CSEP 561 – Bits and Links

David Wetherall
djw@cs.washington.edu

Topic

- How do we send a message across a wire or wireless link?

- The physical/link layers:
  1. Different kinds of media
  2. Fundamental limits
  3. Encoding bits
  4. Model of a link

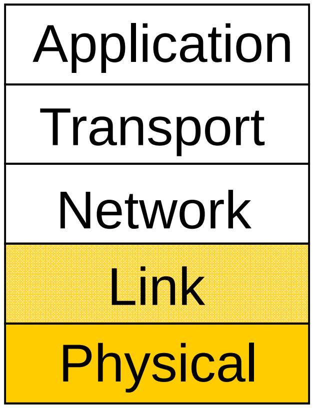

1. Wires

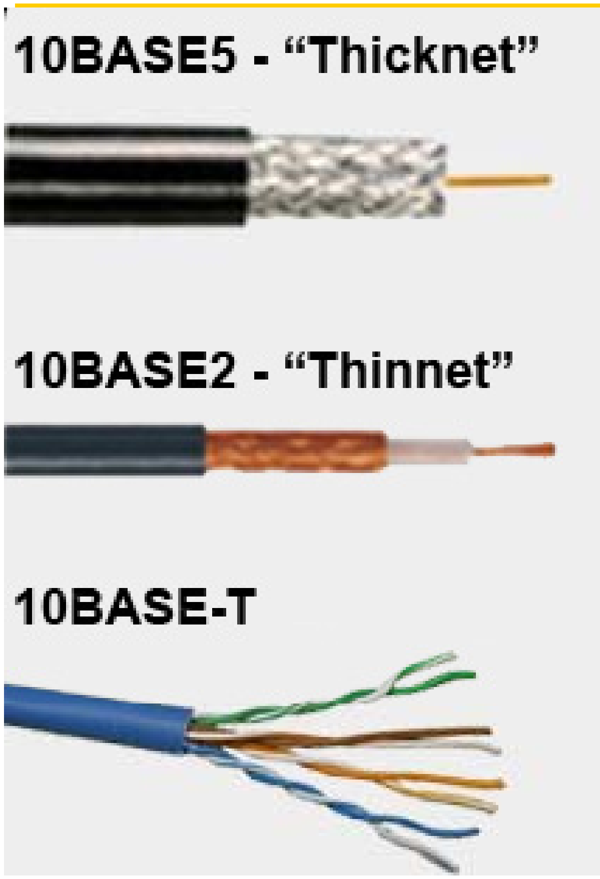

Now Cat 6, Cat 7 for GigE, four pairs

- Twisted pairs: twists reduce RF emission / crosstalk; also shielding can be added
- Coaxial cable: inner and outer ring conductor for superior noise immunity
- Many different specs/grades depending on application
- 100s of MHz for 100s of meters

Fiber Optic Cable

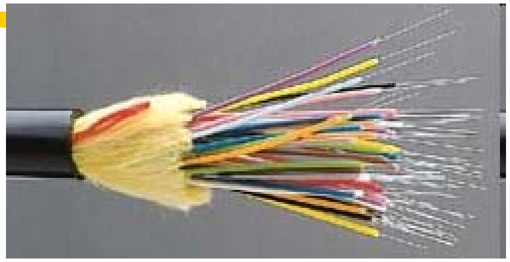

- Long, thin, pure strand of glass
  - light propagated with total internal reflection
  - enormous bandwidth available (terabits)

Multi-Mode

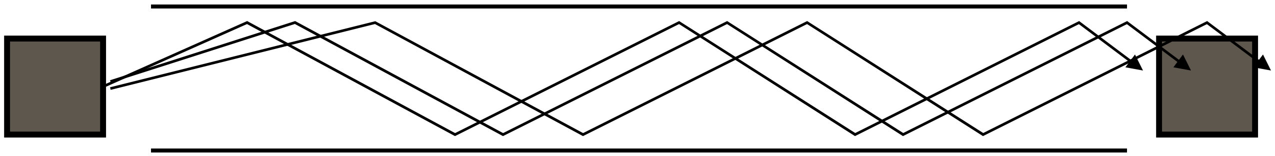

Light source (LED, laser)

Light detector (photodiode)

Single-Mode

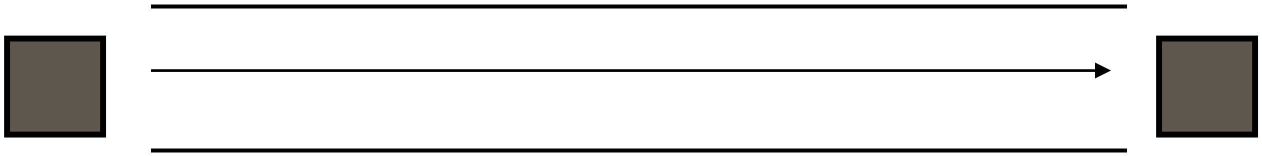

Light source (LED, laser)

