ECSE-4730: Computer Communication Networks (CCN)

Chapter 5:
Data Link Layer: Part II

Shivkumar Kalyanaraman: shivkuma@ecse.rpi.edu
Biplab Sikdar: sikdab@rpi.edu
http://www.ecse.rpi.edu/Homepages/shivkuma

Rensselaer

RSVP

Summary of MAC protocols

- What do you do with a shared media?
  - Channel Partitioning: time, frequency or code
    - Time Division,Code Division, Frequency Division
  - Random partitioning (dynamic),
    - ALOHA, S-ALOHA, CSMA, CSMA/CD
    - carrier sensing: easy in some technoligies (wire), hard in others (wireless)
    - CSMA/CD used in Ethernet
  - Taking Turns
    - polling from a central cite, token passing

LAN technologies

Data link layer so far:

- services, error detection/correction, multiple access

Next: LAN technologies

- addressing
- Ethernet
- hubs, bridges, switches
- 802.11
- PPP
- ATM

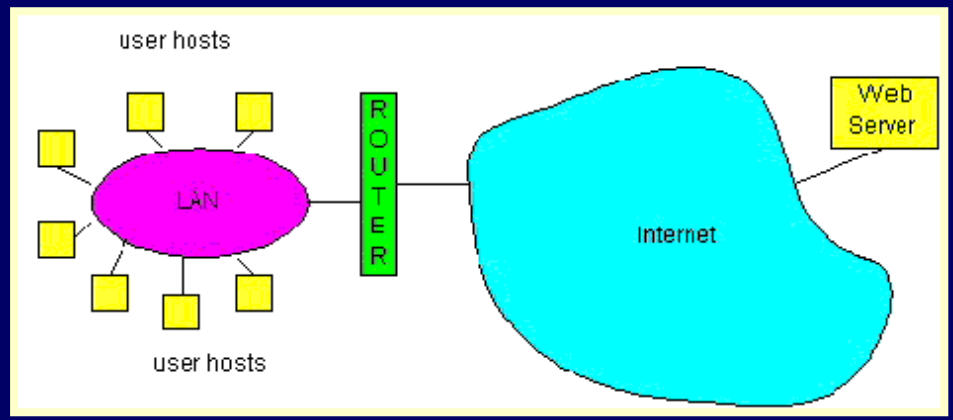

LAN Addresses and ARP - 1

32-bit IP address:

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

LAN (or MAC or physical) 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**

LAN Addresses and ARP - 2

Each adapter on LAN has unique LAN address

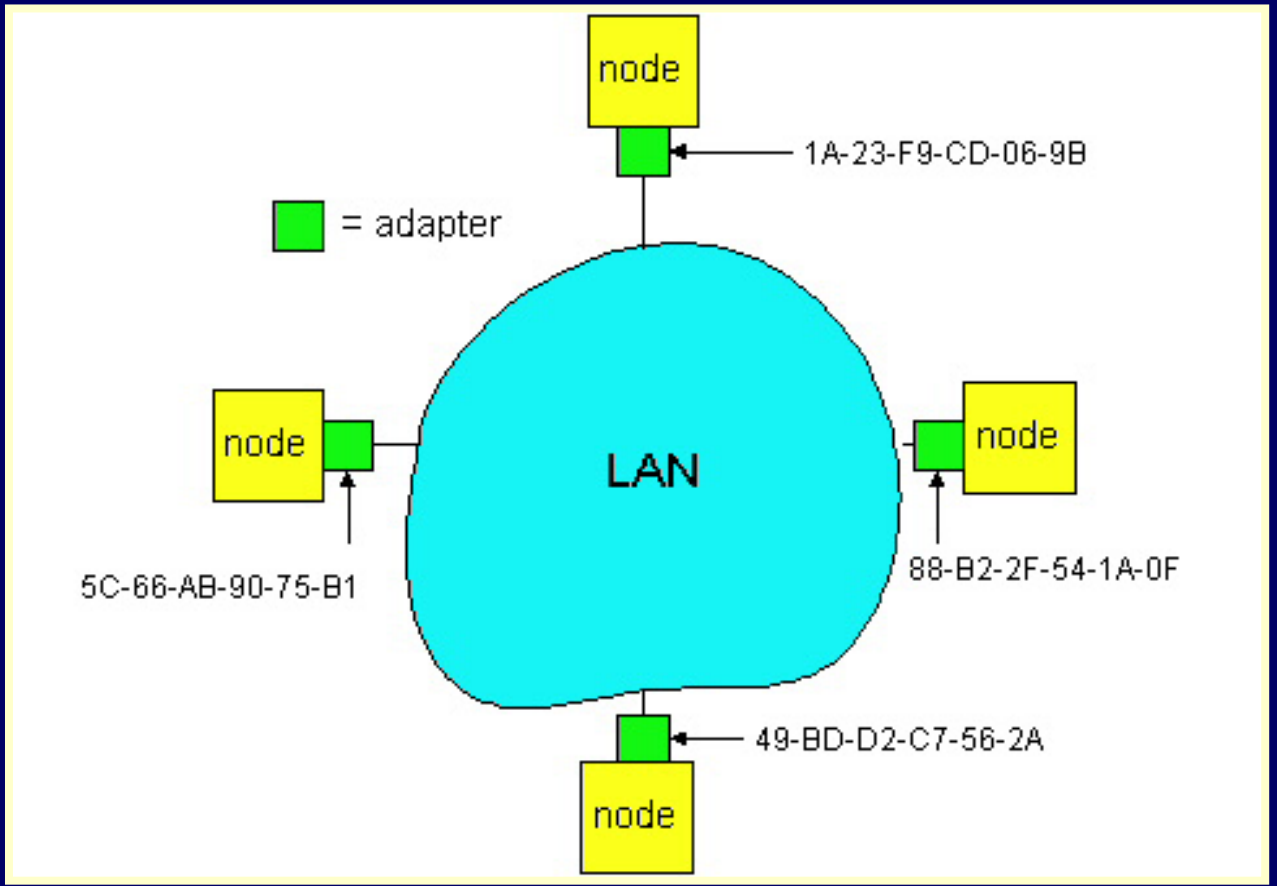

LAN Address (more)

- MAC address allocation administered by IEEE
- manufacturer buys portion of MAC address space (to assure uniqueness)
- 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 network to which one attaches

