Last time

- Multiple access protocols
  - Channel partitioning MAC protocols
    - TDMA, FDMA
  - Random access MAC protocols
    - Slotted Aloha, Pure Aloha, CSMA, CSMA/CD
  - "Taking turns" MAC protocols
    - Polling, token passing
- Link-layer addressing

This time

- Ethernet
- Hubs and switches

Link Layer

Ethernet

Ethernet is the “dominant” wired LAN technology:

- Cheap: $20 for 100Mbps!
- The first widely used LAN technology
- Simpler, cheaper than token LANs
- Kept up with speed race: 10 Mbps – 10 Gbps

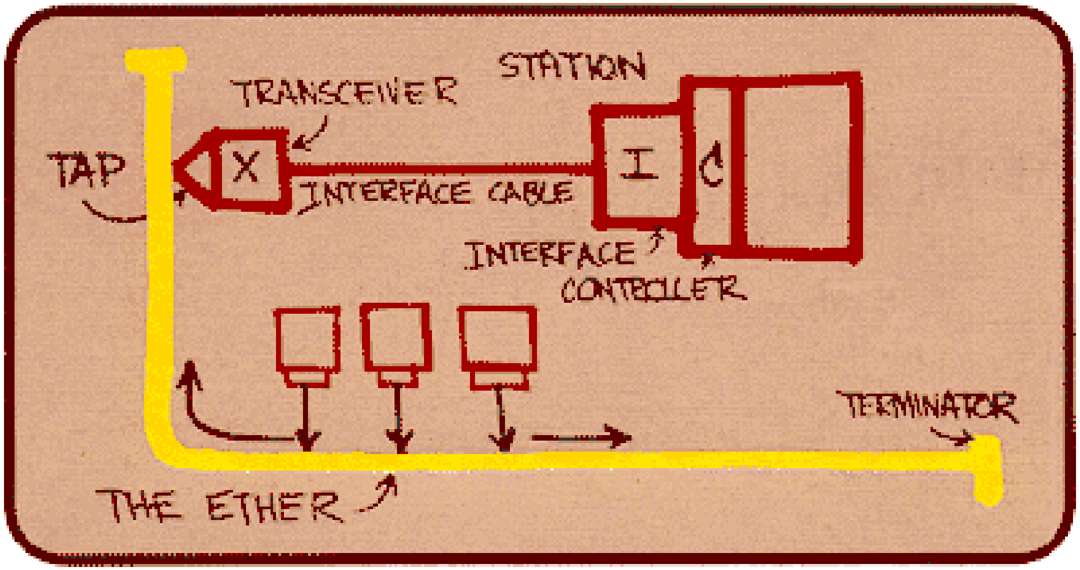

Metcalfe’s original Ethernet sketch

Star topology

- Bus topology popular through mid 90s
- Now star topology prevails
- Connection choices: hub or switch (more later)