Light detector (photodiode)

Wireless

- Different frequencies have different properties
- Signals subject to atmospheric/environmental effects

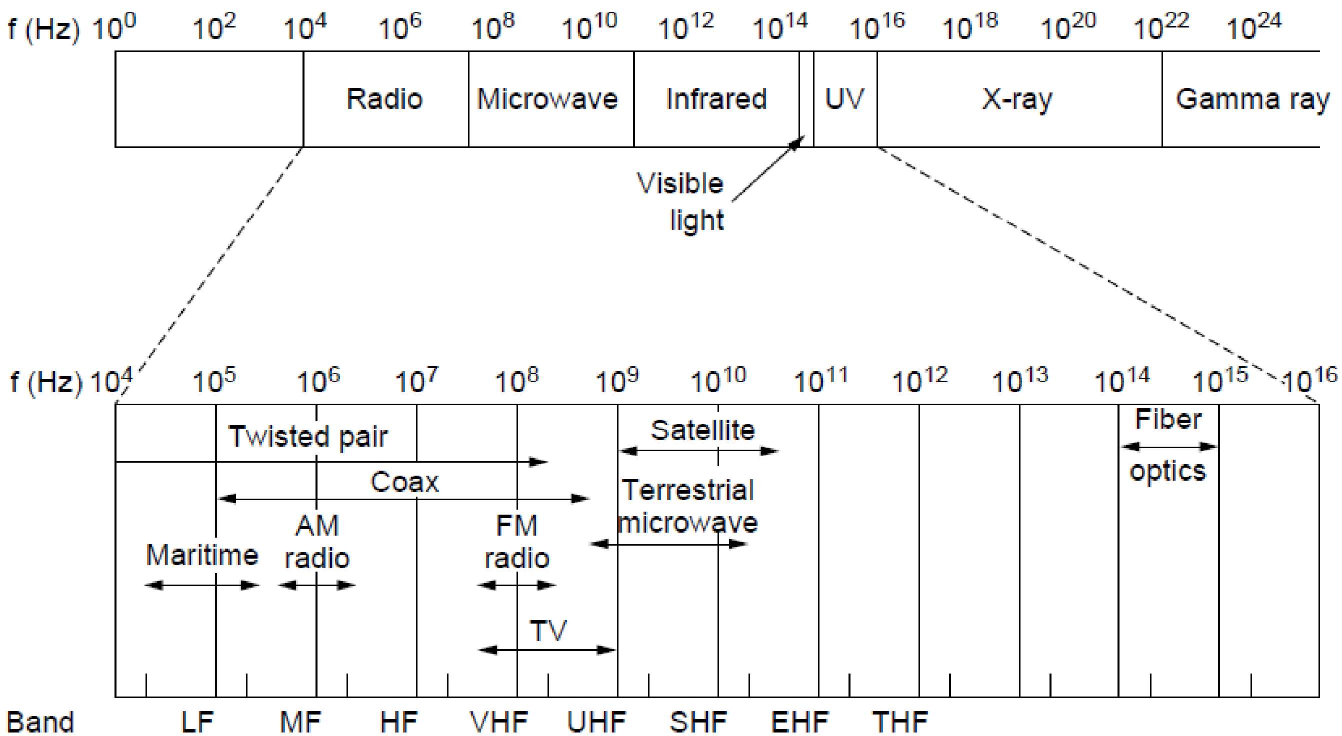