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**

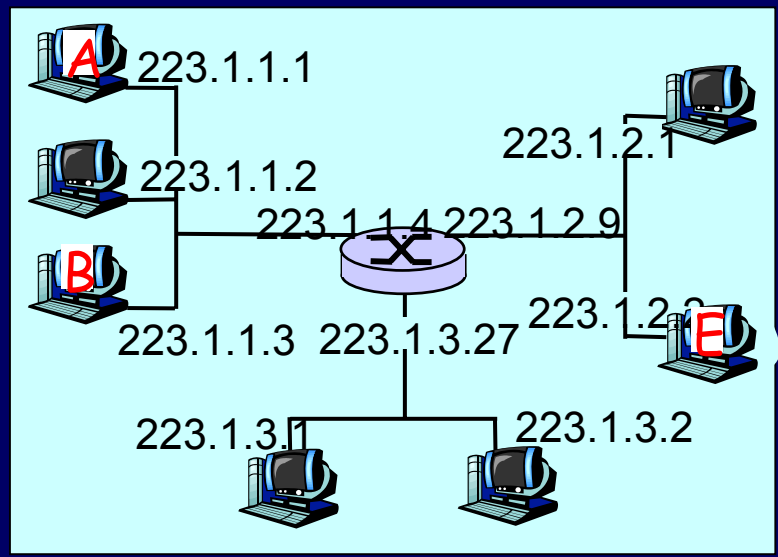

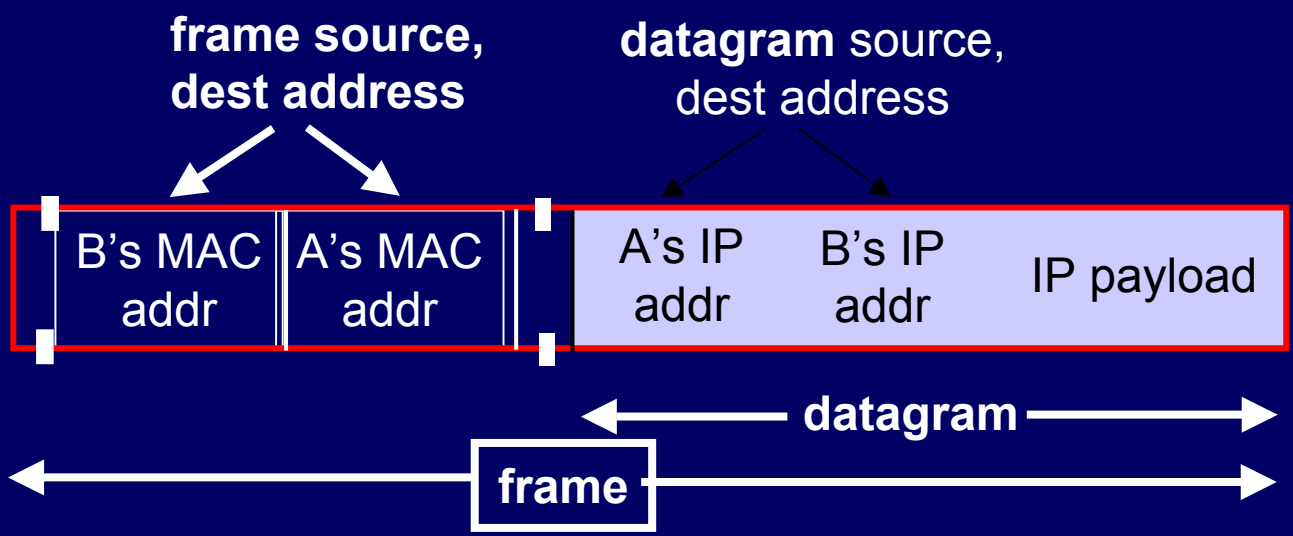

ARP: Address Resolution Protocol

Question: how to determine MAC address of B given B's IP address?

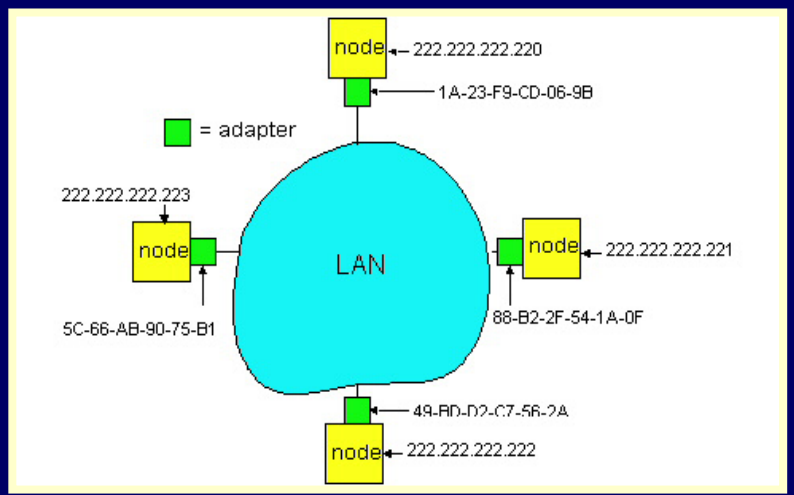

- Each IP node (Host, Router) on LAN has ARP module, table
- ARP Table: IP/MAC address mappings for some LAN nodes

< IP address; MAC address; TTL>

< ..... >

- TTL (Time To Live): time after which address mapping will be forgotten (typically 20 min)

ARP protocol

- A knows B's IP address, wants to learn physical address of B
- A broadcasts ARP query pkt, containing B's IP address
  - all machines on LAN receive ARP query
- B receives ARP packet, replies to A with its (B's) physical layer address
- A caches (saves) IP-to-physical address pairs until information becomes old (times out)
  - soft state: information that times out (goes away) unless refreshed