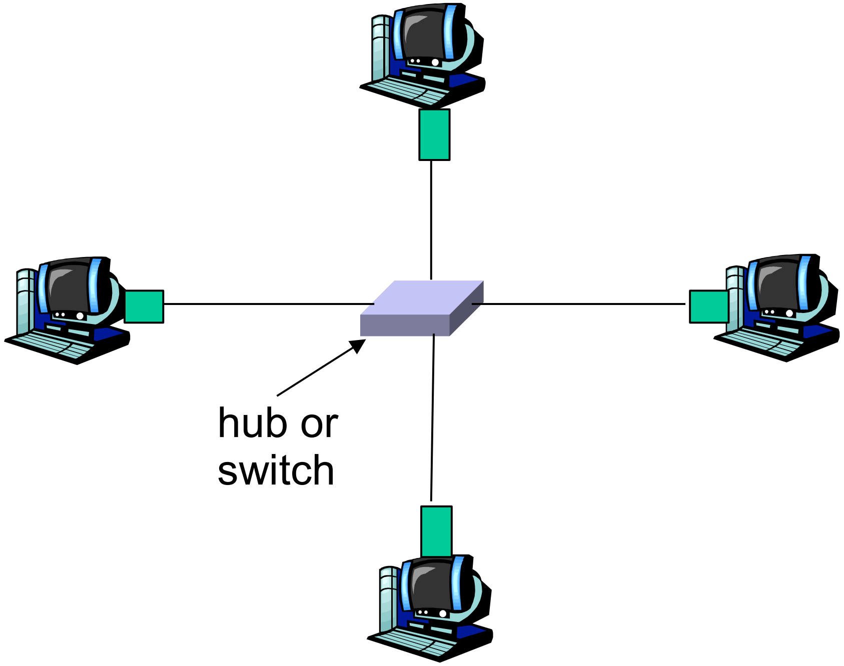

Ethernet Frame Structure

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

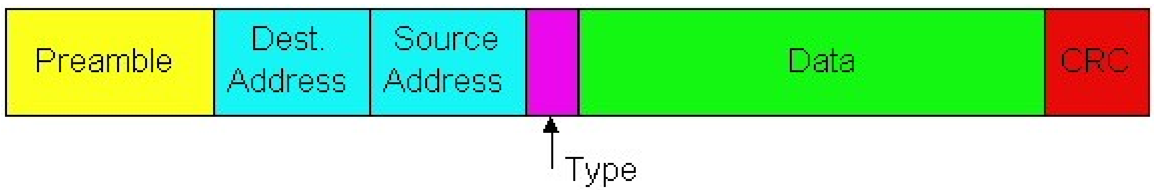

Preamble:

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

Ethernet Frame Structure (more)

- Addresses: 6 bytes
  - if adapter receives frame with matching destination address, or with broadcast address, it passes data in frame to net-layer protocol
  - otherwise, adapter discards frame
- 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

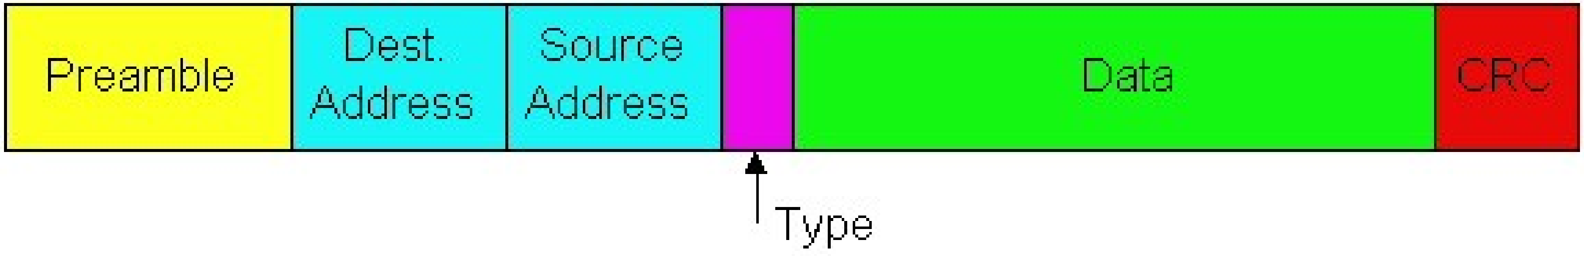

Unreliable, connectionless service

- Connectionless: No handshaking between sending and receiving adapter.
- Unreliable: receiving adapter doesn't send acks or nacks to sending adapter
  - stream of datagrams passed to network layer can have gaps
  - gaps will be filled if app is using TCP
  - otherwise, app will see the gaps
  - receiving adapter may not even exist – no checks in Ethernet

Ethernet uses CSMA/CD

- No slots
- Adapter doesn't transmit if it senses that some other adapter is transmitting, that is, carrier sense
- Transmitting adapter aborts when it senses that another adapter is transmitting, that is, collision detection
- Frames may arrive for transmission at any time, and before attempting a retransmission, adapter waits a random time

Ethernet CSMA/CD algorithm

1. Adaptor receives datagram from net layer & creates frame
2. If adapter senses channel idle, it starts to transmit frame. If it senses channel busy, waits until channel idle and then transmits
3. If adapter transmits entire frame without detecting another transmission, the adapter is done with frame !
4. If adapter detects another transmission while transmitting, aborts and sends jam signal
5. After aborting, adapter enters **exponential backoff**: after the mth collision, adapter chooses a K at random from $\{0,1,2,\ldots,2^m-1\}$. Adapter waits K·512 bit times and returns to Step 2

Ethernet's CSMA/CD (more)

Jam Signal: make sure all other transmitters are aware of collision; 48 bits

Bit time: .1 microsec for 10 Mbps Ethernet ;
for K=1023, wait time is about 50 msec

See/interact with Java applet on UW-ACE: highly recommended !

