ECE 333: Introduction to Communication Networks
Fall 2001

Lecture 18: Medium Access Control VI

- More on token ring networks
- LAN bridges and switches.

More on token rings

In the last lecture we began discussing token ring networks, in particular the IEEE 802.5 token ring standard. In such networks, a token is passed from node to node. Each node can only transmit when it posses the token. In the description of 802.5 last time we specified that a node will release the token only after it has finishing sending a packet **and** the first part of the packet has propagated all the way around the ring. In this case the maximum throughput of the token ring at high loads is given by:

$$max\ throughput \approx \begin{cases} 1/a & if\ a > 1 \\ 1 & if\ a \le 1 \end{cases}$$

$$\text{where } a = \frac{rotation\ time}{transmission\ time\ of\ \ packet}$$

Here the ***rotation time*** refers to the time it takes a signal to propagate around the ring, where this includes the delay incurred at each station. The above suggests that, as with CSMA/CD, if we increase the data rate, and want to still have good performance, then either the length of the ring must be decreased or the minimum packet size must increase.

For example consider a token ring where the rotation time is 1msec. If an average frame contains 2000 bytes (16000 bits), then at 4 Mbps we have

$$a = \frac{1\times10^{-3}}{(16000)/(4\times10^{6})} = .25$$

Notice in this case at any time the ring contains at most $(4\times10^{6})(1\times10^{-3}) = 4\times10^{3}$ bits or 1/4 of a frame. However at 100Mbps, we have

$$a = \frac{1\times10^{-3}}{(16000)/(100\times10^{6})} = 6.25$$

This results in a total throughput at high loads of approximately (100 Mbps)/(6.25) = 16 Mbps. In this case the ring could contain up to 6.25 frames at one time. But because of the protocol, it can only contain 1. This is because each node waits to release the token until the start of the frame has returned. A simple solution to this is to allow nodes to release the token immediately after sending a packet - this is called ***early token release***. (After a node releases the token, it must still remove the message it transmitted from the ring.)

Early Token Release

Early token release allows the next station to append its message to the tail of the previous message. In this case multiple messages can be circulating around the ring at any time. The efficiency under high loads with early token release will be approximately 1, independent of whether a >1 or not.

Early token is an option for 802.5 LANs. The advantage of **not** doing early token release is that without it the management of the ring is easier, also the priority schemes used in 802.5 work better without early token release.

FDDI - Fiber Distributed Data Interface

FDDI is an ANSI standard, purposed in 1986; it is basically a bigger and faster token ring network. It supports data rates of 100 Mbps, with a maximum distance of 200 km and up to 1000 stations. The main physical layer for FDDI is fiber optic cable, but the standard does provide for using twisted pairs over short distances. Until recently, FDDI was the preferred standard for interconnecting LANs - but Gigabit Ethernet and ATM switches are gradually displacing FDDI.

Because of the high speed and large distance, early token release is used in FDDI. An option is also provided for having multiple tokens circulating the ring at any time. The basic frame format and ring maintenance issues are similar to those in 802.5.

One different feature in FDDI is how it allocates capacity. Instead of the priority scheme used in 802.5, FDDI uses a protocol based on timers. This protocol can be used to provide guaranteed delays to synchronous (delay sensitive) traffic and also support asynchronous, bursty traffic.

FDDI - timed token rotation

In 802.5, there is a fixed token holding time for all stations on the network, this guarantees a worst-case delay for a nodes. If the token holding time is set assuming the maximum number of stations on the ring, but many stations are inactive, time is wasted forwarding the token. Ideally, the token holding time should increase as the number of active stations decreases.

In FDDI, a ***target token rotation time (TTRT)*** is established for the ring and stations compare this to the actual ***token rotation time (TRT***). The ***token holding time (THT)*** is given by THT=TTRT - TRT. If THT <0, the station can only transmit it's synchronous traffic. If THT >0, the station can transmit it synchronous traffic and any asynchronous traffic it can during the THT.

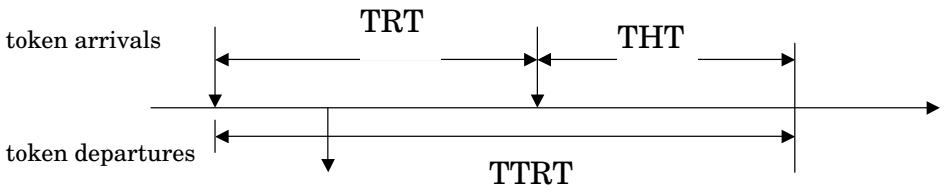

Suppose a station transmits for a long time before releasing its token by sending a lot of asynchronous traffic. The next time the token arrives, its THT will be small, and will not be able to send much asynchronous traffic.