Routing to another LAN - 1

Walkthrough: routing from A to B via R

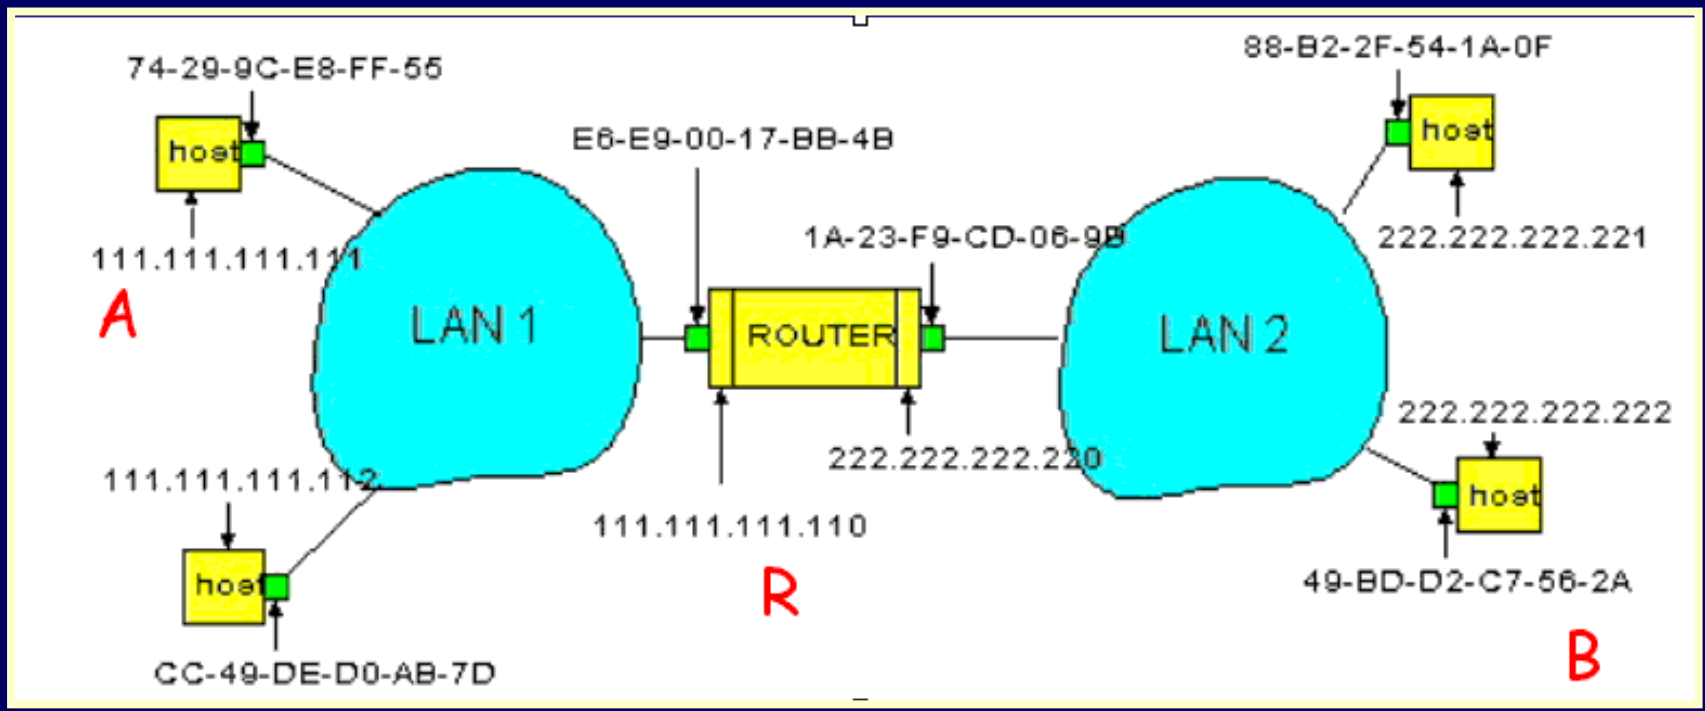

Routing to another LAN - 2

- A creates IP packet with source A, destination B
- A uses ARP to get R's physical layer address for 111.111.111.110
- A creates Ethernet frame with R's physical address as dest, Ethernet frame contains A-to-B IP datagram
- A's data link layer sends Ethernet frame

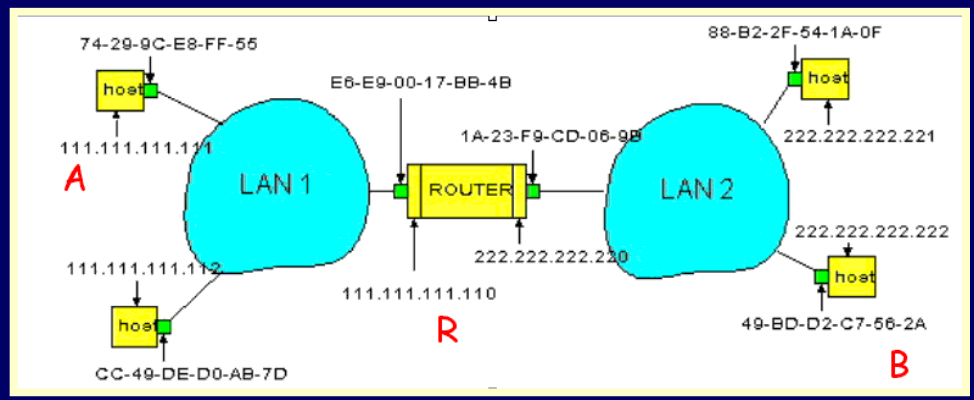

Routing to another LAN - 3

- R's data link layer receives Ethernet frame
- R removes IP datagram from Ethernet frame, sees its destined to B
- R uses ARP to get B's physical layer address
- R creates frame containing A-to-B IP datagram sends to B

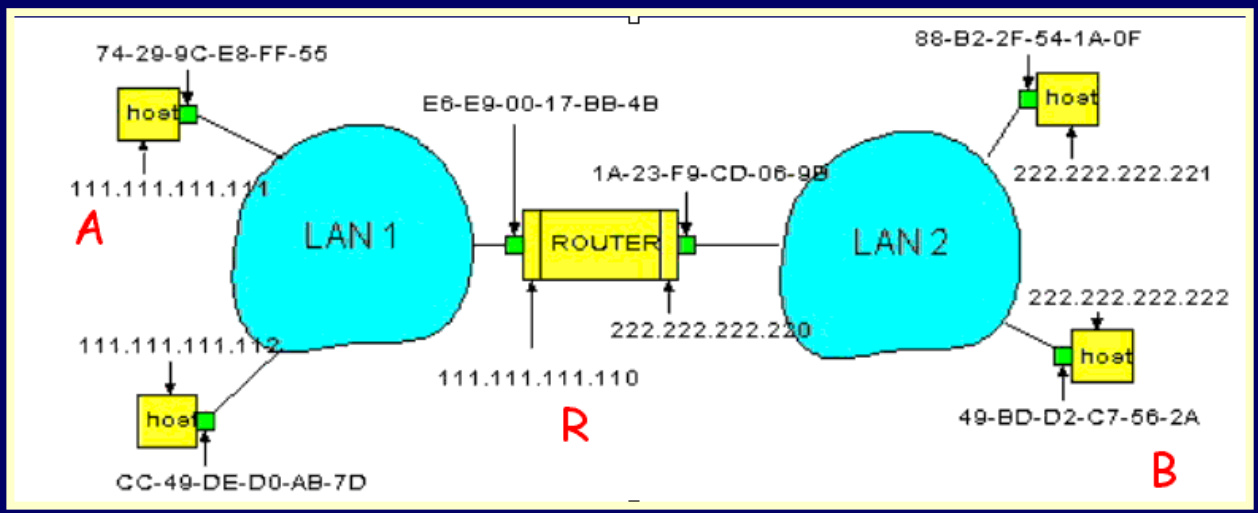

Ethernet

"Dominant" LAN technology:

- Cheap $20 for 100Mbs!
- First widely used LAN technology
- Simpler, cheaper than token LANs and ATM
- Kept up with speed race: 10, 100, 1000 Mbps

