Washington University in St. Louis

Engineering

6. Ethernet and the Internet

- Classical Ethernet and CSMA/CD
- Modern Ethernet
- Address Resolution Protocol (ARP)
- Multicast addresses

Jon Turner

Classical Ethernet

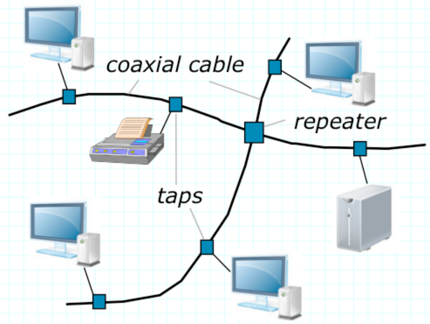

- Passive broadcast bus
  - » 10 Mb/s data rate
  - » datagram packets
  - » hosts identified by "flat" 48 bit addresses
    - assigned by manufacturer
  - » support for unicast, broadcast and multicast
- Access to the bus controlled using CSMA/CD protocol
  - » Collision Sense Multiple Access with Collision Detection
  - » sending host first checks that bus is "idle"
  - » then transmits, while monitoring for simultaneous transmission by other hosts
  - » on collision, back off and try again

More about CSMA/CD

- Signals propagate at finite speed
  - about 200-300 meters per μs
  - taps, repeaters and network adapters also add delay
- Ethernet allows hosts to be fairly far apart
  - 1 km separation means a delay of 3-5 μs
  - and at 10 Mb/s, a host can send 30-50 bits in that time
  - in a network with hundreds of hosts, collisions can easily occur
    - probability of collisions depends very much on traffic load
- Hosts detect collisions by comparing received bits to bits being sent
  - then, send "jamming signal" to ensure that all recognize collision
  - colliding hosts then wait for a random time before trying again
  - on each successive collision, double the random waiting time
    - called *exponential backoff*

Collisions in Classical Ethernet

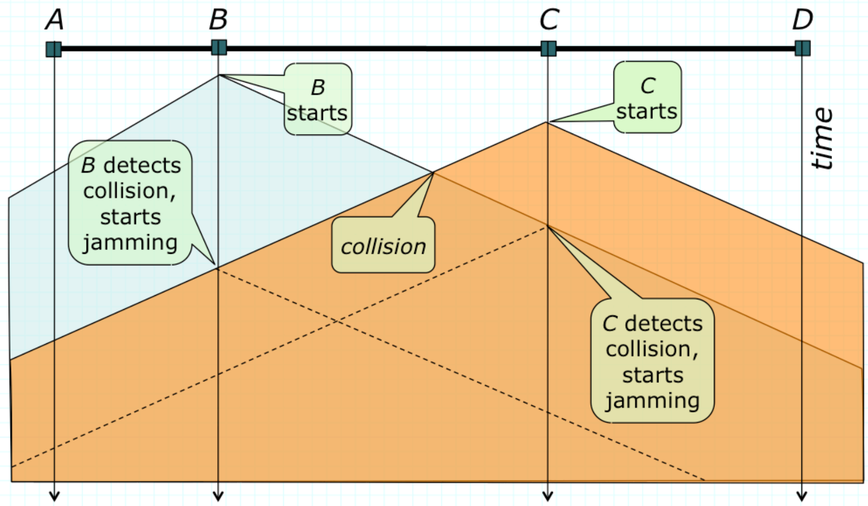

Approximate Efficiency of CSMA/CD

- Approximate efficiency

$$\frac{(d_{trans}/d_{prop})}{5+(d_{trans}/d_{prop})}$$

where d_{trans} is time to transmit a packet and d_{prop} is time for a packet to propagate across net

