



Links

Reading: Chapter 2

CS 375: Computer Networks
Thomas Bressoud


1

Goals of Today's Lecture

- **Link-layer services**
 - Encoding, framing, and error detection
 - Error correction and flow control
- **Sharing a shared media**
 - Channel partitioning
 - Taking turns
 - Random access
- **Ethernet protocol**
 - Carrier sense, collision detection, and random access
 - Frame structure
 - Hubs and switches

2

Message, Segment, Packet, and Frame




Link Layer Protocol for Each Hop

- IP packet transferred over multiple hops
 - Each hop has a link layer protocol
 - May be different on different hops
- Analogy: trip from Denison to Lausanne
 - Shuttle: Denison to CMH
 - Plane: CMH to LaGuardia; LaGuardia to Geneva
 - Train: Geneva to Lausanne
- Refining the analogy
 - Tourist == packet
 - Transport segment == communication link
 - Transportation mode == link-layer protocol
 - Travel agent == routing algorithm

4

Adaptors Communicating



- Link layer implemented in *adapter* (network interface card)
 - Ethernet card, PCMCIA card, 802.11 card
- Sending side:
 - Encapsulates datagram 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

5


Link-Layer Services

- Encoding
 - Representing the 0s and 1s
- Framing
 - Encapsulating packet into frame, adding header, trailer
 - Using MAC addresses, rather than IP addresses
- Error detection
 - Errors caused by signal attenuation, noise.
 - Receiver detecting presence of errors
- Error correction
 - Receiver correcting errors without retransmission
- Flow control
 - Pacing between adjacent sending and receiving nodes

6

Encoding

- Signals propagate over physical links
 - Source node encodes the bits into a signal
 - Receiving node decodes the signal back into bits
- Simplify some electrical engineering details
 - Assume two discrete signals, high and low
 - E.g., could correspond to two different voltages
- Simple approach
 - High for a 1, low for a 0



Problem With Simple Approach

- Long strings of 0s or 1s introduce problems
 - No transitions from low-to-high, or high-to-low
- Receiver keeps average of signal it has received
 - Uses the average to distinguish between high and low
 - Long flat strings make receiver sensitive to small change
- Transitions also necessary for clock recovery
 - Receiver uses transitions to drive its own clock
 - Long flat strings do not produce any transitions
 - Can lead to clock drift at the receiver
- Alternatives (see Section 2.2)
 - Non-return to zero inverted, and Manchester encoding

8

Framing

- Break sequence of bits into a frame
 - Typically implemented by the network adaptor
- Sentinel-based
 - Delineate frame with special pattern (e.g., 01111110)

01111110	Frame contents	01111110
----------	----------------	----------

- Problem: what if special patterns occurs within frame?
- Solution: escaping the special characters
 - E.g., sender always inserts a 0 after five 1s
 - ... and receiver always removes a 0 appearing after five 1s
- Similar to escaping special characters in C programs

9

Framing (Continued)

- Counter-based
 - Include the payload length in the header
 - ... instead of putting a sentinel at the end
 - Problem: what if the count field gets corrupted?
 - Causes receiver to think the frame ends at a different place
 - Solution: catch later when doing error detection
 - And wait for the next sentinel for the start of a new frame
- Clock-based
 - Make each frame a fixed size
 - No ambiguity about start and end of frame
 - But, may be wasteful

10

Error Detection

- Errors are unavoidable
 - Electrical interference, thermal noise, etc.
- Error detection
 - Transmit extra (redundant) information
 - Use redundant information to detect errors
 - Extreme case: send two copies of the data
 - Trade-off: accuracy vs. overhead
- Techniques for detecting errors
 - Parity checking
 - Checksum
 - Cyclic Redundancy Check (CRC)

11


Error Detection Techniques

- Parity check
 - Add an extra bit to a 7-bit code
 - Odd parity: ensure an odd number of 1s
 - E.g., 0101011 becomes 01010111
 - Even parity: ensure an even number of 1s
 - E.g., 0101011 becomes 01010110
- Checksum
 - Treat data as a sequence of 16-bit words
 - Compute a sum of all the 16-bit words, with no carries
 - Transmit the sum along with the packet
- Cyclic Redundancy Check (CRC)
 - See Section 2.4.3

12

Point-to-Point vs. Broadcast Media

- Point-to-point
 - PPP for dial-up access
 - Point-to-point link between Ethernet switch and host
- Broadcast (shared wire or medium)
 - Traditional Ethernet
 - 802.11 wireless LAN



13

Multiple Access Protocol

- Single shared broadcast channel
 - Avoid having multiple nodes speaking at once
 - Otherwise, collisions lead to garbled data
- Multiple access protocol
 - Distributed algorithm for sharing the channel
 - Algorithm determines which node can transmit
- Classes of techniques
 - Channel partitioning: divide channel into pieces
 - Taking turns: passing a token for the right to transmit
 - Random access: allow collisions, and then recover

14

Channel Partitioning: TDMA

TDMA: time division multiple access

- Access to channel in "rounds"
 - Each station gets fixed length slot in each round
- Time-slot length is packet transmission time
 - Unused slots go idle
- Example: 6-station LAN with slots 1, 3, and 4




15

Channel Partitioning: FDMA

FDMA: frequency division multiple access

- Channel spectrum divided into frequency bands
 - Each station assigned fixed frequency band
- Unused transmission time in bands go idle
- Example: 6-station LAN with bands 1, 3, and 4



16


“Taking Turns” MAC protocols

Polling

- Master node “invites” slave nodes to transmit in turn
- Concerns:
 - Polling overhead
 - Latency
 - Single point of failure (master)