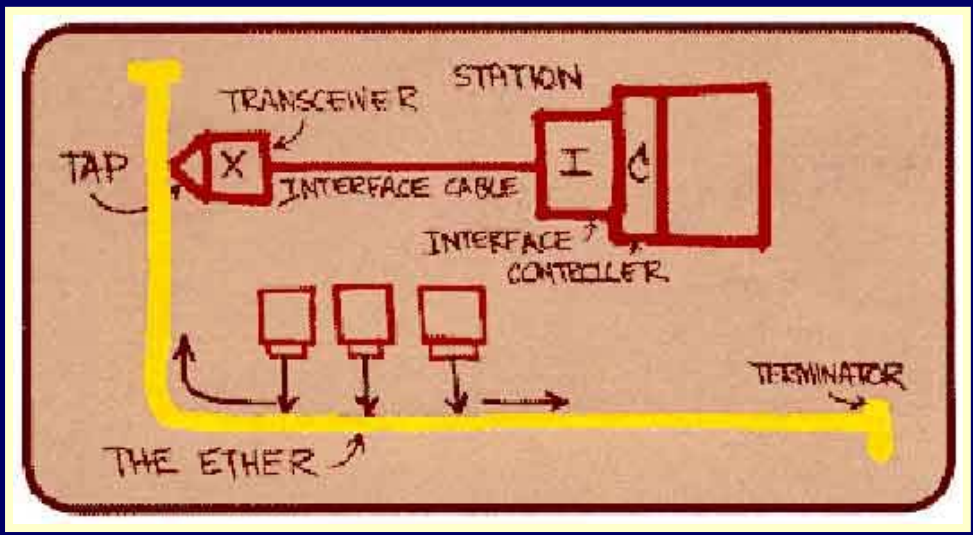

Metcalfe's Ethernet sketch

Ethernet Frame Structure - 1

Sending adapter encapsulates IP datagram (or other network layer protocol packet) in Ethernet frame

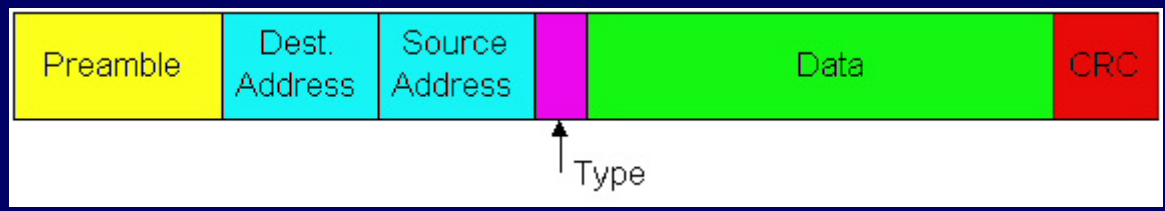

Preamble:

- 7 bytes with pattern 10101010 followed by one byte with pattern 10101011
- Used to synchronize receiver, sender clock rates

Ethernet Frame Structure - 2

- **Addresses:** 6 bytes, frame is received by all adapters on a LAN and dropped if address does not match
- **Type:** indicates the higher layer protocol, mostly IP but others may be supported such as Novell IPX and AppleTalk)
- **CRC:** checked at receiver, if error is detected, the frame is simply dropped

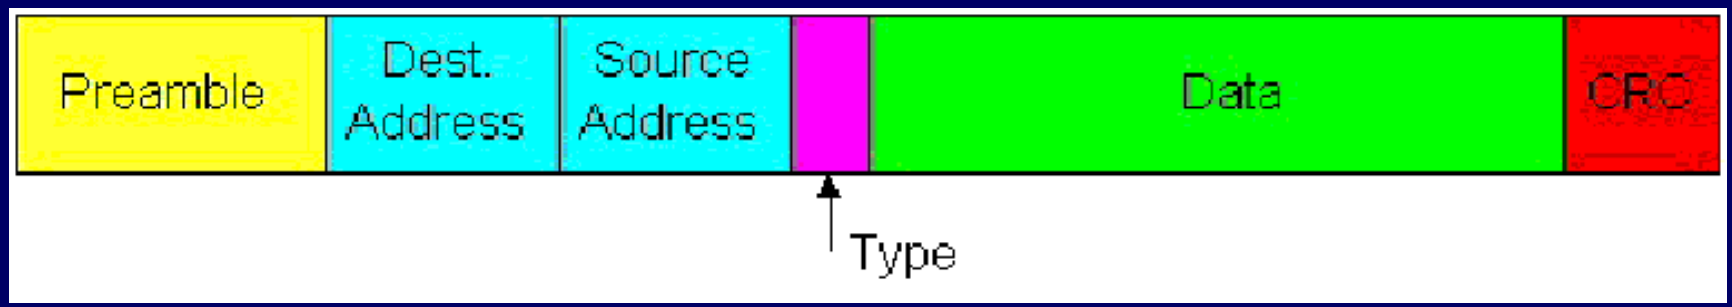

Ethernet

- Ethernet uses 1-persistent CSMA/CD on coaxial cable at 10 Mbps (802.3 allows other speeds & media)
- The maximum cable length allowed: 500m
- Longer distances covered using repeaters to connect multiple "segments" of cable
- No two stations can be separated by more than 2500 meters and 4 repeaters
- Including the propagation delay for 2500m and the store and forward delay in 4 repeaters, the maximum time for a bit to travel between any two stations is τ_{max}=25.6μse (one way)

Ethernet: uses CSMA/CD

```
A: sense channel, if idle
   then {
          transmit and monitor the channel;
          If detect another transmission
           then {
                abort and send jam signal;
               update # collisions;
               delay as required by exponential backoff
                 algorithm;
               goto A
               }
            else {done with the frame; set collisions to
              zero}
      }
   else {wait until ongoing transmission is over and
      goto A}
```

Ethernet's CSMA/CD

- In order to ensure that every collision i s "heard" by all stations, when a station detects a collision , it jams the channel for
- Example
  - Two stations, A and B, are close together
  - A third station, C, is far away
  - A and B will detect each other's transmission very quickly and shut off
  - This will only cause a short blip which may not be detected by C but will still cause enough errors to destroy C's packet

Ethernet's CSMA/CD

- When collisions occur, Ethernet uses a random retransmission scheme called exponential backoff:
  1. If your packet is in a collision, set K=2
  2. Pick a number k at random from {0, 1,..,K-1}
  3. After τ_{max} seconds, sense channel, transmit if idle
  4. If collision occurs, let K=2 x K, go to step 2
- After 10 repeats, stop doubling K
- After 16, give up and tell layer above "I give up"
- "Fixes" random access stability problem by passing it to the layer above!

