ECE 435 – Network Engineering
Lecture 28

Vince Weaver

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

vincent.weaver@maine.edu

8 April 2026

Announcements

- HW#9 due Friday
- Question on modern Token Ring, you'd think there'd be homebrew projects around it but couldn't really find any. Most recent hardware is regular PCI
- Aside, FDDI, a token-ring like protocol was proposed at the time 100Mb ethernet came out, but lost out to fast ethernet

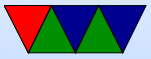

"Classic" Ethernet Overview

- Not really used anymore, but a classic example of what a relatively easy-to-understand link-level interface is like

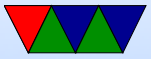

Ethernet MAC

- CSMA/CD "Carrier sense multiple access with collision detection"
- First senses cable (how? – see later)
- If busy, waits
- Sends. If collision, jams the cable aborts transmission, waits random back off time before retrying.
- Exponential backoff. Randomly choose time from 0 to

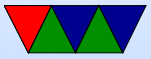

$2^k - 1$ where k is number of tries (capping at 10). Time slot is 512 bits for 10/100, 4096 for 1Gbs

- on newer full-duplex links no need for carrier sense and collision detection not needed

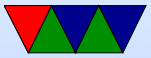

Ethernet Collisions

- 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

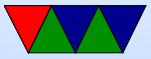

Physical Layer Encoding Background: NRZ/NRZI

For initial comparison, these are the simplest encodings.

NRZ (non-return to zero) is just 0 low, 1 high

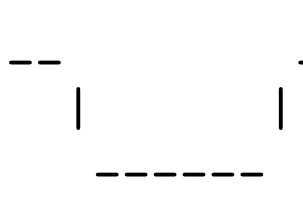

NRZI (NRZ Inverted) (1 flip, 0 same)

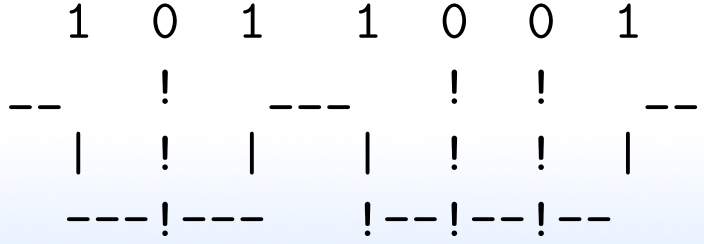

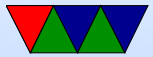

! ! !

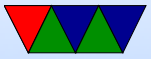

Physical Layer – Voltages

- Could just use 0V for 0 and 5V for 1
- What if you want to be able to detect idle/unused?
- Have Idle be 0V, and +1V/-1V for 1/0
- This also avoids having a DC bias to your signal

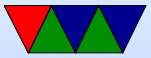

Manchester Encoding (Ethernet)

- 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

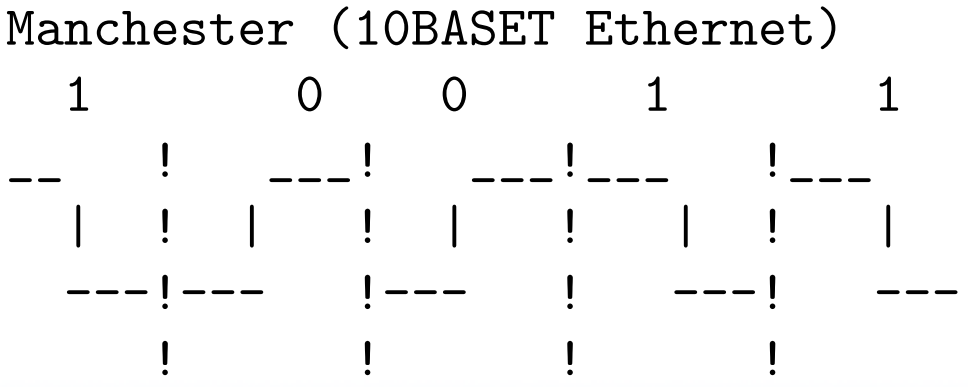

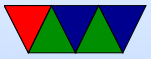

Differential Manchester (Token Ring)

- transition at start of interval means 0
- lack of transition means 1
- Still transition in the middle
- More complex but better noise handling