US spectrum allocations – regulation

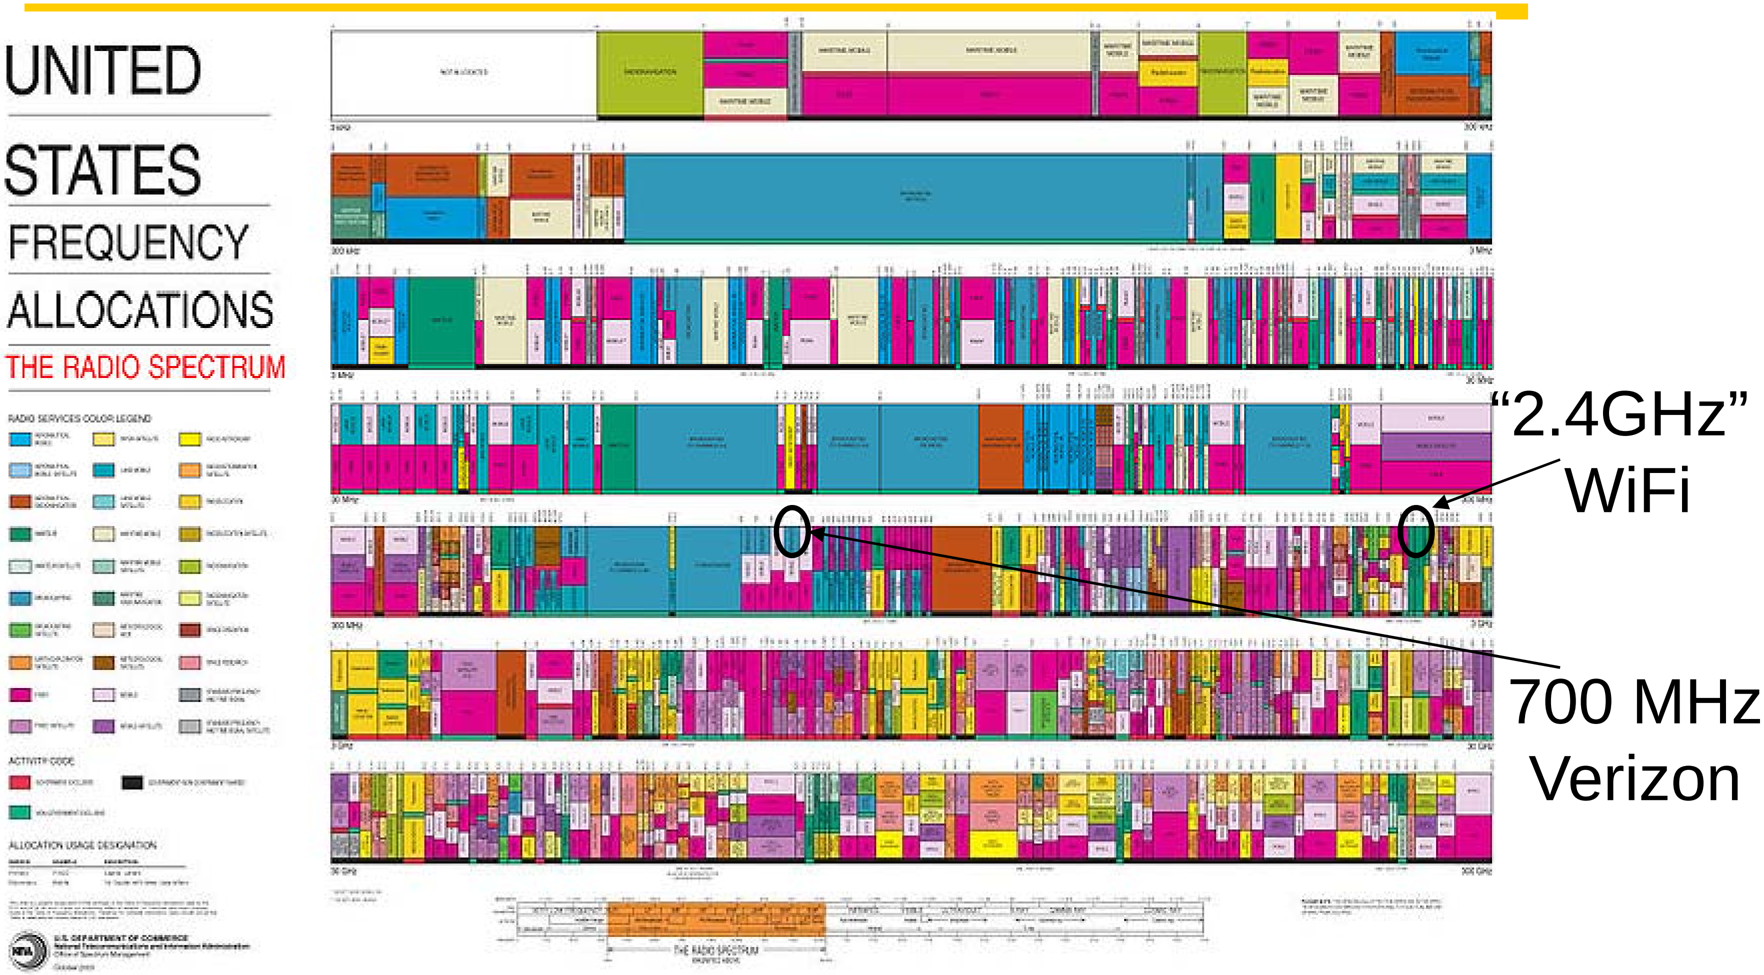

Model of a wire

- Frequencies beyond cutoff highly attenuated
  - Bandwidth = passband (Hz)