Ethernet Technologies: 10Base2

- 10: 10Mbps; 2: under 200 meters max cable length
- Thin coaxial cable in a bus topology

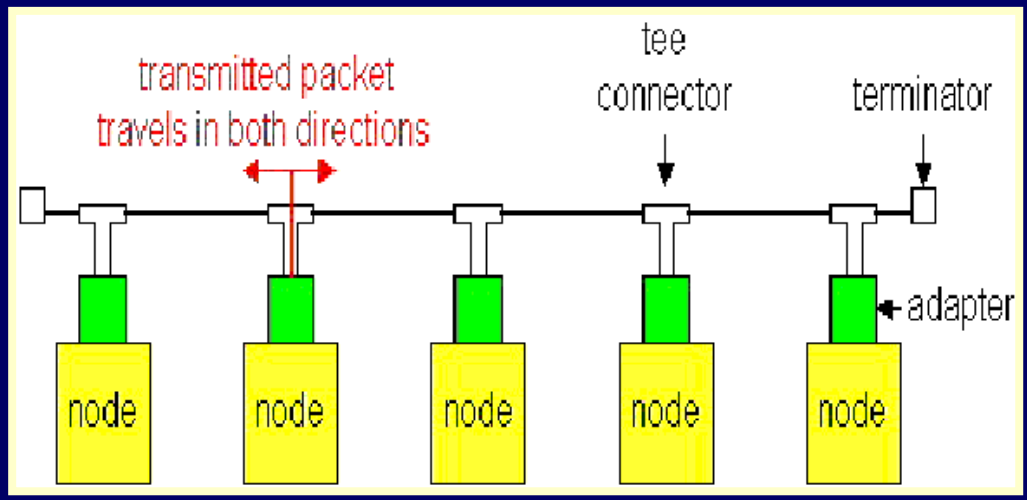

- Repeaters used to connect up to multiple segments
- Repeater repeats bits it hears on one interface to its other interfaces: physical layer device!

10BaseT and 100BaseT - 1

- 10/100 Mbps rate; latter called “fast ethernet”
- T stands for Twisted Pair
- Hub to which nodes are connected by twisted pair, thus “star topology”
- CSMA/CD implemented at hub

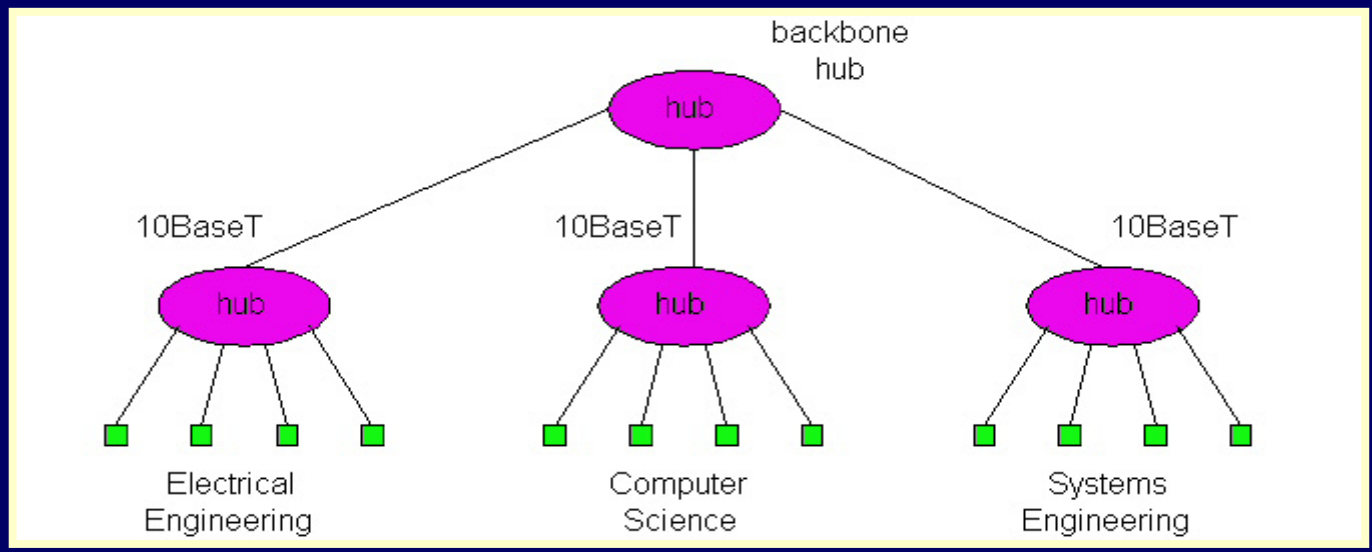

10BaseT and 100BaseT - 1

- Max distance from node to Hub is 100 meters
- Hub can disconnect “jabbering adapter
- Hub can gather monitoring information, statistics for display to LAN administrators

Gbit Ethernet

- Use standard Ethernet frame format
- Allows for point-to-point links and shared broadcast channels
- In shared mode, CSMA/CD is used; short distances between nodes to be efficient
- Uses hubs, called here “Buffered Distributors”
- Full-Duplex at 1 Gbps for point-to-point links

Token Rings (IEEE 802.5)

- A ring topology is a single unidirectional loop connecting a series of stations in sequence
- Each bit is stored and forwarded by each station's network interface

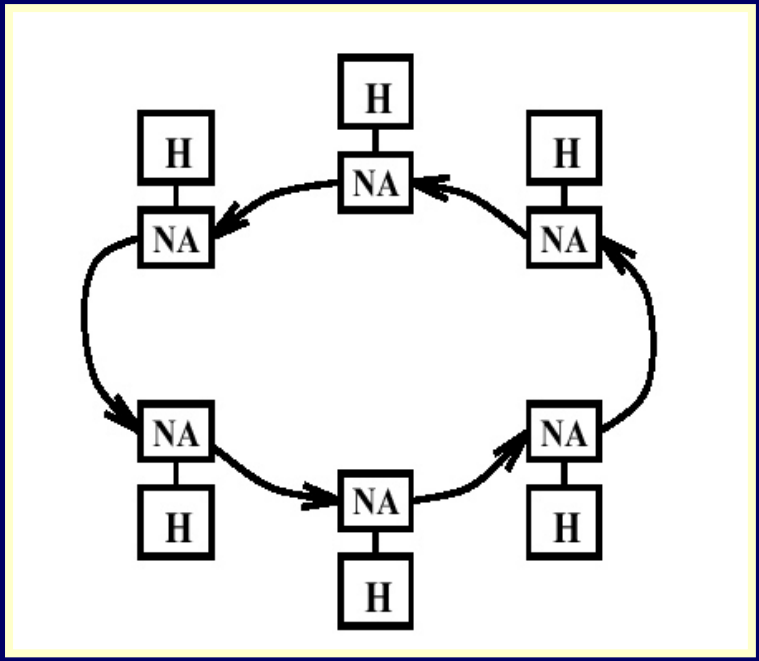

Token Rings: IEEE 802.5 -1

- Versions that operate at 1, 4, and 16 Mbps over shielded twisted pair copper wire
- Max token holding time: 10 ms, limiting frame length