Exponential Backoff:

- *Goal*: adapt retransmission attempts to estimated current load
  - heavy load: random wait will be longer
- first collision: choose K from {0,1}; delay is K· 512 bit transmission times
- after second collision: choose K from {0,1,2,3}…
- after ten collisions, choose K from {0,1,2,3,4,…,1023}

CSMA/CD efficiency

- t_{prop} = max prop time between 2 nodes in LAN
- t_{trans} = time to transmit max-size frame

$$\text{efficiency} = \frac{1}{1 + 5t_{prop} / t_{trans}}$$

- Efficiency goes to 1 as t_{prop} goes to 0
- Goes to 1 as t_{trans} goes to infinity
- Much better than ALOHA, but still decentralized, simple, and cheap

10BaseT and 100BaseT

- 10/100 Mbps rate; latter called “fast ethernet”
- T stands for Twisted Pair
- Nodes connect to a hub: “star topology”; 100 m max distance between nodes and hub

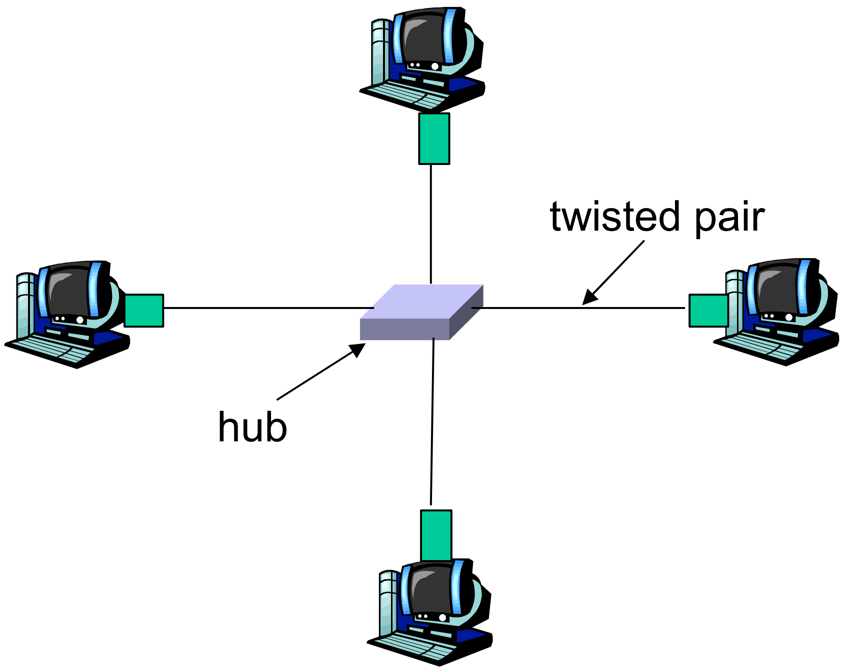

Gbit Ethernet

- Uses 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 required for efficiency
- Uses hubs, called here “Buffered Distributors”
- Full-Duplex at 1 Gbps for point-to-point links
- 10 Gbps now !

Link Layer

- 5.1 Introduction and services
- 5.2 Error detection and correction
- 5.3Multiple access protocols
- 5.4 Link-Layer Addressing
- 5.5 Ethernet
- 5.6 Interconnections: Hubs and switches
- 5.7 PPP
- 5.8 Link Virtualization: ATM

Hubs

Hubs are essentially physical-layer repeaters:

- bits coming from one link go out all other links
- at the same rate
- no frame buffering
- no CSMA/CD at hub: adapters detect collisions

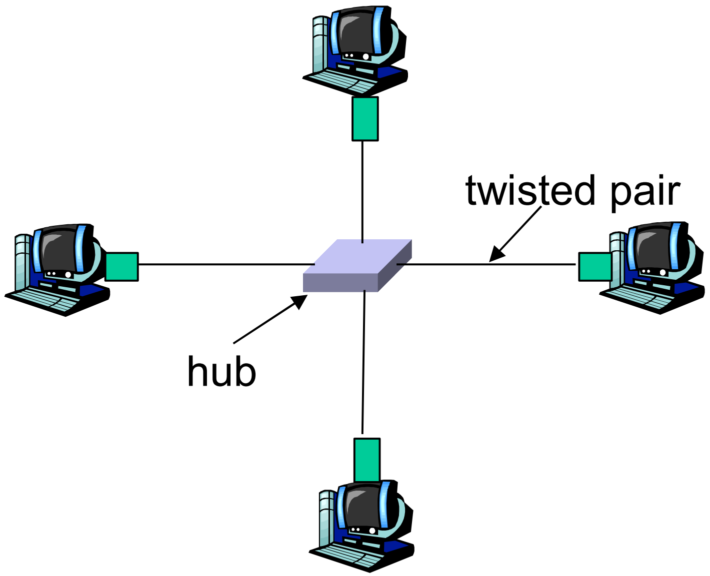

Interconnecting with hubs

- Backbone hub interconnects LAN segments
- Extends maximum distance between nodes
- But individual segment collision domains become one large collision domain
- Can't interconnect 10BaseT & 100BaseT