Bigger (interconnected) LANs

In many cases organizations have multiple LANs and desire to connect them. Reasons for having multiple LANs (instead of one large LAN) include:

- Capacity requirements
- Medium length constraints
- Reliability and maintainability considerations
- Network latency concerns
- Inter-floor, inter-building, inter-campus coverage (cost)
- Security

There are several approaches to interconnecting LAN's; these approaches can be categorized by the layer at which the LANs are connected as follows:

- Hub (physical layer approach)
- Bridges (MAC layer approach)
- Routers (Network Layer approach)

(Note: this terminology varies somewhat in the literature and tends to evolve over time.)

Hub or Repeater: A hub or repeater is a node with multiple ports that simply retransmits what is received on one port on all the others. Repeaters deal with the Physical Layer only (signaling, amplification, detecting, transmitting, receiving, etc.). They do not need to have buffer space and introduce little packet delay. Attaching two LANs to a hub converts them into a single LAN. Both LANs must use same protocol. For Ethernet, when connected via a hub, two nodes are still in the same ***collison domain.*** This is not very useful approach in addressing the above issues.

Routers: A Router is a device for interconnecting two or more LANs at the network layer. Thus routers perform Physical layer thru Network layer functions. (We will talk more about these next week).

Bridges or LAN switches: A bridge or LAN switch refers to a device that interconnects two or more LANs at the MAC layer. The resulting network is often called ***extended LAN***. A bridge is an older term than a LAN switch; some people make various distinctions between Bridges and LAN switches, while others uses these as synonymous terms. Will not distinguish between these here.

Bridge functionality

Basically, a bridge performs a MAC layer relay, that is:

- It reads all the packets transmitted on every LAN to which it is connected.
- It retransmits the packet on the LAN that owns the respective destination nodes (Routing.)
- It is responsible for all MAC layer functions (e.g., collision detection, carrier sensing, exponential back-off in CSMA/CD, issuing free tokens, removing data, etc.)

Bridges need some buffer capacity and introduce some processing delay but operate at very high speed (specified by manufacturer).

As we noted in lecture 16, for Ethernet, two LANs connected by a bridge or switch are divided into separate collision domains.

Bridge Examples: Token ring bridge

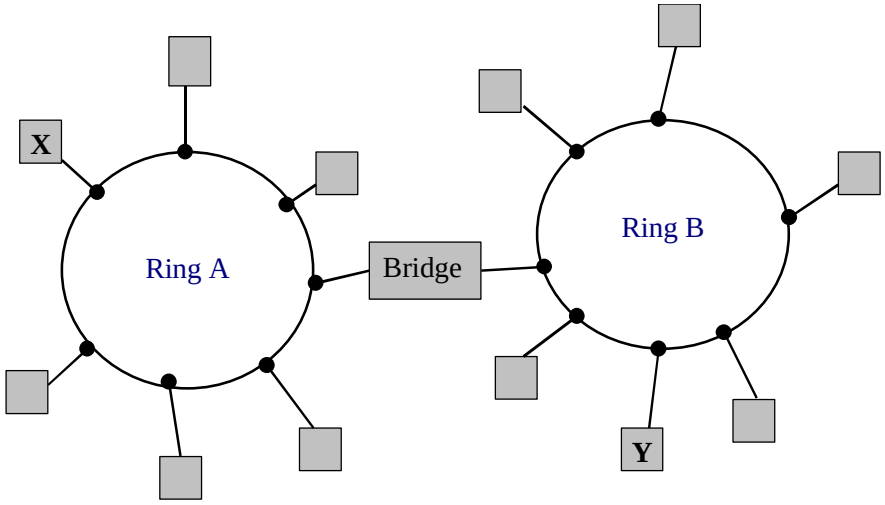

For X -> Y packet, bridge must receive packet from ring A and acknowledge packet, acquire token on ring B, and forward packet to Y.