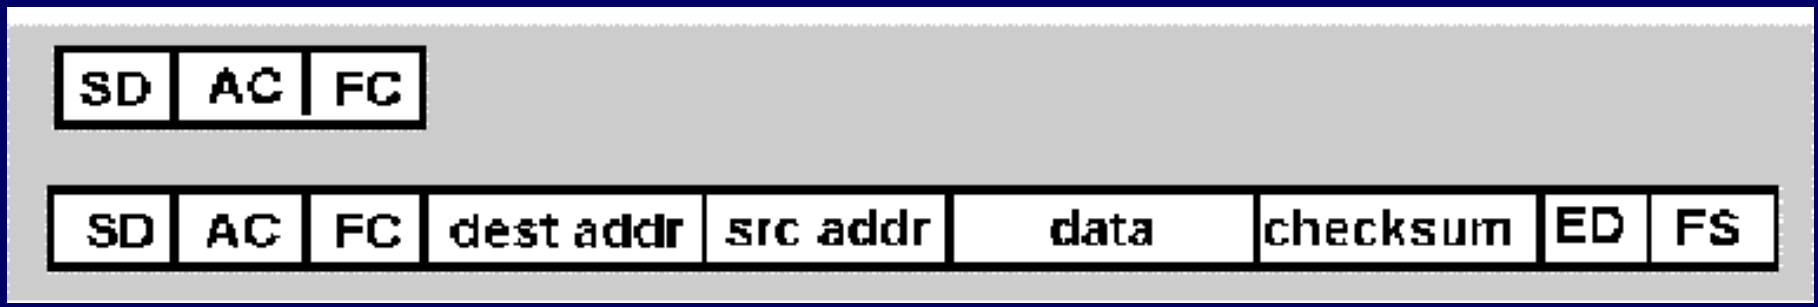

- SD, ED mark start, end of packet

Token Ring: IEEE 802.5 - 2

AC: access control byte:

- **Token bit: value 0 means token can be seized, value 1 means data follows FC**
- **Priority bits: priority of packet**
- **Reservation bits: station can write these bits to prevent stations with lower priority packet from seizing token after token becomes free**

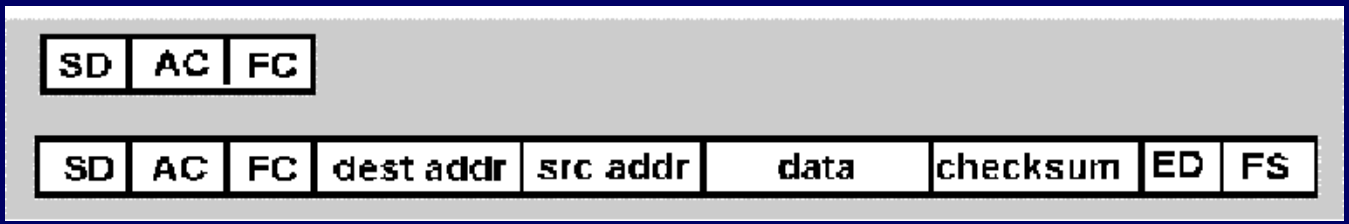

Token Ring: IEEE 802.5 - 3

- **FC: frame control used for monitoring and maintenance**
- **Source, destination address: 48 bit physical address, as in Ethernet**
- **Data: packet from network layer**
- **Checksum: CRC**
- **FS: frame status: set by dest., read by sender**
  - **set to indicate destination up, frame copied OK from ring**
  - **DLC-level ACKing**

Token Ring: IEEE 802.5 - 4

- After transmitting one or more packets (depending on the rules of the protocol), the node transmits a new token to the next node in one of 3 ways:
  1. Single Packet Mode: Token is transmitted after receiving the last bit of transmitted packet(s)
  2. Multiple Token Mode: Token is transmitted immediately after the last bit of the packet(s) is transmitted
- In small rings, the last two are the same

Interconnecting LANs

Q: Why not just one big LAN?

- **Limited amount of supportable traffic: on single LAN, all stations must share bandwidth**
- **Limited length: 802.3 specifies maximum cable length**
- **Large "collision domain" (can collide with many stations)**
- **Limited number of stations: 802.5 have token passing delays at each station**

Hubs - 1

- Physical Layer devices: essentially repeaters operating at bit levels: repeat received bits on one interface to all other interfaces
- Hubs can be arranged in a hierarchy (or multi-tier design), with backbone hub at its top

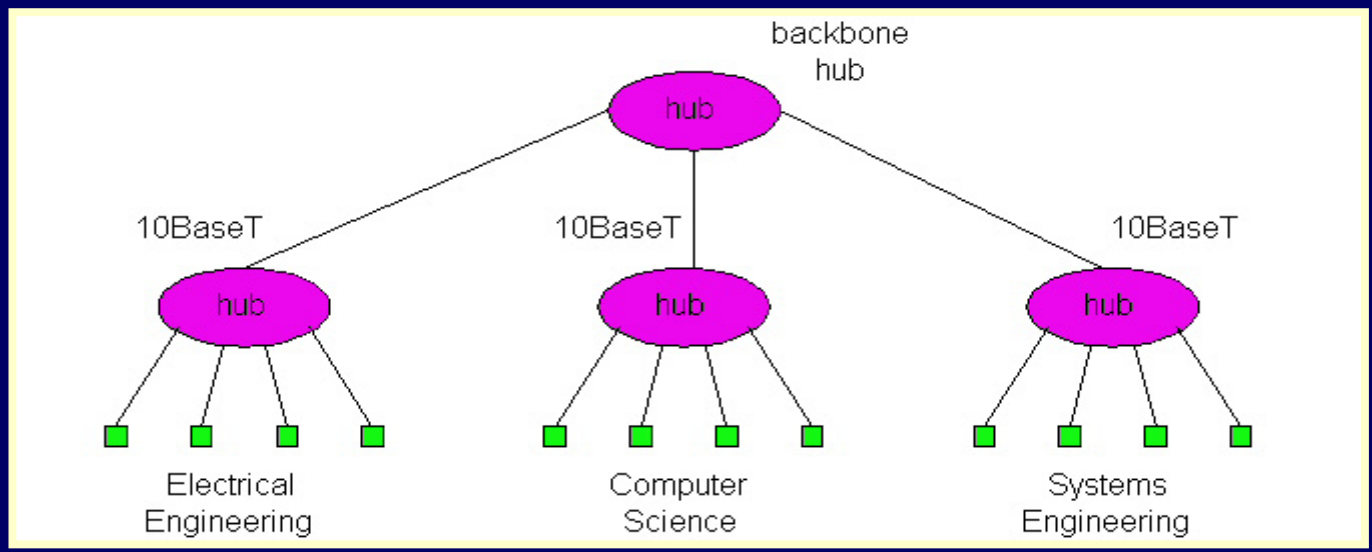

Hubs - 2

- Each connected LAN referred to as LAN segment
- Hubs do not isolate collision domains: node may collide with any node residing at any segment in LAN
- Hub Advantages:
  - simple, inexpensive device
  - Multi-tier provides graceful degradation: portions of the LAN continue to operate if one hub malfunctions
  - extends maximum distance between node pairs (100m per Hub)

Hub limitations

- Single collision domain results in no increase in max throughput
  - multi-tier throughput same as single segment throughput
- Individual LAN restrictions pose limits on number of nodes in same collision domain and on total allowed geographical coverage
- Cannot connect different Ethernet types (e.g., 10BaseT and 100baseT)

Bridges - 1

- **Link Layer devices: operate on Ethernet frames, examining frame header and selectively forwarding frame based on its destination**
- **Bridge isolates collision domains since it buffers frames**
- **When frame is to be forwarded on segment, bridge uses CSMA/CD to access segment and transmit**

Bridges - 2

- **Bridge advantages:**
  - **Isolates collision domains resulting in higher total max throughput, and does not limit the number of nodes nor geographical coverage**
  - **Can connect different type Ethernet since it is a store and forward device**
  - **Transparent: no need for any change to hosts LAN adapters**

Bridges: frame filtering, forwarding

- Bridges filter packets
  - Same-LAN -segment frames not forwarded onto other LAN segments

- Forwarding:
  - How to know which LAN segment on which to forward frame?
  - Looks like a routing problem (more shortly!)