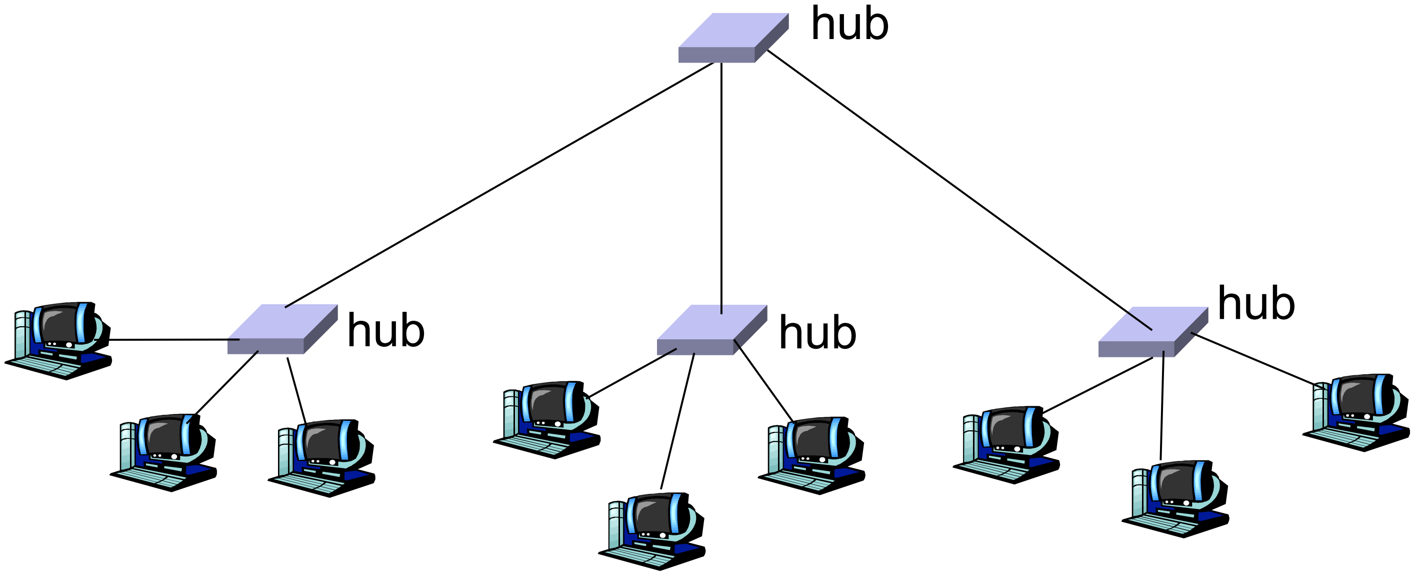

Switches

- Link layer device
  - stores and forwards Ethernet frames
  - examines frame header and selectively forwards frame based on MAC dest address
  - when frame is to be forwarded on segment, uses CSMA/CD to access segment
- Transparent
  - hosts are unaware of presence of switches
- Plug-and-play, self-learning
  - switches do not need to be configured

