EECS 122, Lecture 7

Kevin Fall

kfall@cs.berkeley.edu

Link Layer Networks

- Achieve station-to-station connectivity
- May be point-to-point or multi-access
  - point-to-point may not require addresses
  - multi-access requires addresses and sharing
- Error detection
- Addressing

Sharing

- How to share a broadcast media (e.g. wire, air)?
  - what to do with two simultaneous speakers
  - approaches: centralized and distributed
- Centralized (polling)
- Decentralized (speaking when media idle)

Comparing Approaches

- Centralized Approach
  - Polling requires speaker to await moderator, even if others idle
  - Problems if moderator's connection fails
- Decentralized Approach
  - no moderator wait, but subject to collisions
  - collisions could continue forever
- The *Multiple-Access* Problem

Contexts for the MA Problem:

- Wired LANs (e.g. Ethernet, FDDI)
- Wireless LANs (RF or Light)
- Packet Radio (e.g. Ricochet)
- Cellular Telephones
- Satellite Communications

Data Model

- Packet Mode vs Circuit Mode
  - Smooth, continuous traffic suggests circuit mode access (telephony)
  - Variable-demand (bursty) sources suggest packets
- Sometimes packet mode used for circuit establishment

Wireless Link Constraints

- Spectrum scarcity
  - ISM bands 902-928Mhz & 2.4-2.48Ghz in US
- Air link properties
  - fading (signal attenuation)
  - multipath interference
  - hidden-terminal problem
  - near-far problem

Other Constraints

- The "a" parameter:
  - a = D / T
  - D = max prop delay between 2 stations
  - T = time to xmit average packet
- How much can be placed on wire before farthest station receives first bit
- small *a* (.01, LANs); big *a* (100, Sats)

Design Issues

- *a* impacts what happens during simultaneous transmission:
  - *a* small -> early collision detection
  - *a* large -> late detection, want to avoid
- Performance issues
  - goodput, mean delay, stability, fairness, cost

Base Technologies

- Simple Techniques (FDMA, TDMA)
- Code Division Multiple Access (CDMA)
  - uses freq. Hopping or direct sequence *spread spectrum*
  - *SS* -> multi-bit codewords across multiple frequencies; users get orthogonal words for isolation

CDMA Benefits/Problems

- Benefits of CDMA
  - security, noise/jamming immunity
  - no time or freq sync required
  - no hard limit on capacity
  - inactive senders improve others S/N ratio
- Problems
  - complexity, power control, large frequency allocation

CSMA Type Networks

- CSMA -- Carrier Sense Multiple Access
  - detect when medium is busy
  - Persistent (send immediately)
  - Non-persistent (send some time soon)
- Approach to collisions
  - *p*-Persistence
  - detection with backoff

p-Persistent CMSA

- p = Prob(send|idle)
- E(# stations xmit after idle) = *np* [*n*: # total stations ready to send]
- If *np* > 1, likely secondary collision, so want *p* < 1/*n*
- *n* increases with system load, so want smaller *p* with high load
- smaller *p* affects message delay

Detection and Backoff

- Determine if frame transmitted successfully, if not, wait
- Detection via ACKs or *collision detect ion*
- Wait using *exponential backoff*
  - wait random on interval [0..2*(max prop)]
  - double 2 on each successive collision
  - even 1-persistent becomes stable; also avoids need to pick optimal *p*

Ethernet

- Most popular form of IEEE 802.3
- Variant of 1-persistent CSMA/CD with exponential backoff on wired LAN
- "Classical" Ethernet is 10Mb/s over 50-ohm Coax wiring
- Newer standards cover UTP wiring, 100Mb/s operation, etc

Names for Ethernet

- Names of form

  [rate][modulation][media or distance]
- Examples:
  - 10Base2 (10Mb/s, baseband, small coax)
  - 10Base-T (10Mb/s, baseband, twisted pair)
  - 100Base-TX (100Mb/s, baseband, 2 pair)
  - 100Base-FX (100Mb/s, baseband, fiber)

Ethernet Properties

- Will discuss "classical" Ethernet primarily
- Single segments up to 500m; with up to 4 repeaters gives 2500m max length
- Baseband signals broadcast, Manchester encoding, 32-bit CRC for error detection
- Max 100 stations/segment, 1024 stations/Ethernet

Detecting Collisions

- CD circuit operates by looking for voltage exceeding a transmitted voltage
- Want to ensure that a station does not complete transmission prior to 1st bit arriving at farthest-away station
- Time to CD can thus take up to 2x{max prop. delay} (but *a* is small!)

