CSE/EE 461 Lecture 4
Media Access

Tom Anderson
tom@cs.washington.edu

Recap:
Error Detection at Multiple Layers

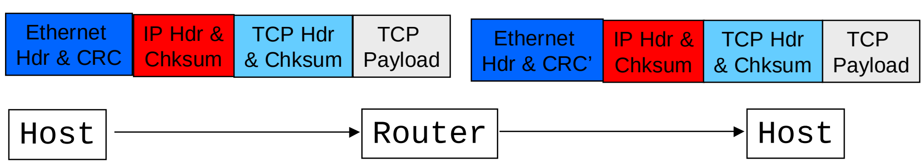

- Ethernet CRC recomputed on each hop
- TCP/IP checksums are “end to end”
  - IP chksum covers the IP header only
  - TCP chksum covers the TCP header, payload
- Each can catch errors the other misses
- Neither work against malicious spoofing

Aloha to Ethernet to wireless

How do multiple parties share access to a communication channel (wire or wireless)?

- Delivery: when packet is broadcast, how does the receiver know intended destination?
  - put destination address in frame header
    - ex: globally unique Ethernet MAC address
  - discard if not intended target
- Arbitration: how do we decide who sends next?

Multiplexing Options

- Frequency division multiple access (FDMA)
  - everyone assigned a different frequency (wavelength)

- Time division multiple access (TDMA)
  - everyone assigned a different time slot

- Code division multiple access (CDMA)
  - multiple senders at a time in each frequency
  - each sender has unique code (1010 vs. 0101 vs. 1100)

Statistical multiplexing

- Assign frequency/time slot on demand
- When is this more efficient than static partitioning? when less efficient?
- How do we arbitrate access?
  - central vs. distributed control?

Digression: Packet Latency

- How long does it take to send a packet?

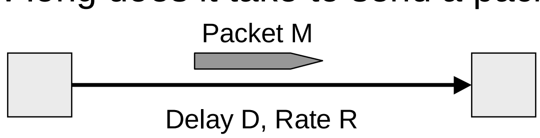

- Two terms:
  - Propagation delay = distance / speed of light in media
  - Transmission delay = message (bits) / rate (bps)
    - – packet takes up space on wire!
- Slow links stretch bits out in time/space

One-way latency examples

- Either a slow link or long wire makes for large latency
- Dialup with a modem:
  - D = 10ms (say), R = 56Kbps, M = 1000 bytes
  - Latency = 10ms + (1024 x 8)/(56 x 1024) sec = 153ms!
- Cross-country with T3 (45Mbps) line:
  - D = 50ms, R = 45Mbps, M = 1000 bytes
  - Latency = 50ms + (1024 x 8) / (45 x 1000000) sec = 50ms!

ALOHA

- Packet radio network in Hawaii, 1970s
- Wanted distributed allocation
  - no special channels or single point of failure
- Aloha protocol:
  - Just send when you have data!
  - There will be some collisions of course …
  - Throw away garbled frames at receiver (using CRC); sender will time out and retransmit
- Simple, decentralized and works well for low load

Carrier Sense Multiple Access

- We can do better by listening before we send (CSMA)
  - good defense against collisions only if “a” is small

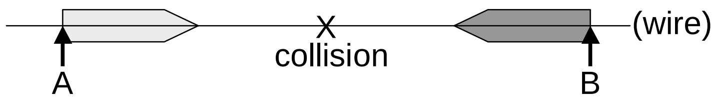

- “a” parameter: number of packets that fit on the wire
  - a = bandwidth * delay / packet size
  - a << 1 for LANs; a >> 1 for satellite

CSMA with Collision Detection

- Even with CSMA there can still be collisions.

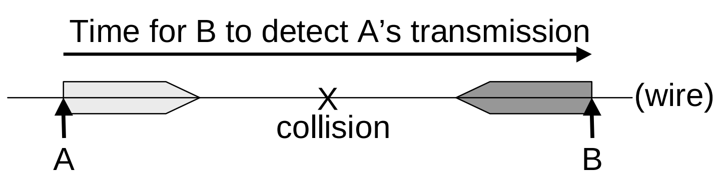

- What if we detect collisions and abort? CSMA/CD
  - Requires a minimum frame size (“acquiring the medium”), 2 x max propagation delay
  - B must continue sending (“jam”) until A detects collision

What if the Channel is Busy?

- 1-persistent CSMA
  - Wait until idle then go for it; what happens?
- Non-persistent CSMA
  - Wait a random time and try again
- p-persistent CSMA
  - Wait until idle, then in each time slot, choose to send with prob p
  - What if p is small? What if p is large?

CSMA/CD with Binary Exponential Backoff

- On collision: jam and exponential backoff
  - Jamming: send bit sequence to ensure collision detection
- Backoff:
  - First collision: wait 0 or 1 frame times at random and retry
  - Second time: wait 0, 1, 2, or 3 frame times
  - Nth time (N<=10): wait 0, 1, …, 2^N-1 times
  - Max wait 1023 frames, give up after 16 attempts
  - Scheme balances average wait with load
    - – what about fairness?