Forwarding

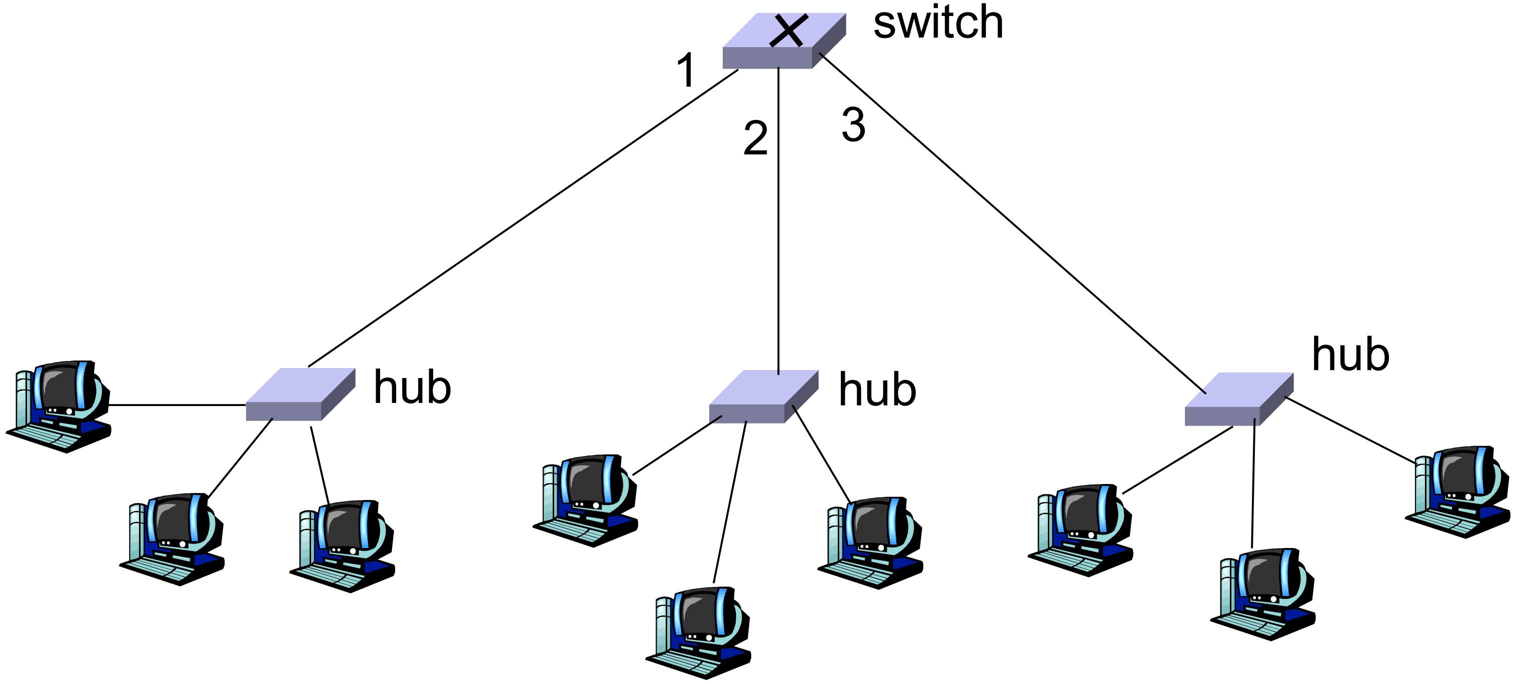

How does the switch determine onto which LAN segment to forward a frame?

Self learning

- A switch has a switch table
- Entries in the switch table look like:
  - (MAC Address, Interface, Time Stamp)
  - stale entries in table dropped (TTL can be 60 min)
- The switch *learns* which hosts can be reached through which interfaces
  - when frame received, switch “learns” location of sender: incoming LAN segment
  - records sender/location pair in switch table

Filtering/Forwarding

When switch receives a frame:

index switch table using MAC dest address

if entry found for destination
  then{
    if dest on segment from which frame arrived
      then drop the frame
      else forward the frame on interface indicated
  }
  else flood

forward on all but the interface on which the frame arrived

Switch example

Suppose C sends frame to D

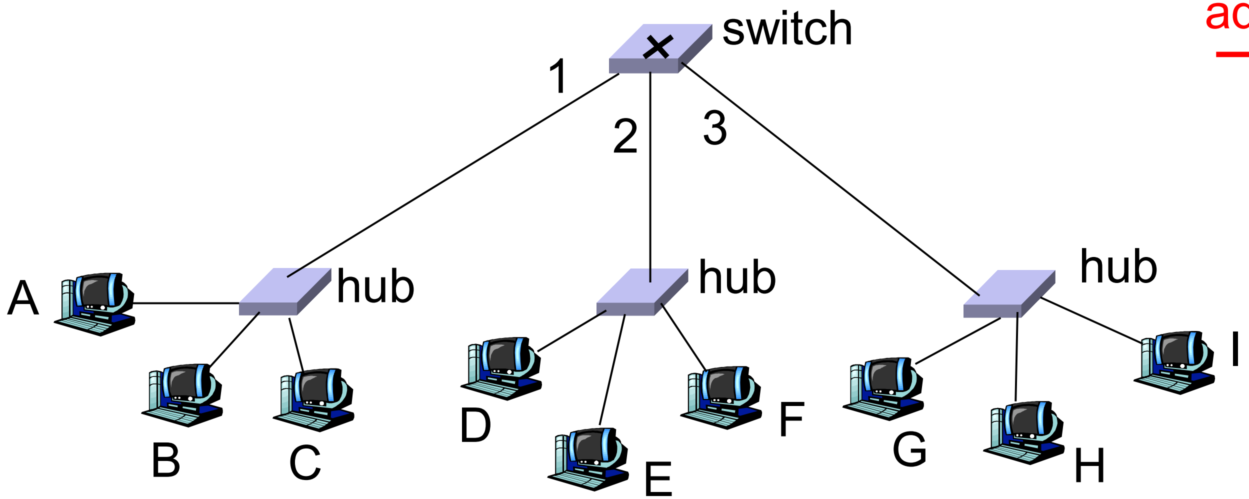

address	interface
A	1
B	1
E	2
G	3

- Switch receives frame from C
  - notes in bridge table that C is on interface 1
  - because D is not in table, switch forwards frame into interfaces 2 and 3
- Frame received by D

Switch example

Suppose D replies back with frame to C.