Backbone Bridge

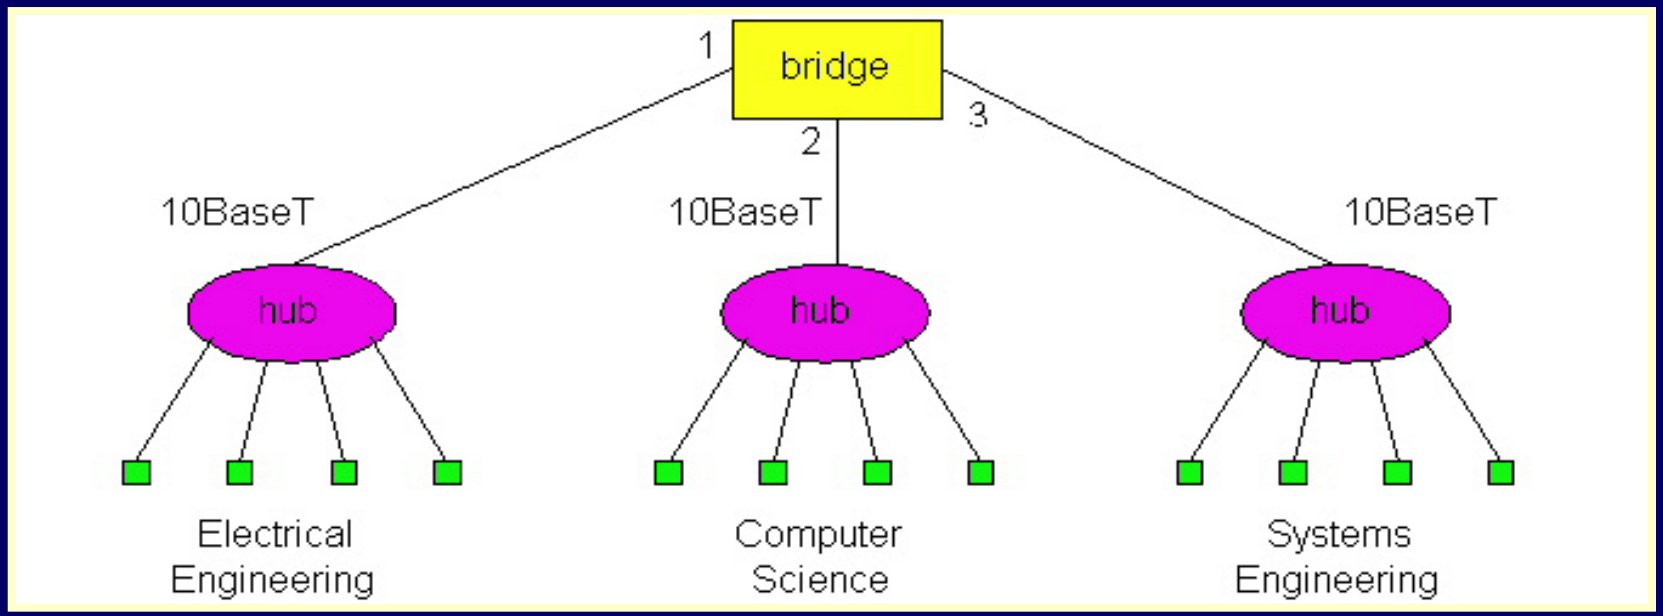

Interconnection Without Backbone

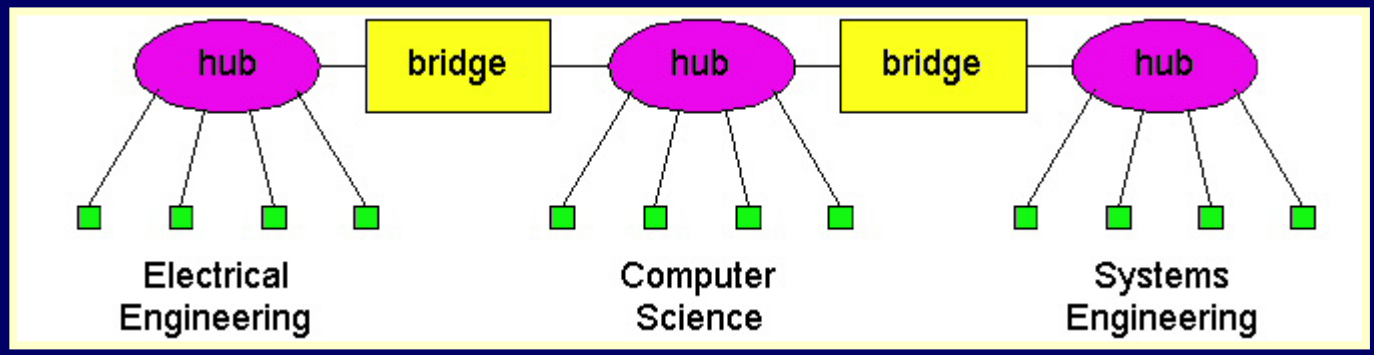

- Not recommended for two reasons:
  - Single point of failure at Computer Science hub
  - All traffic between EE and SE must path over CS segment

Bridge Filtering - 1

- Bridges learn which hosts can be reached through which interfaces: maintain filtering tables
  - when frame received, bridge "learns" location of sender: incoming LAN segment
  - records sender location in filtering table
- Filtering table entry:
  - (Node LAN Address, Bridge Interface, Time Stamp)
  - stale entries in Filtering Table dropped (TTL can be 60 minutes)

Bridge Filtering - 2

- **Filtering procedure:**

```
if destination is on LAN on which frame
   was received
     then drop the frame
     else { lookup filtering table
              if entry found for destination
             then forward the frame on
               interface indicated;
             else flood;    /* forward on all but
               the interface on
               which the frame arrived*/
        }
```