Ethernet

- Ethernet: first practical LAN, Xerox 1980's (802.3)
  - 1-persistent CSMA/CD with binary exponential backoff
  - 10 Mbps over coaxial cable, passive taps
  - Manchester encoding, preamble, 32 bit CRC

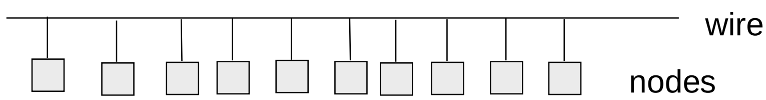

- Newer versions
  - Fast (100 Mbps), gigabit (1 Gbps)
  - Switched, point to point wires

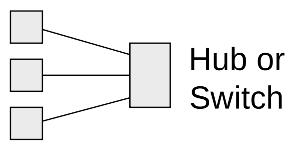

Ethernet Frames

Preamble (8)	Source (6)	Dest (6)	Len (2)	Payload (var)	Pad (var)	CRC (4)

- Min frame 64 bytes, max 1500 bytes
- CSMA/CD jam period is 48 bits
- Max length 2.5km, max between stations 500m (repeaters)
- Addresses unique per adaptor; globally assigned
- Broadcast media:
  - ARP, multicast, promiscuous mode monitoring

Ethernet Evaluation

- Fairness -- backoff favors latest arrival
  - max limit to delay
  - no history -- unfairness averages out
- Stable performance under increasing load
  - Much better than Aloha!
  - Works very well in practice
- Source of protocol inefficiency: collisions
  - What happens as bit rates increase?
    - – Need to shorten wires and increase frame size

Token Ring (802.5)

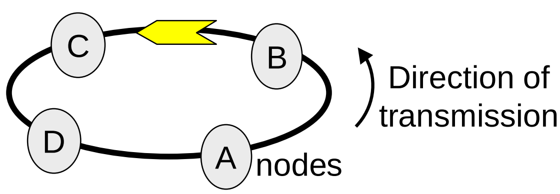

- “Permission to send” token rotates around loop
- When token arrives
  - if no data to send, forward the token
  - if data to send, remove token and inject packet
  - remove packet and re-inject token at **sender**
  - Maximum token holding time (THT) bounds access time

FDDI
Fiber Distributed Data Interface

- Roughly a large, fast token ring
  - 100 Mbps and 200km vs 4/16 Mbps LAN
  - Dual counter-rotating rings for better reliability
  - Complex token holding policies for real-time traffic
- Token ring advantages
  - No contention, bounded access delay
  - Support fair, reserved, priority access
- Disadvantages
  - Complexity, reliability, scalability

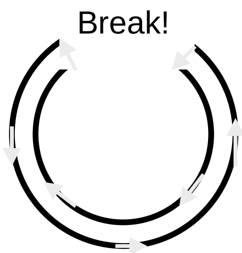

Why Did Ethernet Win?

- Reliablity
  - Token ring failure mode -- network unusable
  - Ethernet failure mode -- node detached
- Cost
  - Passive tap cheaper to build than active forwarder
  - Volume => lower cost => volume => lower cost …
- Scalability
  - Repeater: copy all packets across two segments
  - Bridge: selectively repeat packets across two segs
  - Switch: bridge k segments; Hub: repeater for k segs

Wireless Communication

Wireless is more complicated than wired …

- Cannot detect collisions
  - Transmitter swamps co-located receiver
- Different transmitters have different coverage areas
  - Asymmetries lead to hidden/exposed terminal problems

CSMA with Collision Avoidance

- Since we can't detect collisions, we avoid them
  - CSMA/CA as opposed to CSMA/CD
- If medium busy, choose random backoff interval
  - Wait for that many idle timeslots to pass before sending
- If a collision is inferred, retransmit with binary exponential backoff (like Ethernet)
  - Use CRC and ACK from receiver to infer "no collision"
  - Again, exponential backoff helps us adapt "p" as needed

Hidden and Exposed Terminals

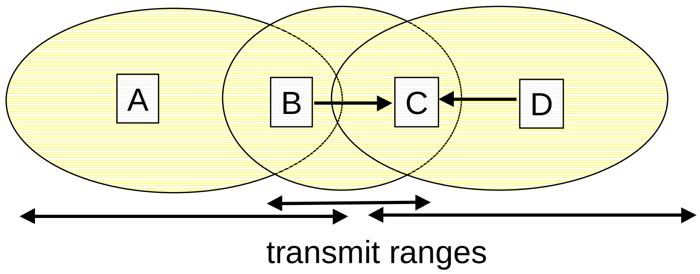

- Hidden terminals: B and D can both send to C but can't hear each other
- Exposed terminals: B, C can hear each other but can safely send to A, D

RTS / CTS Protocols (MACA)

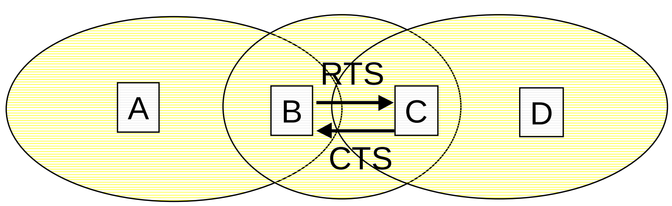

- B asks C: Request To Send (RTS)
- A hears RTS and defers to allow the CTS
- C replies to B with Clear To Send (CTS)
- D hears CTS and defers to allow the data
- B sends to C

802.11 Wireless LANs

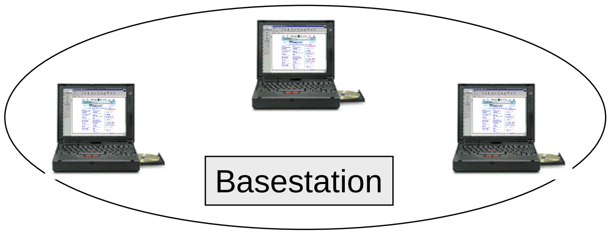

- Emerging standard: wireless plus wired system or pure wireless (ad hoc)
- Avoids collisions (CSMA/CA (p-persistence), RTS/CTS
- Built on new links (spread spectrum, or diffuse infrared)