Token passing

- Control token passed from one node to next sequentially
- Token message
- Concerns:
 - Token overhead
 - Latency
 - Single point of failure (token)



17

Random Access Protocols

- When node has packet to send
 - Transmit at full channel data rate R .
 - No a priori coordination among nodes
- Two or more transmitting nodes → “collision”,
- Random access MAC protocol specifies:
 - How to detect collisions
 - How to recover from collisions
- Examples
 - ALOHA and Slotted ALOHA
 - CSMA, CSMA/CD, CSMA/CA

18

Key Ideas of Random Access

- **Carrier Sense (CS)**
 - Listen before speaking, and don't interrupt
 - Checking if someone else is already sending data
 - ... and waiting till the other node is done
- **Collision Detection (CD)**
 - If someone else starts talking at the same time, stop
 - Realizing when two nodes are transmitting at once
 - ...by detecting that the data on the wire is garbled
- **Randomness**
 - Don't start talking again right away
 - Waiting for a random time before trying again

19

Slotted ALOHA

Assumptions


- All frames same size
- Time divided into equal slots (time to transmit a frame)
- Nodes start to transmit frames only at start of slots
- Nodes are synchronized
- If two or more nodes transmit, all nodes detect collision

Operation

- When node obtains fresh frame, transmits in next slot
- No collision: node can send new frame in next slot
- Collision: node retransmits frame in each subsequent slot with probability p until success

20

Slotted ALOHA



Pros

- Single active node can continuously transmit at full rate of channel
- Highly decentralized: only slots in nodes need to be in sync
- Simple

Cons

- Collisions, wasting slots
- Idle slots
- Nodes may be able to detect collision in less than time to transmit packet
- Clock synchronization

21

CSMA (Carrier Sense Multiple Access)


- CSMA: listen before transmit
 - If channel sensed idle: transmit entire frame
 - If channel sensed busy, defer transmission
- Human analogy: don't interrupt others!

22

CSMA Collisions

Collisions *can* still occur:
propagation delay means
two nodes may not hear
each other's transmission

Collision:
entire packet transmission
time wasted




23

CSMA/CD (Collision Detection)

- CSMA/CD: carrier sensing, deferral as in CSMA
 - Collisions detected within short time
 - Colliding transmissions aborted, reducing wastage
- Collision detection
 - Easy in wired LANs: measure signal strengths, compare transmitted, received signals
 - Difficult in wireless LANs: receiver shut off while transmitting
- Human analogy: the polite conversationalist

24

CSMA/CD Collision Detection



25


Three Ways to Share the Media

- Channel partitioning MAC protocols:
 - Share channel efficiently and fairly at high load
 - Inefficient at low load: delay in channel access, 1/N bandwidth allocated even if only 1 active node!
- “Taking turns” protocols
 - Eliminates empty slots without causing collisions
 - Vulnerable to failures (e.g., failed node or lost token)
- Random access MAC protocols
 - Efficient at low load: single node can fully utilize channel
 - High load: collision overhead

26

Ethernet

- Dominant wired LAN technology:
- First widely used LAN technology
- Simpler, cheaper than token LANs and ATM
- Kept up with speed race: 10 Mbps – 10 Gbps



Metcalf's Ethernet sketch


27

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 a random time before trying again
 - After m^{th} collision, choose K randomly from $\{0, \dots, 2^m-1\}$
 - ... and wait for $K \cdot 512$ bit times before trying again

28


Limitations on Ethernet Length



- **Latency depends on physical length of link**
 - Time to propagate a packet from one end to the 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 $t+2d$

29

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)

30

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

31

Ethernet Frame Structure (Continued)

- **Addresses:** source and destination MAC addresses
 - Adaptor passes frame to network-level protocol
 - If destination address matches the adaptor
 - Or the destination address is the broadcast address
 - Otherwise, adaptor discards frame
- **Type:** indicates the higher layer protocol
 - Usually IP
 - But also Novell IPX, AppleTalk, ...
- **CRC:** cyclic redundancy check
 - Checked at receiver
 - If error is detected, the frame is simply dropped



32


Unreliable, Connectionless Service

- **Connectionless**
 - No handshaking between sending and receiving adapter.
- **Unreliable**
 - Receiving adapter doesn't send ACKs or NACKs
 - Packets passed to network layer can have gaps
 - Gaps will be filled if application is using TCP
 - Otherwise, the application will see the gaps

33

Hubs: Physical-Layer Repeaters


- Hubs are physical-layer repeaters
 - Bits coming from one link go out all other links
 - At the same rate, with no frame buffering
 - No CSMA/CD at hub: adapters detect collisions



34

Interconnecting with Hubs

- Backbone hub interconnects LAN segments
- All packets seen everywhere, forming one large collision domain
- Can't interconnect Ethernets of different speeds



35


Switch

- 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

36

Switch: Traffic Isolation

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



37

Benefits of Ethernet

- Easy to administer and maintain
- Inexpensive
- Increasingly higher speed
- Moved from shared media to switches
 - Change everything except the frame format
 - A good general lesson for evolving the Internet

38

Conclusions

- IP runs on a variety of link layer technologies
 - Point-to-point links vs. shared media
 - Wide varieties within each class
- Link layer performs key services
 - Encoding, framing, and error detection
 - Optionally error correction and flow control
- Shared media introduce interesting challenges
 - Decentralized control over resource sharing
 - Partitioned channel, taking turns, and random access
 - Ethernet as a wildly popular example

39
