial reason, why we require the broadcast that is to discover the MAC addresses of the different computers.