Ethernet Switch (bridge)

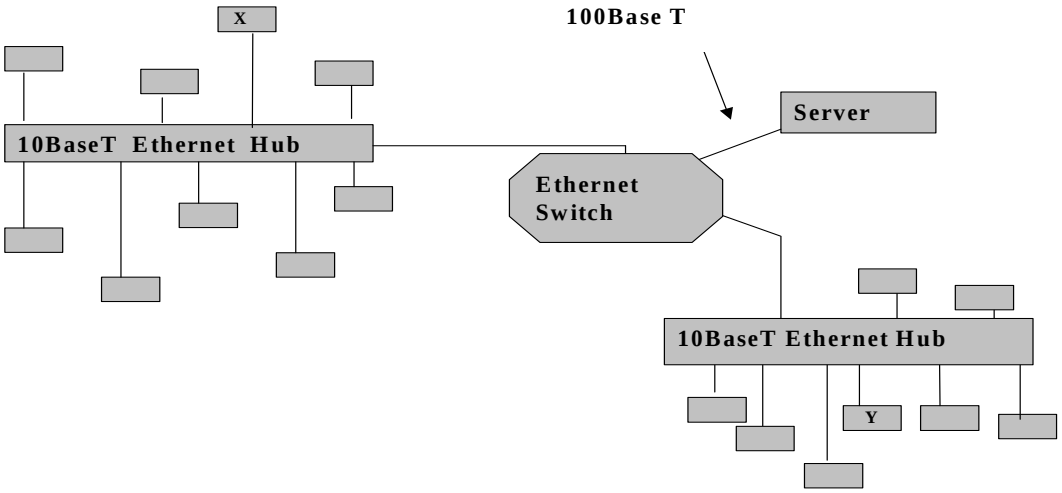

Packet from X->Y must contend with other nodes on left Ethernet LAN, to be received by the switch. The switch then contends to transmit packet on the right LAN.

The protocol stack that is executed in a bridge is shown below:

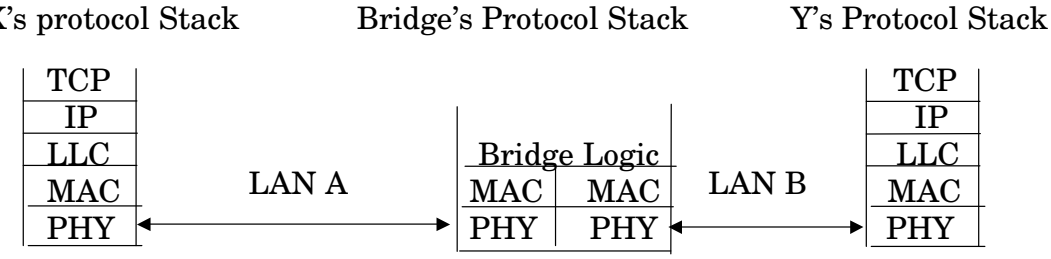

For a router, the network layer in the protocol stack is also executed; an example with two LANs connected across a WAN is shown below.

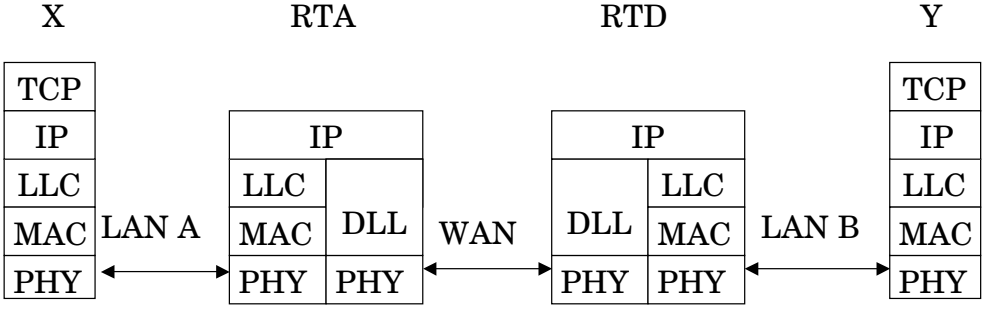

When a bridge receives a packet, it must know whether or not to forward the packet. This is a basic example of a **routing problem**. Consider two LANs connected by a bridge below. If the bridge receives a transmission on LAN 1 for node D, it should retransmit that packet on LAN 2, but it should not retransmit a packet for node A received from LAN 1. There are two basic approaches to this problem. One approach is to simply have a look-up table, which specifies which nodes are on which LANs, such a table is called a ***forwarding table*** or ***forwarding database***. When a packet is received the bridge looks for the destination in the table and decides to forward the packet or not. Alternatively, each packet could have additional information in the header that identifies the destination LAN. With either of these approaches, a key issue is how the information (either in the forwarding table or in the packet header) is known. One approach would be to have a system administrator keep all of this information up-to-date, but for large LANs this is undesirable approach. Instead various algorithms are used to acquire this information.