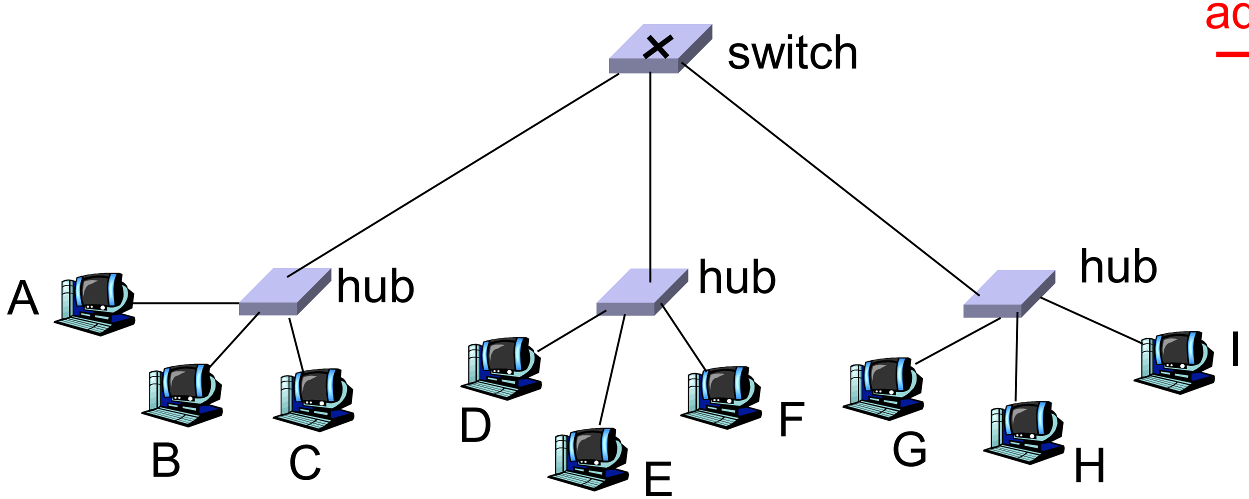

address	interface
A	1
B	1
E	2
G	3
C	1

- Switch receives frame from from D
  - notes in bridge table that D is on interface 2
  - because C is in table, switch forwards frame only to interface 1
- Frame received by C

Switch: traffic isolation

- Switch installation breaks subnet into LAN segments
- Switch filters packets:
  - same-LAN-segment frames not usually forwarded onto other LAN segments
  - segments become separate collision domains

