Links

Mike Freedman

COS 461: Computer Networks

Lectures: MW 10-10:50am in CS 104

http://www.cs.princeton.edu/courses/archive/spr14/cos461/

Protocol Layers

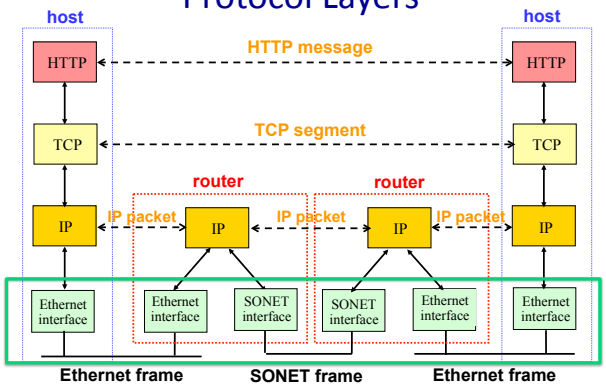

2

Link = Medium + Adapters

3

What is a Link?

Communication Medium

Network Adapter

Broadcast Links: Shared Media

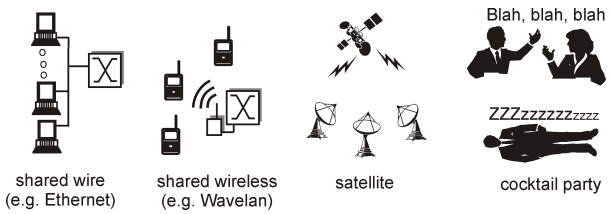

5

Adaptors Communicating

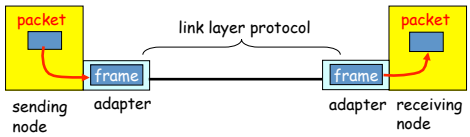

- Sending side
  - Encapsulates packet in a frame
  - Adds error checking bits, flow control, etc.
- Receiving side
  - Looks for errors, flow control, etc.
  - Extracts datagram and passes to receiving node

6

Link-Layer Services

- Encoding
  - Represent the 0s and 1s
- Framing
  - Encapsulate packet into frame, adding header/trailer
- Error detection
  - Receiver detecting errors with checksums
- Error correction
  - Receiver optionally correcting errors
- Flow control
  - Pacing between sending and receiving nodes

7

Addresses

8

Medium Access Control Address

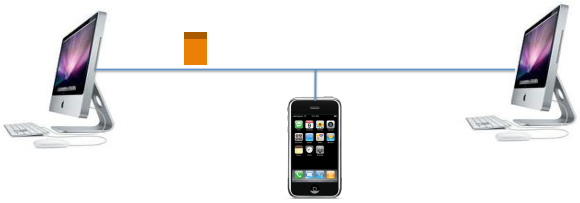

- Identify the sending and receiving adapter
  - Unique identifier for each network adapter
  - Identifies the intended receiver(s) of the frame
  - … and the sender who sent the frame

9

Medium Access Control Address

- MAC address (e.g., 00-15-C5-49-04-A9)
  - Numerical address used within a link
  - Unique, hard-coded in the adapter when it is built
  - Flat name space of 48 bits
- Hierarchical allocation: Global uniqueness!
  - Blocks: assigned to vendors (e.g., Dell) by the IEEE
  - Adapters: assigned by the vendor from its block
- Broadcast address (i.e., FF-FF-FF-FF-FF-FF)
  - Send the frame to *all* adapters

10

As an Aside: Promiscuous Mode

- Normal adapter: receives frames sent to
  - The local MAC address
  - Broadcast address FF-FF-FF-FF-FF-FF
- Promiscuous mode
  - Receive *everything*, independent of destination MAC
- Useful for packet sniffing
  - Network monitoring
  - E.g., wireshark, tcpdump

11

Why Not Just Use IP Addresses?