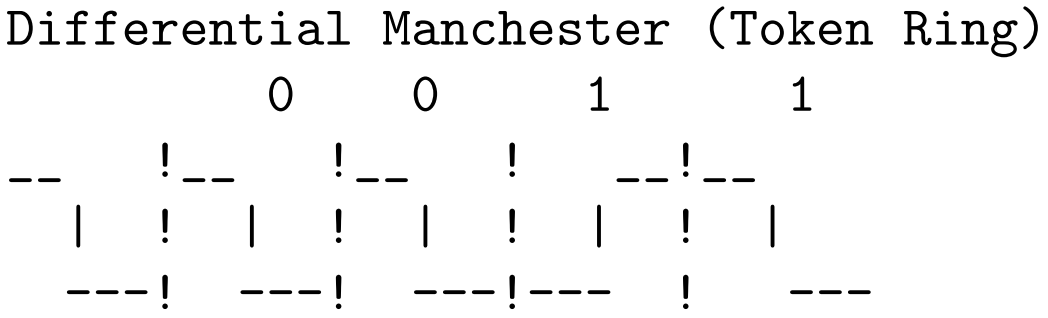

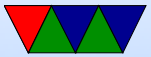

Why worrying about transitions?

- There's no clock on these signals
- To keep things synchronized have to watch for transitions to keep things together
- Often rules about how many 0s in a row for same reason

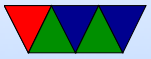

Ethernet on the Wire

- Manchester Encoding
- High 0.85V and low -0.85V
- Makes it hard to bit-bang ethernet with GPIOs

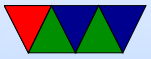

Ethernet Frame layout

This is Ethernet II/DIX standard, the most common

Preamble	SFD	DA	SA	T/L	Data	FCS
7 bytes	1	6	6	2	46-1500	4

Frame size is variable. Often first two fields are excluded and said that Ethernet frames are between 64 and 1518 bytes long

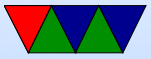

Ethernet Frame Preamble

- Fixed 1010...1010 in transmission order (LSB, least significant bit first)
- On original Ethernet this was 10MHz 6.4us pulse used to synch clocks
- PHY might do other things (100BASE-X uses 4B/5B stuff, so different pattern)

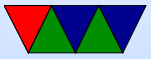

Ethernet Start Frame Delimiter (SFD)

- SFD - indicates the start of the frame
- value 10101011 in transmission order
- Original Ethernet declared 8 bytes of same pattern, but on modern first 7 bytes might be different

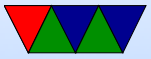

Ethernet MAC addresses

- DA = 48 bit destination MAC address
- SA = 48 bit source MAC address
  - First 3 bytes the OUI (organization unique identifier)
  - Next 3 bytes supposed to be a unique ID
- Ethernet packets put on the wire least-significant bit first (as if shifted right out of a shift register)
- Multicast if the "first" bit (meaning 0x1, not 0x80) is set in the first octet (e.g. 01-80-C2-00-00-00)
- Broadcast if all bits set ff:ff:ff:ff:ff:ff

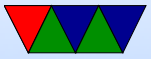

Ethernet Type/Length Field

- Originally type field
- 802.3 makes it length of *data* (not length of frame)
- In 1997 802.3 approved as type too, so dual meaning
- How tell difference?
  - Max len 1500, value bigger than 0x600 (1536) is type
  - 0x0800 = IPv4, 0x86dd = IPv6, 0x0806 = ARP
  - How tell length if type? Detect end of signal or inter-frame gap, or valid checksum (this is most common)
  - How tell type if length? Will have 802.2 header?

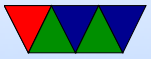

Ethernet Frame – Data

- 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)

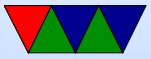

Interesting Article on Frame Sizes

- `https://www.potaroo.net/ispcol/2024-10/packet-sizes.html`
- Small frames have high overhead
- Big frames much more likely to have errors

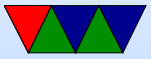

Ethernet Frame Check Sequence (FCS)

- FCS – a 32-bit CRC code.
- To avoid problems with all-zeros, bits are complemented when running. These means a 'good' result isn't 0 but 0xc704dd7b
- Somewhat complicated to calculate, a bit beyond this class
  - Hard to get a clear description of what it looks like w/o a lot of math
  - You can implement with linear feedback shift register