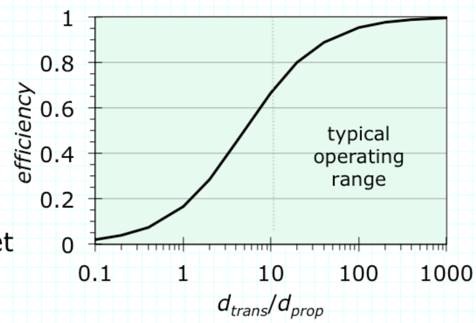

- Ethernet is most efficient when $d_{trans} >> d_{prop}$
  - » for 10 Mb/s Ethernet d_{trans} is at least 50 μs and d_{prop} is generally under 10 μs
- If we increase bit rate, d_{trans} drops proportionally
  - » so to maintain efficiency at higher rates, need to reduce d_{prop}, which requires reducing the maximum distance between nodes

5

Ethernet Frame Format

preamble (7 bytes)
start of frame
destination address
source address
type (2 bytes)
data (46-1500 bytes)
CRC

- *Preamble* and *start of frame* used to synchronize sender, receiver
- Address *fields* identify source and destination – 6 bytes each
  - » globally unique addresses, assigned by manufacturers of network adapters
  - » no location information
    - receiving hosts match dest addr to their own
  - » dest addr of all 1's is *broadcast address*
- *Type* identifies type of payload data
  - » such as IP (0x800) or AppleTalk (0x809B)
- Minimum data field length of 46 bytes
  - » required to prevent undetected collisions
- *Cyclic Redundancy Code* (CRC) is used to detect errors (4 bytes)

Modern Ethernet

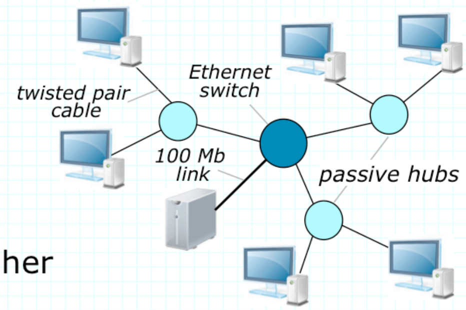

- Passive hubs work like classical Ethernet
  - » allow use of inexpensive twisted pair cable
  - » hub wiring more convenient large office installations
- Ethernet switches offer higher performance
  - » 100 Mb/s, 1 Gb/s, 10 Gb/s available now
  - » switches can send/receive packets on multiple links at same time
    - so, no collisions on switch links
  - » most unicast packets do not need to be broadcast
    - switches "learn" location of different hosts by observing packets
    - still requires use of a tree-structured topology
- Modern Ethernet very different from classical Ethernet
  - » but, still carries the Ethernet "brand-name"

Exercises

1. Suppose hosts *A* and *B* are 400 meters apart in a "classical" 10 Mb/s Ethernet. If *A* starts sending a packet at time 0, when will its first bit reach *B* (assume that signals propagate at roughly 300 m/μs)? If *B* starts sending a packet exactly 1 μs after *A* does, when will the first bit of *A*'s packet collide with the first bit of *B*'s? When does *A* know that a collision has occurred? When does *B* know?
2. Ethernet packets have a minimum payload length of 46 bytes. Explain the impact that a shorter packet length could have on the operation of the CSMA/CD protocol.
3. Estimate the operating efficiency of a classical Ethernet network carrying packets with an average payload length of 200 bytes, and having a maximum propagation delay of 10 μs.
4. Suppose that 2 bits that are 20 bits apart in an Ethernet packet are corrupted while the packet is being transmitted across a LAN. Will the receiver be able to detect these errors using the CRC in all cases? What if the two bits are of 40 bits apart from each other?

Address Resolution Protocol – ARP

- How does one send an IP packet across an Ethernet LAN?
  - sender knows only the *IP address* of the destination host
  - but, needs Ethernet address (aka MAC addr) of destination host
- ARP protocol serves as glue between IP and L2 networks
- Basic operation
  - sending host sends an *ARP request* with target's IP address in an Ethernet *broadcast frame*, with type=0x806
  - all hosts in subnet receive frame and deliver to their ARP protocol module
  - target host responds with reply packet containing its MAC addr
  - now sending host sends its IP packet using provided MAC addr
  - also, saves MAC address in an *ARP table* so that later packets to same destination can avoid lookup
    - ARP table entries eventually timeout, allowing IP addresses to change

