ECE 435 – Network Engineering
Lecture 6

Vince Weaver

http://web.eece.maine.edu/~vweaver

vincent.weaver@maine.edu

19 September 2016

Announcements

- Don't forget HW#2

Example Link Protocols

- Obsolete/Fading
  - Token Ring
  - HIPPI
  - FDDI – fiber distributed data interface
  - Fibre Channel
  - ATM
  - ISDN
  - X.25
- Current

- Ethernet (802.3)
- PPP (? fading)
- WLAN (802.11)
- Bluetooth (802.15)
- FiberChannel (?)
- DSL
- LTE/WiMAX (802.16)

HDLC

- SDLC protocol used on IBM mainframes (Synchronous Data Link Control)
- Standardized as high-level data link control
- normal response mode
- asynchronous response mode
- asynchronous balance mode

- 01111110 and bit stuffing
  flag(8) address(8) control(8) data fcs(16) flag(8)

- Flag is 0111 1110. On idle channel this is sent continuously.

- Control used for sequence numbers, acks, etc.

- Address is address if not point-to-point

- Data may be any length but checksum less effective the longer it is.

- fcs 16-bit CRC-CCITT complex
- sliding window, 3-bit sequence. Up to 7 unacknowledged frames can be outstanding
- Control has seq field, next field (piggyback acknowledge). Instead of sending last frame received correctly sends next frame expected
- P/F poll final used when polling a group of terminals
- info, supervisory, and unnumbered frames

- Supervisory:
  - Type 0: acknowledge frame (RECEIVE READY)
  - Type 1: negative acknowledge frame (REJECT)
  - Type 2: RECEIVE NOT READY
  - Type 3: SELECTIVE REJECT
- Unnumbered, control or data when unreliable

PPP

- point-point protocol, RFC1661/1662/1663
- link-control-protocol needs to authenticate before it can transfer packets
- a network control protocol
- state diagram. dead, establish, authenticate terminate, network
- HDLC-like format, but byte oriented (not bit)

- no medium access control (point-to-point)
- no flow control, left to upper layer
- PPPoE (over Ethernet). Encapsulate PPP over Ethernet

Ethernet History

- Proposed by Bob Metcalfe in 1973 (went on to found 3Com)
- Metcalfe, Boggs, Thacker, and Lampson listed on patent
- Inspired by ALOHAnet, a wireless network in Hawaii, allow users on various islands to connect to server on Oahu
- Various competing local networks, Ethernet won in the end

- Token Ring, dates back to 1970s
  Standardized by IBM, 1984, IEEE802.5
  4Mbps, eventually shielded twisted pair, eventually 16Mbps, 100Mbps and 1Gbps
  3-byte frame passed around gives permission to transmit
  More complex, no crossover cable (direct connect two machines), Supports multiple identical MAC addresses
  Deterministic time to get to transmit
  Frames can have different access priorities
  Empty token passed around. If data to transmit, put in.

Then passes around until it gets to receiver, removed, and back to passing empty token. When gets back to originator it knows it has been received.

- Token Bus, GM, IEEE802.4 (withdrawn)
  like ring, but virtual ring. Needed to know neighbors to pass token. Guaranteed worst case transmit time.
- FDDI, ATM, DQDB.
- Why did it win? Simpler and thus cheaper.
- Why simpler? No priority mechanism, no QoS, no central control
- Could use cheaper twisted pair cable

- Token ring cards generally a lot more expensive than Ethernet

- Current Ethernet is very different than original

- IEEE 802.3

- Big deal was its MAC: "Carrier sense multiple access with collision detection" (CSMA/CD)

Ethernet Timeline

- Standardized in 1981
- Timeline: 1972 – experimental 3Mbps
  1981 – DIX (DEC/Intel/Xerox) ver 1 (10Mbps)
  1982 – DIX ver 2
  1983 – IEEE 802.3/10BASE5 "Thick Ethernet", up to 500m, vampire taps, often yellow or orange (standard suggests yellow). Looks like garden hose. AUI connector, drill into cable, at 2.5m intervals (to avoid reflections), terminated on each end. One bad connection could ruin

for all

1985 – 10BASE2 – “thin net”, thinner connections, BNC connectors, T connectors (185m, rounded up to 200m), 50 ohm terminator, grounding loops, one and only one must be tied to ground. How to detect network problem? Send pulse, look for echo

1990 – 10BASE-T – twisted pair (Cat3) , needed hub

1993 – 10BASE-F – fiber

1995 – 100BASE-T

100BASE-TX 4B5B MLT-3 cat5 two twisted pairs

1997 – Full-duplex

1998/1999 – 1000BASE-TX PAM-5, four twisted pairs, can transfer in both directions on one pair using DSP/echo cancelation
2006 – 10GBASE-T
2010 – 40G and 100G
2017 – 400GB

- Power over Ethernet?
- Naming: Speed/BROAD,BASE,PASS/PHY. Almost all is baseband. PHY originally was distance could travel (in 100m) but now medium type.