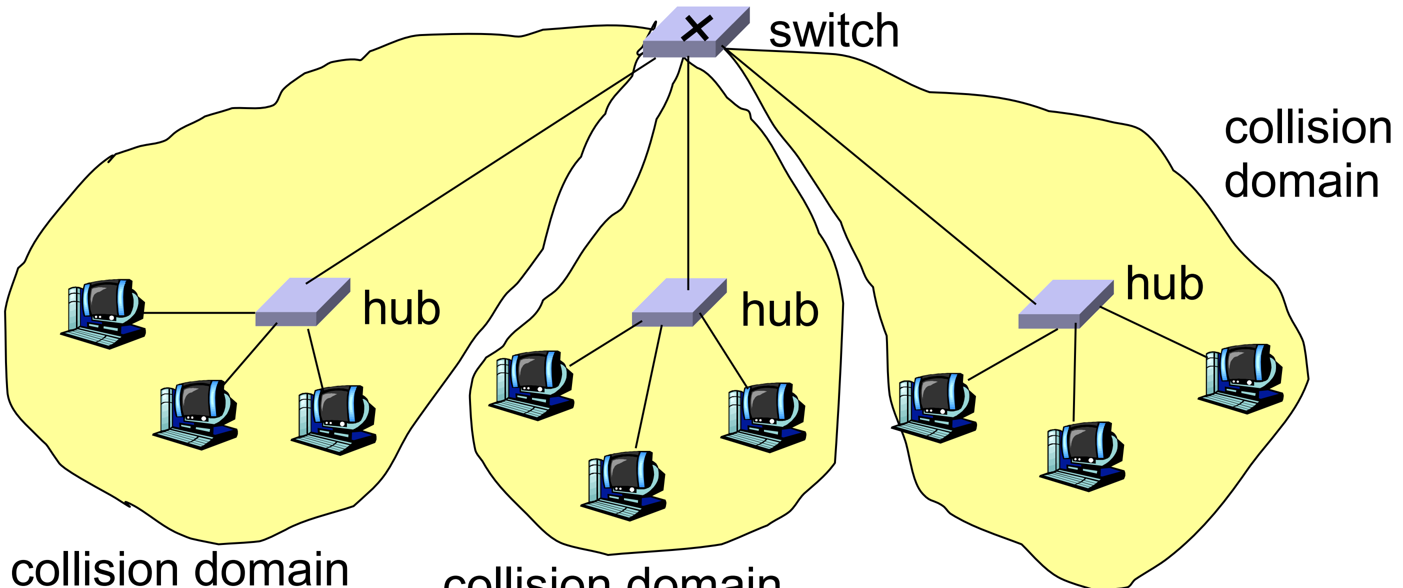

Switches: dedicated access

- Switch with many interfaces
- Hosts have direct connection to switch
- No collisions; full duplex

Switching: A-to-A' and B-to-B' simultaneously, no collisions

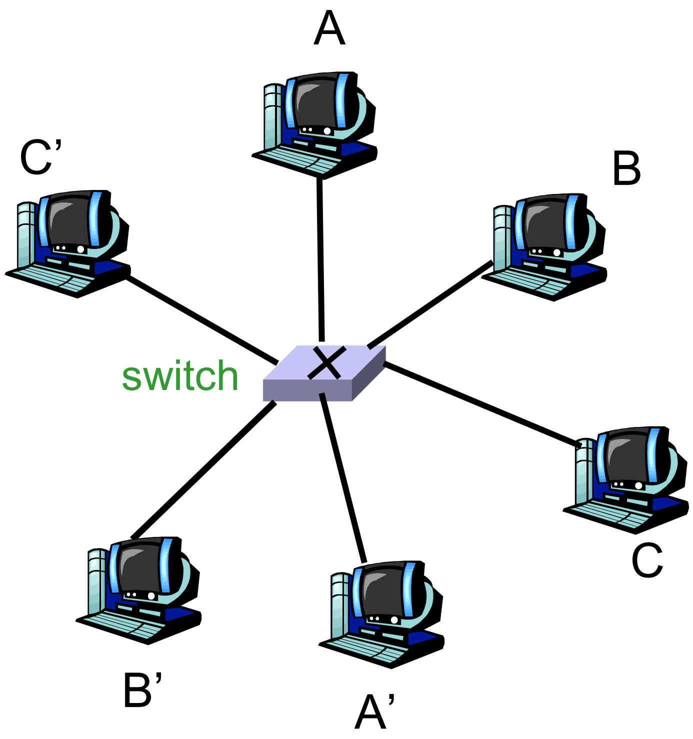

More on Switches

- Cut-through switching: frame forwarded from input to output port without first collecting entire frame
  - slight reduction in latency
- Combinations of shared/dedicated, 10/100/1000 Mbps interfaces