Delivering Packets Between Subnets

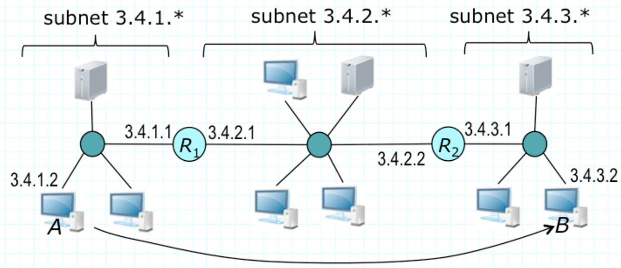

- Host *A* sending to *B*; first sends packet to router R_1, since destination is on different subnet
  - » may need to ARP for R_1's MAC address
- Router R_1's routing table specifies router R_2 as next hop
  - » so, R_1 sends to R_2, possibly using ARP to get R_2's MAC address
- Router R_2's routing table has no specific next-hop IP, so packet sent directly to 3.4.3.2 (possibly after ARP)

Exercises

1. In the example on page 10, assume that initially, all hosts and routers have empty ARP tables. How many packets are sent on each of the three subnets before the packet from *A* reaches *B*? What are the contents of the ARP tables at each of the hosts and routers after the packet reaches *B*?
2. Consult RFC 826. What is the length of an ARP packet used to discover the Ethernet address associated with a given IP address? What value should appear in the ar$pro field? What value should appear in the ar$op field? What value appears in the destination address field of the Ethernet packet? How can you tell the difference between an ARP request and an ARP reply?

Washington University in St. Louis — Engineering

Multicast in Ethernet

- Ethernet addresses in which the first eight bits form an odd number are called *multicast addresses*
- Hosts can configure their network adapters to accept packets with specific multicast addresses
- Multicast allows a host to send data to many receivers without having to send data many times
  - » can be used for real-time video delivery and group collaboration applications
  - » used for some Ethernet and IP control protocols
- IP multicast packets are sent in Ethernet frames with multicast addresses
  - » the low order 23 bits of IP destination address are copied into low order 23 bits of Ethernet destination address
  - » high order bits of Ethernet address are set to x01:00:5E

Internet Group Management Protocol

- IGMP allows hosts within a subnet to communicate using multicast
- Three message types (IGMPv2)
  - » membership query
    - broadcast by router to learn which multicast groups are "of interest" to hosts in this subnet (allowing router to decide which multicast packets to forward to subnet)
  - » membership report
    - sent by a host identifying a multicast address for which it wants to receive packets
  - » leave group message
    - sent by host to drop out of a group
  - » IGMPv3 allows source-specific multicast delivery
- IGMP snooping – Ethernet switches often observe IGMP messages & use them to improve efficiency of multicast

Multicast in Small Networks

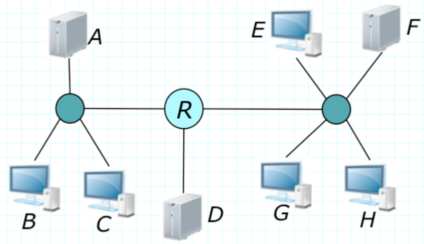

- Router acts as querier for all three networks
  - » broadcasts query and hosts respond with membership packets
  - » suppose *A*, *C*, *D*, *F*, *G* express interest in multicast group *X*, while *C*, *E* and *G* express interest in group *Y*
- Suppose *B* sends packet with multicast address *X*
  - » left switch delivers to *A*, *C* and *R*
  - » router forwards to right switch and *D*
  - » right switch delivers to *F* and *G* (assuming IGMP snooping)

14

Washington University in St. Louis | Engineering

Exercises

Refer to RFCs 2236 and 3376.

1. IGMP specifies two kinds of query messages, a group-specific query and a general query. Explain the difference.
2. In IGMP, one router on each subnet plays the role of "Querier". How does the protocol ensure that there is only one Querier? What does a Querier do that other routers don't do? Why do you think it's important that there be just one querier?
3. When a host gets a general query, it responds with a membership report for each of the groups it belongs to, but it waits a random amount of time before sending each response. Why do you think this is done?
4. In IGMPv3, how does a host report that it wants to receive packets from all senders in a given multicast group?