- Links can support *any* network protocol
  - Not just for IP (e.g., IPX, Appletalk, X.25, …)
  - Different addresses on different kinds of links
- An adapter may move to a new location
  - So, cannot simply assign a static IP address
  - Instead, must reconfigure the adapter's IP address
- Must identify the adapter during bootstrap
  - Need to talk to the adapter to assign it an IP address

12

Who Am I: Acquiring an IP Address

- Dynamic Host Configuration Protocol (DHCP)
  - Broadcast "I need an IP address, please!"
  - Response "You can have IP address 1.2.3.4."

Who Are You: Discovering the Receiver

- Address Resolution Protocol (ARP)
  - Broadcast "who has IP address 1.2.3.6?"
  - Response "0C-C4-11-6F-E3-98 has 1.2.3.6!"

Sharing the Medium

Collisions

- Single shared broadcast channel
  - Avoid having multiple nodes speaking at once
  - Otherwise, collisions lead to garbled data

Multi-Access Protocol

- Divide the channel into pieces
  - In time
  - In frequency

- Take turns
  - Pass a token for the right to transmit

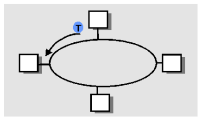

- Punt
  - Let collisions happen
  - … and detect and recover from them

17

Like Human Conversation…

- Carrier sense
  - Listen before speaking
  - …and don't interrupt!

- Collision detection
  - Detect simultaneous talking
  - … and shut up!

- Random access
  - Wait for a random period of time
  - … before trying to talk again!

18

Carrier Sense Multiple Access

- Listen for other senders
  - Then transmit your data
- Collisions can still occur
  - Propagation delay
  - Wasted transmission

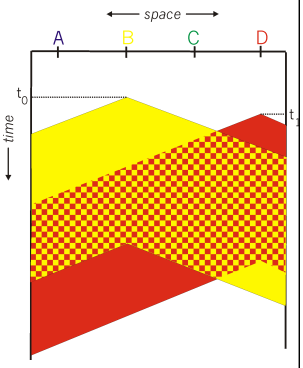

19

CSMA/CD Collision Detection

- Detect collision
  - Abort transmission
  - Jam the link
- Wait random time
  - Transmit again
- Hard in wireless
  - Must receive data while transmitting

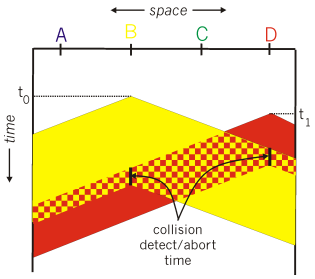

20

Comparing the Three Approaches

- Channel partitioning is
  - (a) Efficient/fair at high load, inefficient at low load
  - (b) Inefficient at high load, efficient/fair at low load
- "Taking turns"
  - (a) Inefficient at high load
  - (b) Efficient at all loads
  - (c) Robust to failures
- Random access
  - (a) Inefficient at low load
  - (b) Efficient at all load
  - (c) Robust to failures

21

Comparing the Three Approaches

- Channel partitioning is
  - (a) Efficient/fair at high load, inefficient at low load
  - (b) Inefficient at high load, efficient/fair at low load
- "Taking turns"
  - (a) Inefficient at high load
  - (b) Efficient at all loads
  - (c) Robust to failures
- Random access
  - (a) Inefficient at low load
  - (b) Efficient at all load
  - (c) Robust to failures

22

Ethernet

23

Ethernet

- Dominant wired LAN technology
- First widely used LAN technology
- Kept up with speed race: 10 Mbps – 40 Gbps