Institutional network

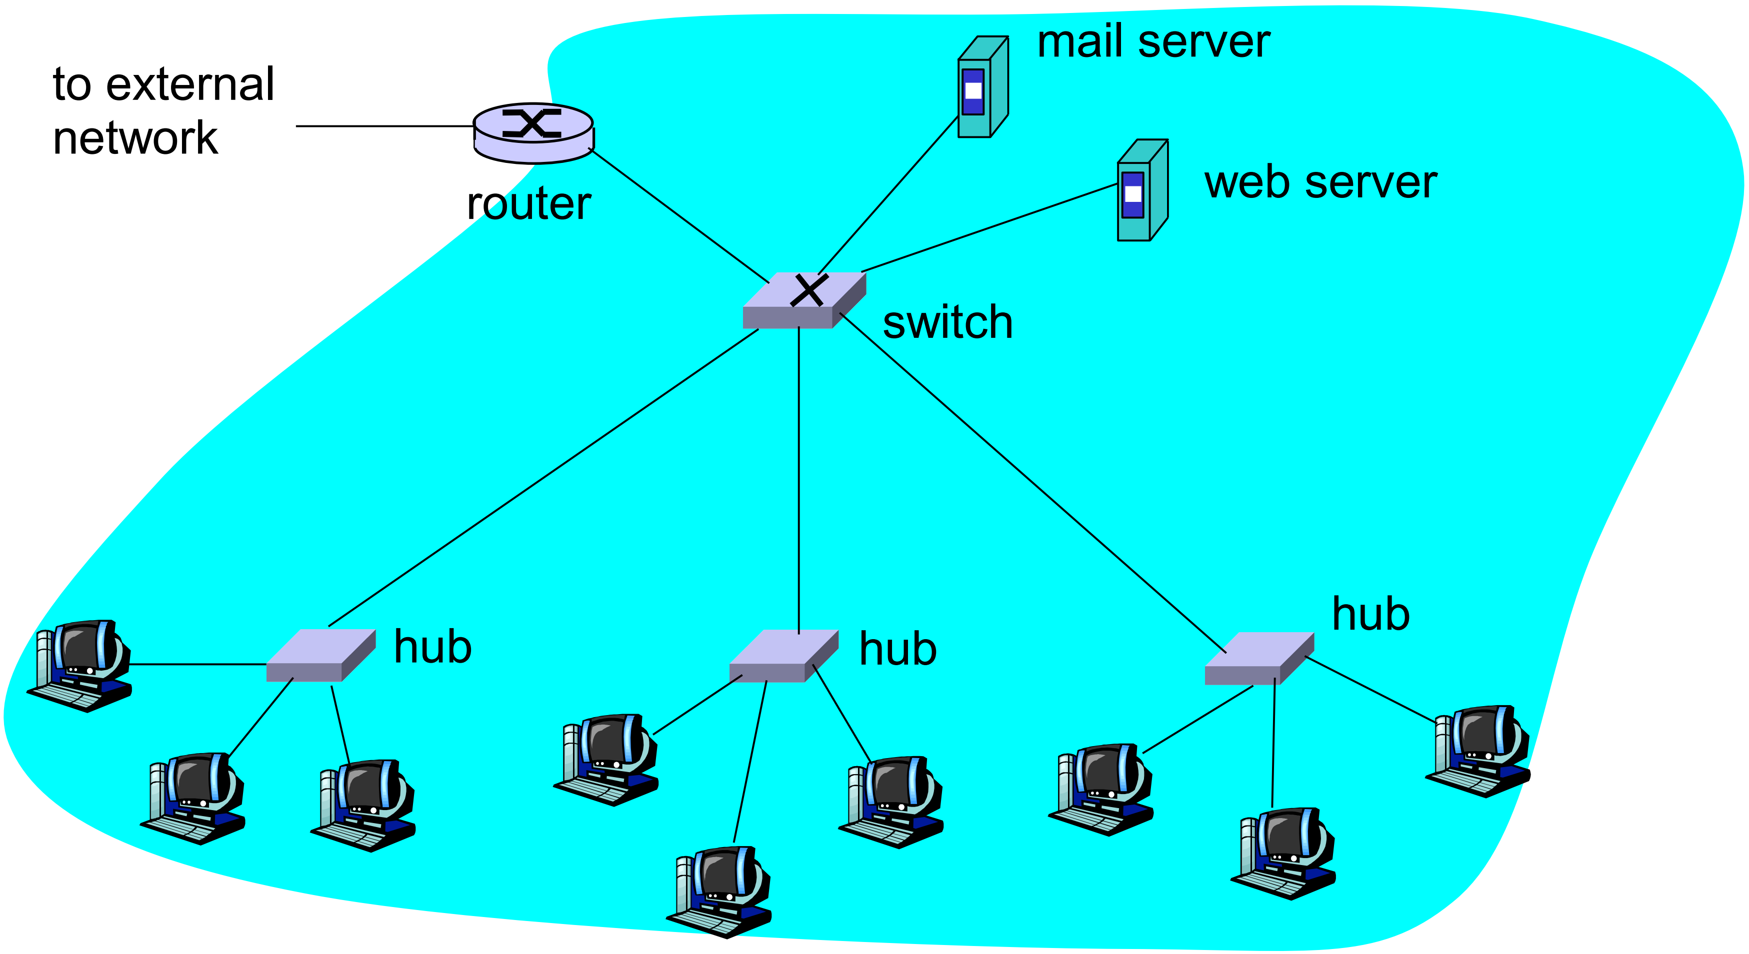

Recap

- Ethernet
  - frame structure
  - CSMA/CD algorithm
- Hubs
  - physical-layer repeaters
  - make one large collision domain
- Switches
  - link-layer devices
  - separates collision domains
  - transparent, plug-and-play, self-learning

Next time

- Wireless link-layer
  - Introduction
  - Characteristics of wireless links
  - 802.11 wireless LANs
  - Cellular Internet access