Minimum Frame Size

- Speed of light is about 3x10^8 m/s in vacuum and about 2x10^8 in copper
- So, max Ethernet signal prop time is about 12.5 µsec, or 25µsec RTT
- With repeaters, etc. 802.3 requires 51usec, corresponding to 512 bit-times
- Thus, minimum frame size is 512 bits (64 bytes); also called *slot time*

Maximum Frame Size

- 1500 byte limitation on maximum frame transmission size
- Later we will call this the *MTU*
- limits maximum buffers at receiver
- allows for other stations to send
  - also requires 96 bit Inter-Packet-Gap (IPG)

Transmitter

- When ready & line idle, await IPG (96 bit times) and send while listening (CD)
- If CD true, send max 48-bit jamming sequence and do exponential backoff
- Jamming sequence used to inform all stations that a collision has occurred

Exponential Backoff

- For retransmission N (1<=N<=10)
  - choose k at random on U(0..2^N-1)
  - wait k * (51.2µsec) to retransmit
  - send on idle; repeat on another collision
  - for (11<=N<=15), use U(0..1023)
  - if N = 16, drop frame
- Longer wait implies lower priority (strategy is not 'fair')

Capture Effect

- Given two stations A & B, unfair strategy can cause A to continue to 'win'
- Assume A & B always ready to send:
  - if busy, both wait, send and collide
  - suppose A wins, B backs off
  - next time, B's chances of winning are halved

Frame Structure

- 7 byte preamble: alternating 1/0 combination producing 10Mhz square wave for 5.6 µsec, used for recv sync
- 1 byte SOF (start of frame) 10101011
- 6 byte dest addr, 6 byte src addr, 2 byte type/length overloaded field
- variable sized data portion followed by 4-byte CRC-32
- sends low-order bit first for 802.3

Ethernet Frame Encapsulation

- Payload contains data of higher layer

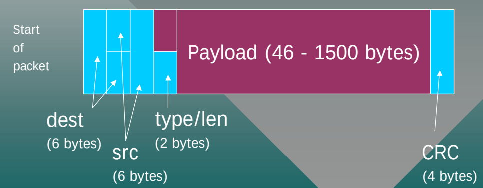

Ethernet Addressing

- 48 bit Ethernet/MAC/Hardware Addresses
- Prefix assigned per-vendor by IEEE
- Unique per-adapter, burned in ID PROM
- Multicast & Broadcast (all 1’s) addresses
- Many adapters support *promiscuous* mode

Multicast Addressing

- Each vendor assignment supports 2^24 individual and group (multicast) addresses
- Each adapter supports multiple group “subscriptions”
  - usually implemented as hash table
  - thus, software may have to filter at higher layer

802.3/Ethernet Type/Len

- 3rd field is 16-bits… overloaded
- Type field (Ethernet)
  - indicates type of data contained in payload
  - issue: what is the length?
- Length field (802.3)
  - type info follows frame header

Field Ambiguity

- So, is it the type or length?
  - “Ethernet”: types have values above 2048 (RFC894 for IP)
  - 802.3: length (RFC1042 for IP)
- If length, next headers are LLC & SNAP (for IP)
  - LLC (3 bytes): DSAP, SSAP, CTL
  - SNAP (5 bytes): org code, type (above)

IEEE 802.3u 100 Mb/s Ethernet

- “Fast Ethernet” (1995) adds:
  - 10x speed increase (100m max cable length retains min 64 byte frames)
  - replace Manchester with 4B/5B (from FDDI)
  - full-duplex operation using switches
  - speed & duplex auto-negotiation

IEEE 802.3{z,ab} 1000 Mb/s Ethernet

- "Gigabit Ethernet" (1998,9) adds:
  - 100x speed increase
  - carrier extension (invisible padding…)
  - packet bursting

Other LAN Technologies

- Ring networks generally more complex
  - IBM 4/16 Mb/s token ring
  - FDDI
- Connection-oriented
  - ATM, HIPPI

Perspective

- Ethernet is wildly successful, partly due to low cost (compare with FDDI or Token Ring--- see text book)
- Some issues:
  - nondeterministic service
  - no priorities
  - min frame size may be large