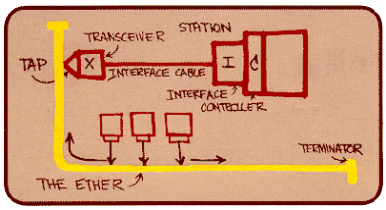

Metcalfe's Ethernet sketch

24

Ethernet Uses CSMA/CD

- Carrier Sense: wait for link to be idle
  - Channel idle: start transmitting
  - Channel busy: wait until idle
- Collision Detection: listen while transmitting
  - No collision: transmission is complete
  - Collision: abort transmission, and send jam signal
- Random Access: exponential back-off
  - After collision, wait random time before trying again
  - After m^{th} collision, choose K randomly from {0, …, 2^m-1}
  - … and wait for K*512 bit times before trying again

25

Limitations on Ethernet Length

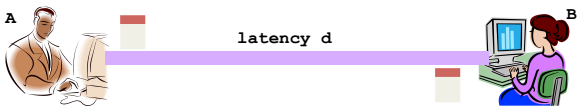

- Latency depends on physical length of link
  - Time to propagate a packet from one end to other
- Suppose A sends a packet at time t
  - And B sees an idle line at a time just before t+d
  - … so B happily starts transmitting a packet
- B detects a collision, and sends jamming signal
  - But A doesn't see collision till t+2d

26

Limitations on Ethernet Length

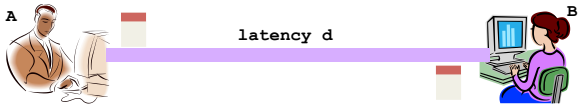

- A needs to wait for time 2d to detect collision
  - So, A should keep transmitting during this period
  - … and keep an eye out for a possible collision
- Imposes restrictions on Ethernet
  - Maximum length of the wire: 2500 meters
  - Minimum length of the packet: 512 bits (64 bytes)

27

Ethernet Frame Structure

- Sending adapter encapsulates packet in frame

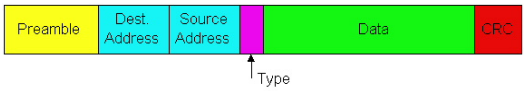

- Preamble: synchronization
  - Seven bytes with pattern 10101010, followed by one byte with pattern 10101011
  - Used to synchronize receiver, sender clock rates

28

Ethernet Frame Structure

- Addresses: source and destination MAC addresses
  - Adaptor passes frame to network-level protocol
    - If destination is local MAC address or broadcast address
  - Otherwise, adapter discards frame
- Type: indicates the higher layer protocol
  - Usually IP
  - But also Novell IPX, AppleTalk, …
- CRC: cyclic redundancy check

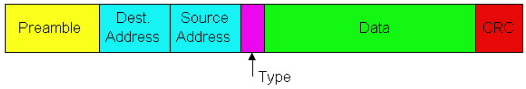

Unreliable, Connectionless Service

- Connectionless
  - No handshaking between send and receive adapter
- Unreliable
  - Receiving adapter doesn't send ACKs or NACKs
  - Packets passed to network layer can have gaps
  - Gaps can be filled by transport protocol (e.g., TCP)
  - Otherwise, the application will see the gaps

Hubs and Switches

Physical Layer: Repeaters

- Distance limitation in local-area networks
  - Electrical signal becomes weaker as it travels
  - Imposes a limit on the length of a LAN
- Repeaters join LANs together
  - Analog electronic device
  - Continuously monitors electrical signals
  - Transmits an amplified copy

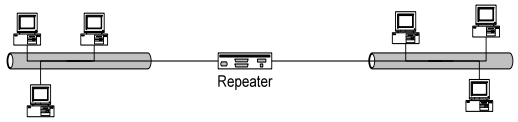

Physical Layer: Hubs

- Joins multiple input lines electrically
  - Designed to hold multiple line cards
  - Do not necessarily amplify the signal
- Very similar to repeaters
  - Also operates at the physical layer

