Housekeeping

- Lab access
- HW turn-in
- Jin ?
- Class preparation for next time: look at the section on CRCs 2.4.3. Be prepared to explain how/why the shift register implements the CRC
- Skip Token Rings section
- New HW assignment next time – no programming

Point-to-Point Links

Outline

Basic Performance Metrics
Encoding
Framing
Error Detection
Sliding Window Algorithm

Basic Performance Metrics

- Bandwidth (throughput)
  - data transmitted per time unit
  - link versus end-to-end
  - notation
    - K**B** = 2^{10} bytes
    - M**b**ps = 10^6 bits per second
- Latency (delay)
  - time to send message from point A to point B
  - one-way versus round-trip time (RTT)
  - components

    Latency = Propagation + Transmit + Queue
    Propagation = Distance / *c*velocityFactor* (Table 2.8)
    Transmit = Size / Bandwidth

Bandwidth versus Latency

- Relative importance
  - 1-byte: 1ms vs 100ms dominates 1Mbps vs 100Mbps
  - 25MB: 1Mbps vs 100Mbps dominates 1ms vs 100ms
- Infinite bandwidth
  - RTT dominates
    - Throughput = TransferSize / TransferTime
    - TransferTime = RTT + 1/Bandwidth x TransferSize
  - 1-MB *file* to 1-Gbps link as 1-KB *packet* to 1-Mbps link

Delay x Bandwidth Product

- Amount of data “in flight” or “in the pipe”
- Example: 100ms x 45Mbps = 560KB
- Why is the delay-bandwidth product important?

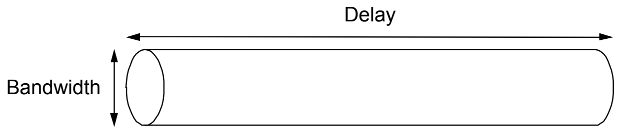

Things that go wrong

- Damaged packets – electrical/optical errors
- Dropped packets – congestion
- Duplicated packets
- Link and node failures
- Messages are delayed
- Messages are deliver out-of-order
- Third parties eavesdrop or inject traffic

Encoding

- Signals propagate over a physical medium
  - modulate electromagnetic waves
  - e.g., vary voltage
- Encode binary data onto signals
  - e.g., 0 as low signal and 1 as high signal
  - known as Non-Return to zero (NRZ)

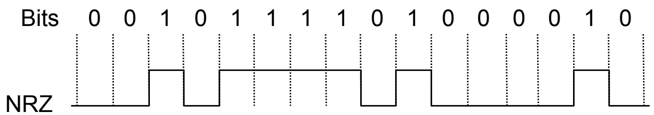

Problem: Consecutive 1s or 0s

- Low signal (0) may be interpreted as no signal
- High signal (1) leads to baseline wander
- Unable to recover clock

Alternative Encodings

- Non-return to Zero Inverted (NRZI)
  - make a transition from current signal to encode a one; stay at current signal to encode a zero
  - solves the problem of consecutive ones
- Manchester
  - transmit XOR of the NRZ encoded data and the clock
  - only 50% efficient.

Encodings (cont)

- 4B/5B
  - every 4 bits of data encoded in a 5-bit code
  - 5-bit codes selected to have no more than one leading 0 and no more than two trailing 0s
  - thus, never get more than three consecutive 0s
  - resulting 5-bit codes are transmitted using NRZI
  - achieves 80% efficiency

Encodings (cont)

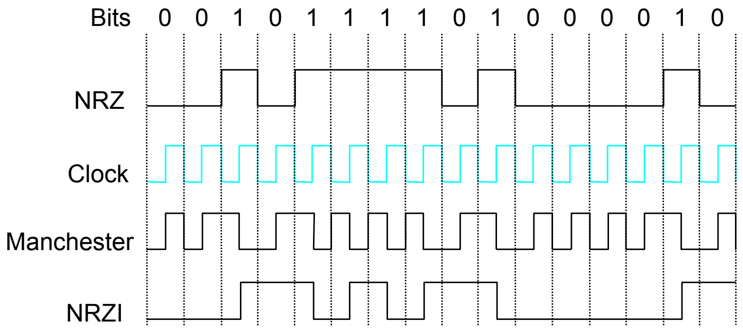

Framing

- Break sequence of bits into a frame
- Typically implemented by network adaptor

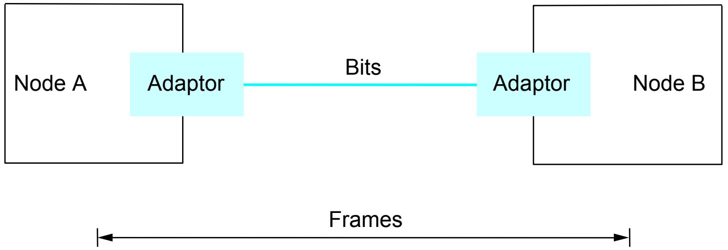

Approaches

- Sentinel-based
  - delineate frame with special pattern: 01111110
  - e.g., HDLC, SDLC, PPP