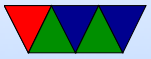

(shift register with xors)

- Possibly CRC calculated MSB first rather than LSB

- If incorrect FCS, silently drops

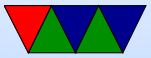

End of Frame / Inter Packet Gap

- At end of frame, drop carrier
- More modern Ethernet might signal this with a symbol
- Inter-packet gap, 96 bits (12 bytes) of idle before sending next frame
- Gives receiver time to handle frame before next starts

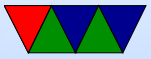

ARP – address resolution protocol

- On local network, how do we find MAC address if we know IP?
- Hard-code mapping `/etc/ethers`?
- Can it be automatically determined somehow?
- ARP – address resolution protocol (IPv4)
- ND – Neighborhood Discovery (IPv6)

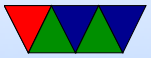

ARP (RFC826)

- Device first checks ARP cache to see if already knows
- Otherwise, broadcasts to ff:ff:ff:ff:ff:ff "who has this IP"
- Device reply with its IP and MAC (unicast)
- These are cached
- Timeout in case you reassign
- ARP announcement: can broadcast when your address changes so they can update (gratuitous ARP)
- Other optimizations(?)

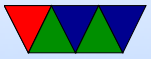

ARP Security

- Can you spoof ARP responses to get frames meant for a different device?

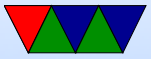

IPV6: Neighborhood Discovery Protocol

- Uses ICMPv6 message format
  - Router Solicitation (Type 133) – used to find router for local network
  - Router Advertisement (Type 134) – routers periodically also send their router info to whole local network
  - Neighbor Solicitation (Type 135) – can request MAC address from IPv6 address of neighbors
  - Neighbor Advertisement (Type 136) – response to a

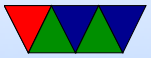

solicitation, or can just send it to everyone if something changed
  - Redirect (Type 137)
- To do request, must create two multicast addresses. Less overhead than ARP as in this case only a small number of hosts will share the multicast address
  - solicited-node multicast address least-significant 24 bits of the number looking up and appending to prefix ff02::1:ff00:0/104
  - solicited-node multicast MAC address least-significant 24 bits of the previous solicited-node multicast address

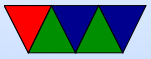

and appending to prefix 33:33:FF:xx:xx:xx

- Secure Neighbor Discovery Protocol (SEND) – uses certificates and stuff to avoid ARP spoofing

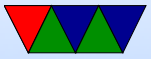

RARP/BOOTP

- Some cases need to do RARP (Reverse ARP) (RFC 903) have own MAC, find IP (netbooting is common reason)
- ARP packets not forwarded, so extension called BOOTP that allowed network booting.
- BOOTP automated by DHCP.
- IPv6 has IND (Inverse Neighborhood Discovery Protocol)

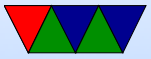

Classic Ethernet Transmission (Review)

- Break 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
- Start 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

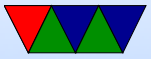

transmitting

- 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

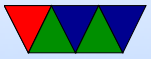

Classic Ethernet Receiving (Review)

- Physical layer receives it, recording 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

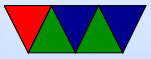

- Frame passed up (minus preamble, SFD, and often crc)

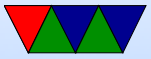

Classic Ethernet Receiving – Security

- Every ethernet card seem every frame that goes by on wire
- You are supposed to discard them once you see you aren't the destination MAC address
- Can you tell the card to keep and analyze them anyway?
- This was called "Promiscuous mode"
- Huge security issue, especially back in the days when the internet was plaintext and e-mail, passwords, chats, etc, in plain text

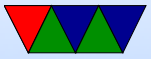

Maximum Frame Rate

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

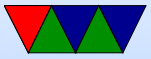

Ethernet Flow Control

- Flow control is optional
- In half duplex a receiver can transmit a "false carrier" of 1010..10 until it can take more.
- Congested receiver can also force a collision, causing a backoff and resend. Sometimes called force collision
- Above schemes called "back pressure"
- For full duplex can send a PAUSE frame that specifies how much time to wait.