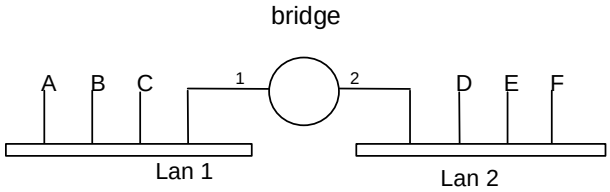

There are two basic types of bridges (both standardized by the (IEEE 802 working group). The main conceptual difference between these approaches is in how the routing problem is handled.

Transparent Bridge (specified in IEEE 802.1)

Accepted as standard for interconnection of 802.x LANs.
Use forwarding table approach.

Source Routing Bridge:

Developed separately by 802.5 committee (Token Ring)
Modified for connecting 802.3 & 802.4
Put forwarding information in packet header.

In the following we will look closer at the transparent bridge approach; for a discussion of source routing bridges see Tanenbaum, Section 4.4.2.

Transparent Bridges

Transparent bridges were designed to be a *plug and play* system - i.e., a transparent bridge can be simply be connected to two or more LANs without making any changes to the stations attached to the LANs or requiring any configuration on the part of the users. Everything should work automatically. Thus the bridge must learn the location of stations in the LANs and build up its forwarding database.

A transparent bridge operates in ***promiscuous mode*** and receives all packets transmitted on the LANs to which it is attached. Each packet contains both source and destination addresses. When a packet is received, the bridge associates the port on which the packet was received with the source addresses to create the forwarding table. (This is called ***backward learning***.)

When the bridge receives a packet, it forwards it if the destination is not reachable on the incoming LAN. If the destination is not in the forwarding table, then the bridge forwards the packet on every LAN, except the one it received the packet on.

Difficulties with Learning

Two things make the above problem somewhat more difficult:

1. Stations can be moved from one LAN to another.

2. There may be loops in the topology.

The first problem is addressed by including an age field in the forwarding table for each entry. This age field is updated whenever a new frame from that station arrives. Information is discarded if it is too old.

Loops

Often extra bridges are used to increase reliability. However, this can create loops in the topology:

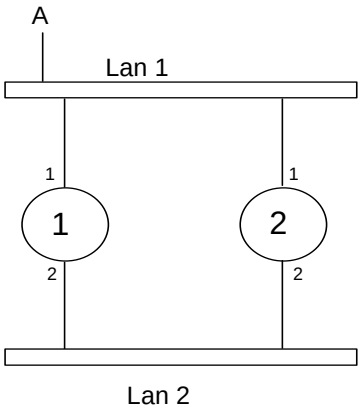

Suppose **node A** sends a packet with a destination address that is not in either bridges forwarding table. In this case, both bridges would forward to LAN 2. Each packet arrives at other bridge on LAN 2 and is forwarded to LAN 1, etc.

Spanning Trees

To avoid packets being broadcast indefinitely as above, transparent bridges use an algorithm to construct a ***spanning tree***. A spanning tree is a loop-free connection of all LANs in the network. We can represent the network as a graph[1], where each node in the graph corresponds to a LAN and arcs connect any two LANs that are connected by a bridge. A spanning tree is a loop-free sub-graph that connects all the nodes.

Example: A network of bridges and LANs is shown on the left below, the corresponding graph is shown on the right. One spanning tree is to include only the arcs that are highlighted. Notice there may be several possible spanning trees for a graph.

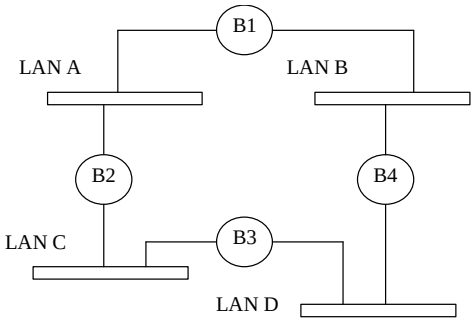

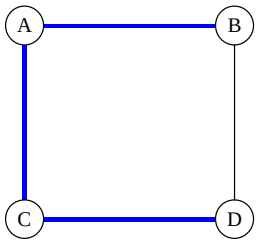

[1] Actually this representation may be a multi-graph, because there may be multiple edges between a pair of nodes.