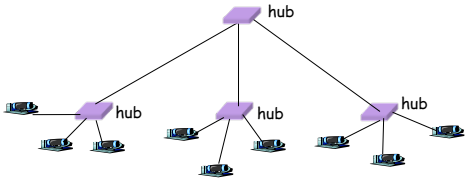

33

Limitations of Repeaters and Hubs

- One large shared link
  - Each bit is sent everywhere
  - So, aggregate throughput is limited
- Cannot support multiple LAN technologies
  - Does not buffer or interpret frames
  - Can't interconnect between different rates/formats
- Limitations on maximum nodes and distances
  - Shared medium imposes length limits
  - E.g., cannot go beyond 2500 meters on Ethernet

34

Link Layer: Bridges

- Connects two or more LANs at the link layer
  - Extracts destination address from the frame
  - Looks up the destination in a table
  - Forwards the frame to the appropriate segment
- Each segment can carry its own traffic

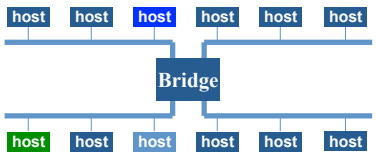

35

Link Layer: Switches

- Typically connects individual computers
  - A switch is essentially the same as a bridge
  - … though typically used to connect hosts
- Supports concurrent communication
  - Host A can talk to C, while B talks to D

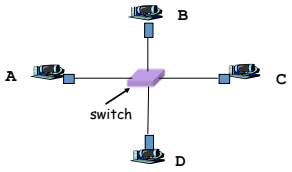

36

Bridges/Switches: Traffic Isolation

- Switch filters packets
  - Frame only forwarded to the necessary segments
  - Segments can support separate transmissions

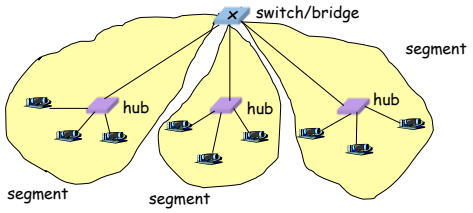

37

Switches vs. Hubs

- Compared to hubs, Ethernet switches support
  (a) Larger geographic span
  (b) Similar span
  (c) Smaller span
- Compared to hubs, switches provides
  (a) Higher load on links
  (b) Less privacy
  (c) Heterogenous communication technologies

38

Switches vs. Hubs

- Compared to hubs, Ethernet switches support
  (a) Larger geographic span
  (b) Similar span
  (c) Smaller span
- Compared to hubs, switches provides
  (a) Higher load on links
  (b) Less privacy
  (c) Heterogenous communication technologies

39

Self Learning: Building the Table

- When a frame arrives
  - Inspect the *source* MAC address
  - Associate the address with the *incoming* interface
  - Store the mapping in the switch table
  - Use a timer to eventually forget the mapping

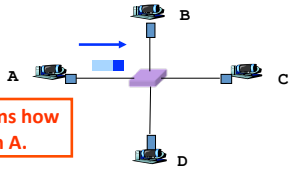

Switch learns how to reach A.

40

Self Learning: Handling Misses

- When frame arrives with unfamiliar destination
  - Forward the frame out all of the interfaces
  - … except for the one where the frame arrived
  - Hopefully, this case won't happen very often!

When in doubt, shout!

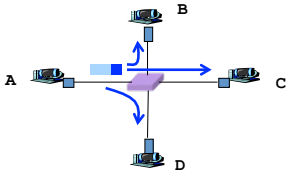

41

Summary: Multiple Layers

- Different devices switch different things
  - Network layer: packets (routers)
  - Link layer: frames (bridges and switches)
  - Physical layer: electrical signals (repeaters and hubs)

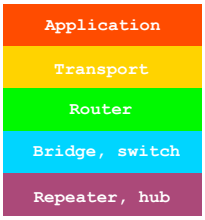

42

Conclusion

- Links
  - Connect two or more network adapters
  - … each with a unique address
  - … over a shared communication medium
- Coming next
  - Friday: Socket Programming "How To"
  - Monday: Network layer (IP)
- Get started
  - On assignment #0 on socket programming

43