- Original Ethernet was a bus on coaxial cable. collisions are common
- 10BASE-T dedicated media between devices (need hub or switch)
- Full-duplex
- Use of single cable – shared, anyone can see any traffic, also bandwidth is shared
- Autonegotiation – how figure out line speed and duplex

Manchester Encoding

- Does not use 0V for 0 and 5V for 1. Why? Idle is 0, so how can you tell how many zeros at beginning of signal?
- Could use +1V/-1V, but still would need way to sync signal on long runs of 0 or 1
- Manchester encoding : 1 is high to low transition. 0 is low to high transition. Always a transition in the middle of an interval.
- Disadvantage, need twice as much bandwidth

- Differential Manchester: transition at start of internal means 0, lack of one means 1. Still transition in the middle. More complex but better noise handling

- Ethernet uses Manchester, Token Ring uses differential Manchester

- Ethernet high 0.85V and low -0.85V

Ethernet frame layout

- Packet format:
  Preamble = 7 bytes
  SFD (start of frame delimit) = 1 byte
  DA (destination address) = 6 bytes
  SA (source address) = 6 bytes
  T/L (type length) = 2 bytes
  Data = 46-1500 bytes
  Frame Check Sequence = 4 bytes

- Preamble is fixed 1010...1010 in transmission order (little-endian)
  On original Ethernet this was 10MHz 6.4us pulse used to synch clocks The PHY might do other things (100BASE-X adds 4B/5B stuff)
- SFD - indicates the start of the frame with value 10101011 in transmission order.
- DA = 48 bit destination MAC address
  first bit unicast or multicast
  with first bit high, multicast, sent to many in a group, all 1s is broadcast, send to everyone

1-3 OUI (organization unique identifier)
Globally Unique

- SA = 48 bit source MAC address

- T/L: Originally type field. 802.3 makes it length of *data* field (not length of frame). Later in 1997 802.3 approved as type too, so dual meaning. How tell difference? Since cannot be longer than 1500, any value bigger than 0x600 (1536) is type.
- Data – data from 46 to 1500 bytes
Why limit 1500B? because RAM was expensive in 1978.

If smaller than 46 bytes padded. Makes sure checksum works. Also if too short, could be done transmitting before a collision can be detected (light travel to furthest node and back)

- FCS – a 32-bit CRC code. If incorrect FCS, silently drops. How can we do this? Up to upper protocol (say TCP/IP) to figure out if need to resend. Makes things simple. No need to wait for ACKs.
- Frame size is variable. Often first two fields are excluded and said that Ethernet packets are between 64 and 1518 bytes long

Ethernet MAC

- CSMA/CD
- First senses cable (how?)
- If busy, waits
- Sends. If collision, jams the cable aborts transmission, waits random back off time before retrying.
- on newer full-duplex links no need for carrier sense and collision detection not needed

Ethernet Transmission

- MAC puts data into frame
- In half-duplex CSMDA/CD senses carrier. Waits until channel clear
- Wait for an inter-frame-gap (IFG) 96 bit times. Allows time for receiver to finish processing
- Star transmitting frame
- In half-duplex, transmitter should check for collision.

Co-ax, higher voltage than normal
For twisted pair, noticing signal on the receive while transmitting

- If no collision, then done
- If collision detected, a *jam* signal is sent for 32-bits to ensure everyone knows. Pattern is unspecified (can continue w data, or send alternating 1s and 0s)
- Abort the transmission
- Try 16 times. If can't, give up

- Exponential backoff. Randomly choose time from 0 to $2^k - 1$ where k is number of tries (capping at 10). Time slot is 512 bits for 10/100, 4096 for 1Gbs

- Wait the backoff time then retry

Ethernet Receiving

- Physical layer receives it
- Frame passed up (minus preamble and SFD)
- Receiving process records bits until signal done. Truncated to nearest byte.
- If too short (less than 512 bits) treated as collision
- If destination is not the receiver, drop it

- If frame too long, dropped and error recorded
- If incorrect FCS, dropped and error recorded
- If frame not an integer number of octets (?) dropped and error recorded
- If everything OK, de-capsulated and passed up

Ethernet Collisions

- In order to work properly, twice round-trip time needs to be less than time needed to transmit minimal (64-byte) frame, otherwise not possible to notice collision in time and frame loss

- This limits network size to collision domain

- Bits wasted is not bad, collision often caught in the preamble

Maximum Frame Rate

- 7+1 byte preamble 64-byte frame, IFG of 12 bytes between transmissions. equals 672 bits. In 100Mbps system 148,800 frames/second

Full Duplex MAC

- Early Ethernet was coaxial in a bus
- Twisted pair has replaced this, usually in a hub/or switch star topology
- 10BASE-T and 100BASE-TX pair for transmit or receive
- inefficient. Since point to point, why do you need arbitration?
- Full-duplex introduced in 1997. Must be able to

transmit/receive w/o interference, and be point to point.

- Full duplex effectively doubles how much bandwidth between. Also it lifts the distance limit imposed by collision detection