Bridge Learning: example - 1

Suppose C sends frame to D and D replies back with frame to C

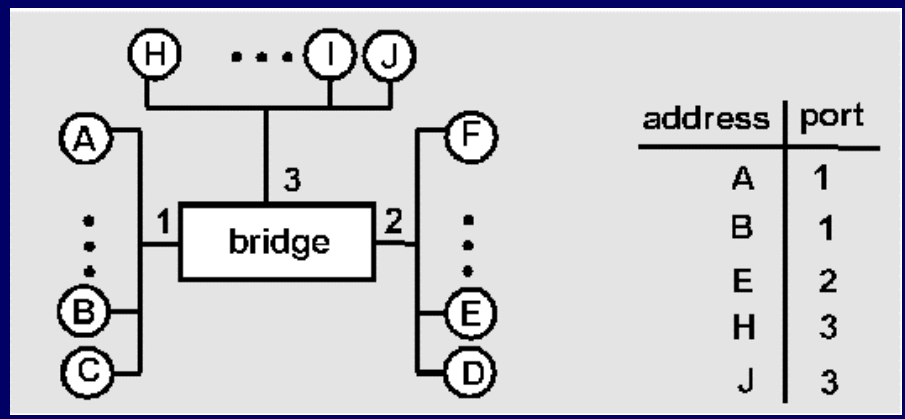

- **C sends frame, bridge has no info about D, so floods to both LANs**
  - **bridge notes that C is on port 1**
  - **frame ignored on upper LAN**
  - **frame received by D**

Bridge Learning: example - 2

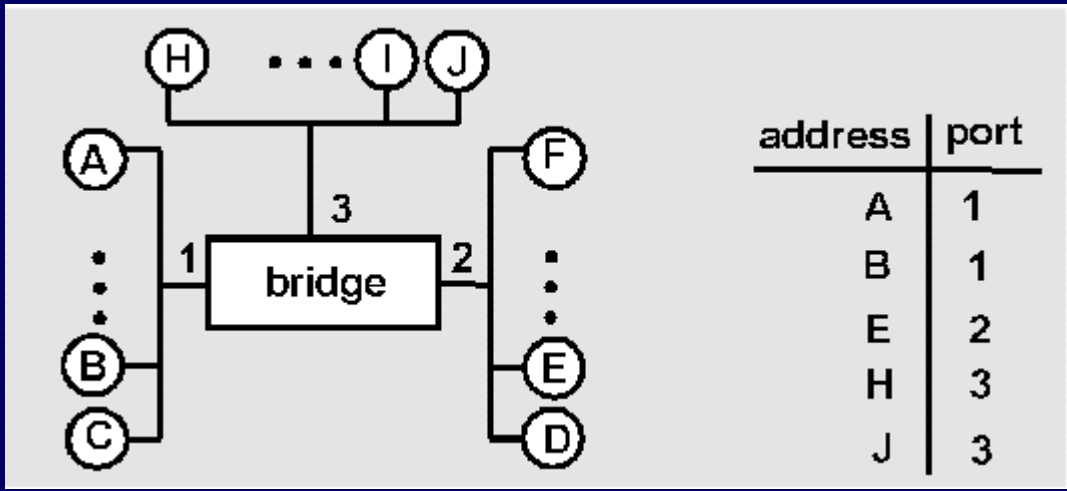

- D generates reply to C, sends
  - bridge sees frame from D
  - bridge notes that D is on interface 2
  - bridge knows C on interface 1, so *selectively* forwards frame out via interface 1

Bridges Spanning Tree

- For increased reliability, desirable to have redundant, alternate paths from source to dest
- With multiple simultaneous paths, cycles result - bridges may multiply and forward frame forever
- Solution: organize bridges in a spanning tree by disabling subset of interfaces

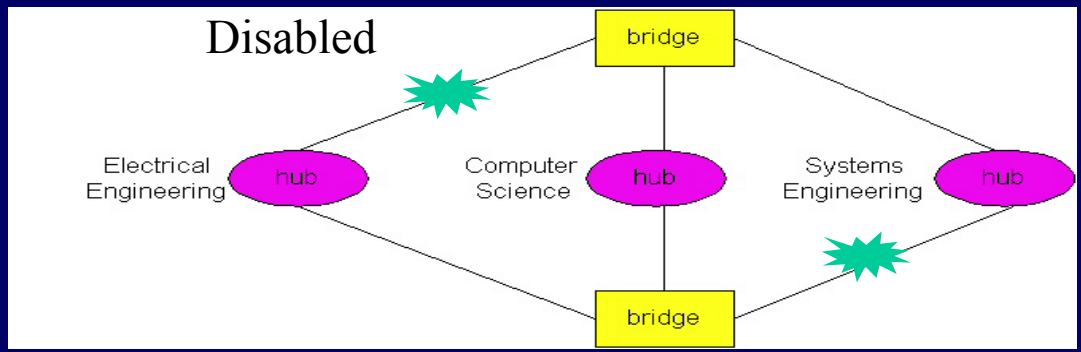

WWF Bridges vs. Routers

- Both store-and-forward devices
  - routers: network layer devices (examine network layer headers)
  - bridges are Link Layer devices
- Routers maintain routing tables, implement routing algorithms
- Bridges maintain filtering tables, implement filtering, learning and spanning tree algorithms

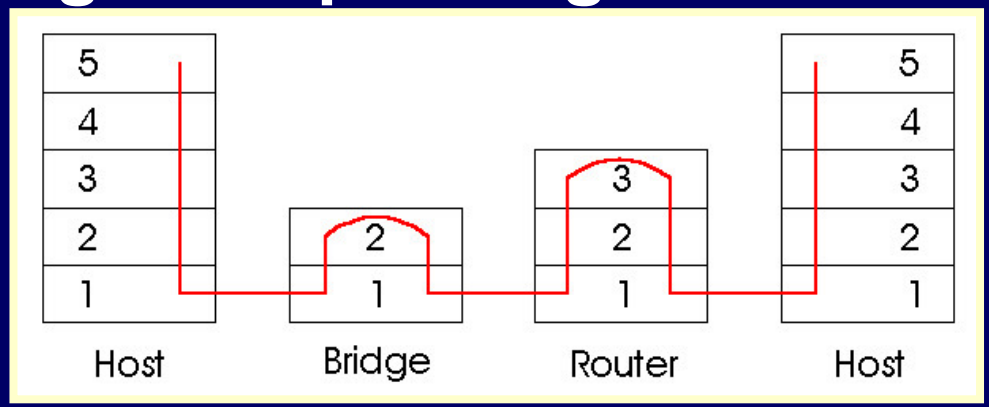

Routers vs. Bridges - 1

Bridges + and -

+ Bridge operation is simpler requiring less processing bandwidth
- Topologies are restricted with bridges: a spanning tree must be built to avoid cycles
- Bridges do not offer protection from broadcast storms (endless broadcasting by a host will be forwarded by a bridge)

Routers vs. Bridges - 2

Routers + and -

+ arbitrary topologies can be supported, cycling is limited by TTL counters (and good routing protocols)
+ provide firewall protection against broadcast storms
- require IP address configuration (not plug and play)
- require higher processing bandwidth
- Bridges do well in small (few hundred hosts) while routers used in large networks (thousands of hosts)