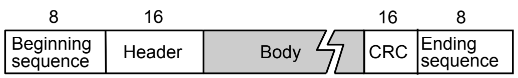

  - problem: special pattern appears in the payload
  - solution: *bit stuffing/byte stuffing*
    - sender: insert 0 after five consecutive 1s
    - receiver: delete 0 that follows five consecutive 1s
  - finite state machine implementation

Approaches (cont)

- Counter-based
  - include payload length in header
  - e.g., DDCMP (*Digital Data Communication Message Protocol)*

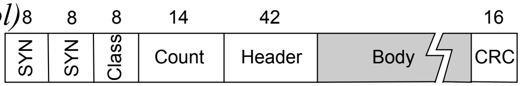

  - problem: count field corrupted
  - solution: catch when CRC fails

Approaches (cont)

- Clock-based
  - each frame is 125us long
  - e.g., SONET: Synchronous Optical Network
  - STS-*n* (STS-1 = 51.84 Mbps)

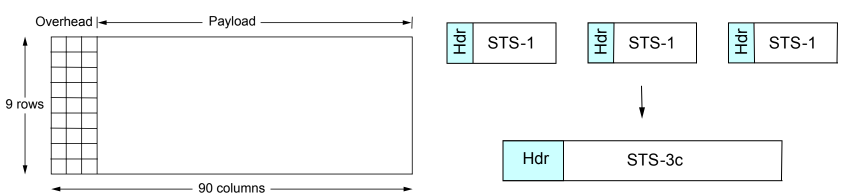

Ethernet Overview

- History
  - developed by Xerox PARC in mid-1970s
  - roots in Aloha packet-radio network
  - standardized by Xerox, DEC, and Intel in 1978
  - similar to IEEE 802.3 standard
- CSMA/CD
  - carrier sense
  - multiple access
  - collision detection
- Frame Format

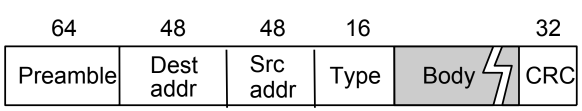

Ethernet (cont)

- Addresses
  - unique, 48-bit unicast address assigned to each adapter
  - example: **`8:0:e4:b1:2`**
  - broadcast: all **`1`**s
  - multicast: first bit is **`1`**
- Bandwidth: 10Mbps, 100Mbps, 1Gbps
- Length: 2500m (500m segments with 4 repeaters)
- Problem: Distributed algorithm that provides fair access

Transmit Algorithm

- If line is idle…
  - send immediately
  - upper bound message size of 1500 bytes
  - must wait 9.6us between back-to-back frames
- If line is busy…
  - wait until idle and transmit immediately
  - called *1-persistent* (special case of *p-persistent*)

Algorithm (cont)

- If collision…
  - jam for 32 bits, then stop transmitting frame
  - minimum frame is 64 bytes (header + 46 bytes of data)
  - delay and try again
    - 1st time: 0 or 51.2us
    - 2nd time: 0, 51.2, or 102.4us
    - 3rd time51.2, 102.4, or 153.6us
    - *nth* time: k x 51.2us, for randomly selected k=0..2^n - 1
    - give up after several tries (usually 16)
    - exponential backoff

Collisions

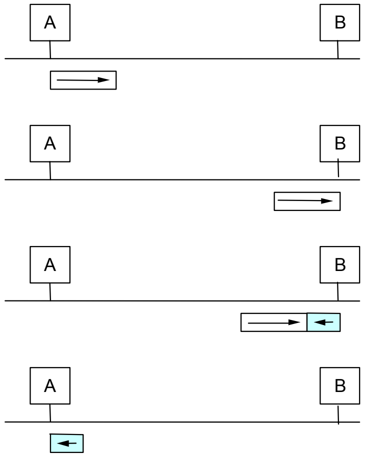

Review

- Physical layer and link layer
  - Performance – concepts extend to multi-hop networks
  - Encoding – primarily physical layer
  - Framing – varies – PPP it is a link layer issue, ethernet it is a physical layer issue