- Signal also subject to:
  - Attenuation (repeaters)
  - Distortion (frequency and delay)
  - Noise (thermal, crosstalk, impulse)

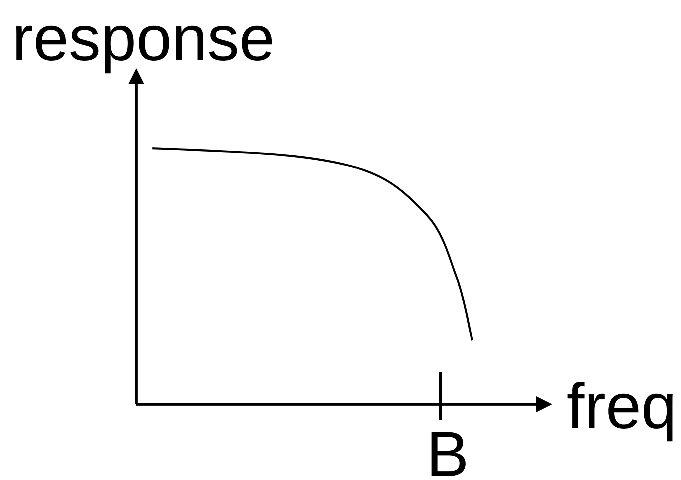

EE: Bandwidth = width of frequency passband, measured in Hz
CS: Bandwidth = information carrying capacity, measured in bits/sec

Attenuation of optic fiber

- Enormous bandwidth in each window

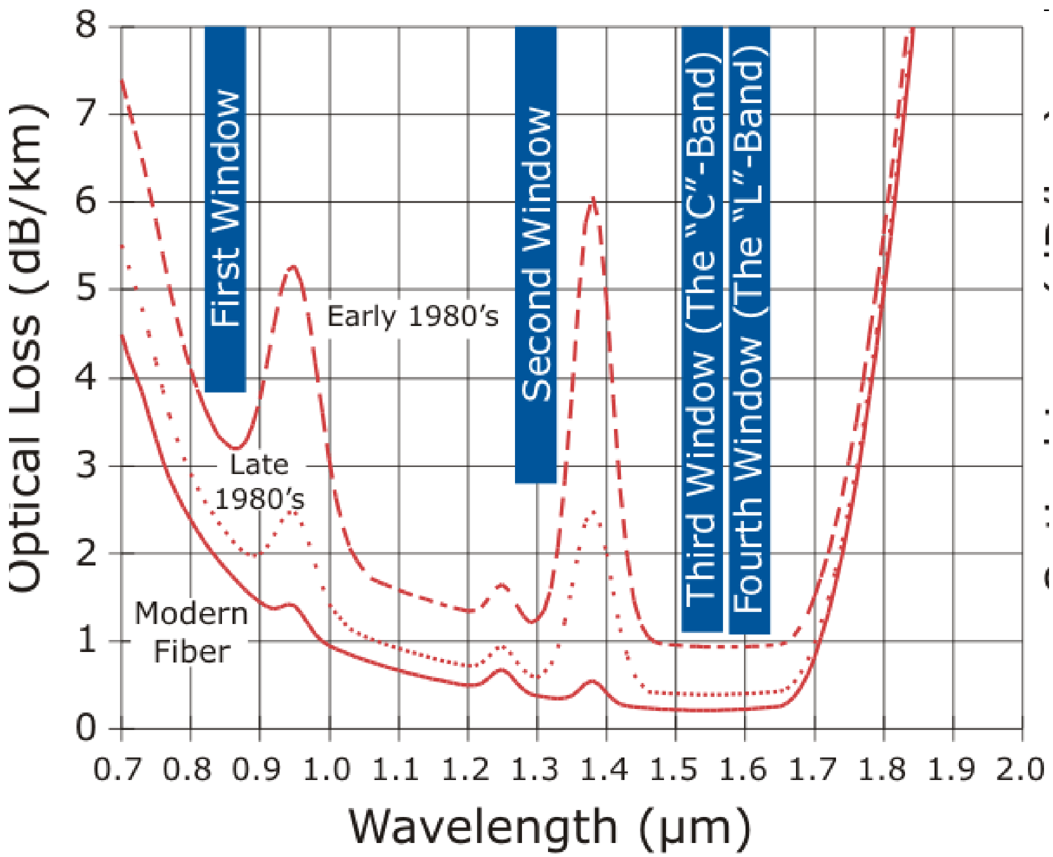

2. Effect of limited bandwidth

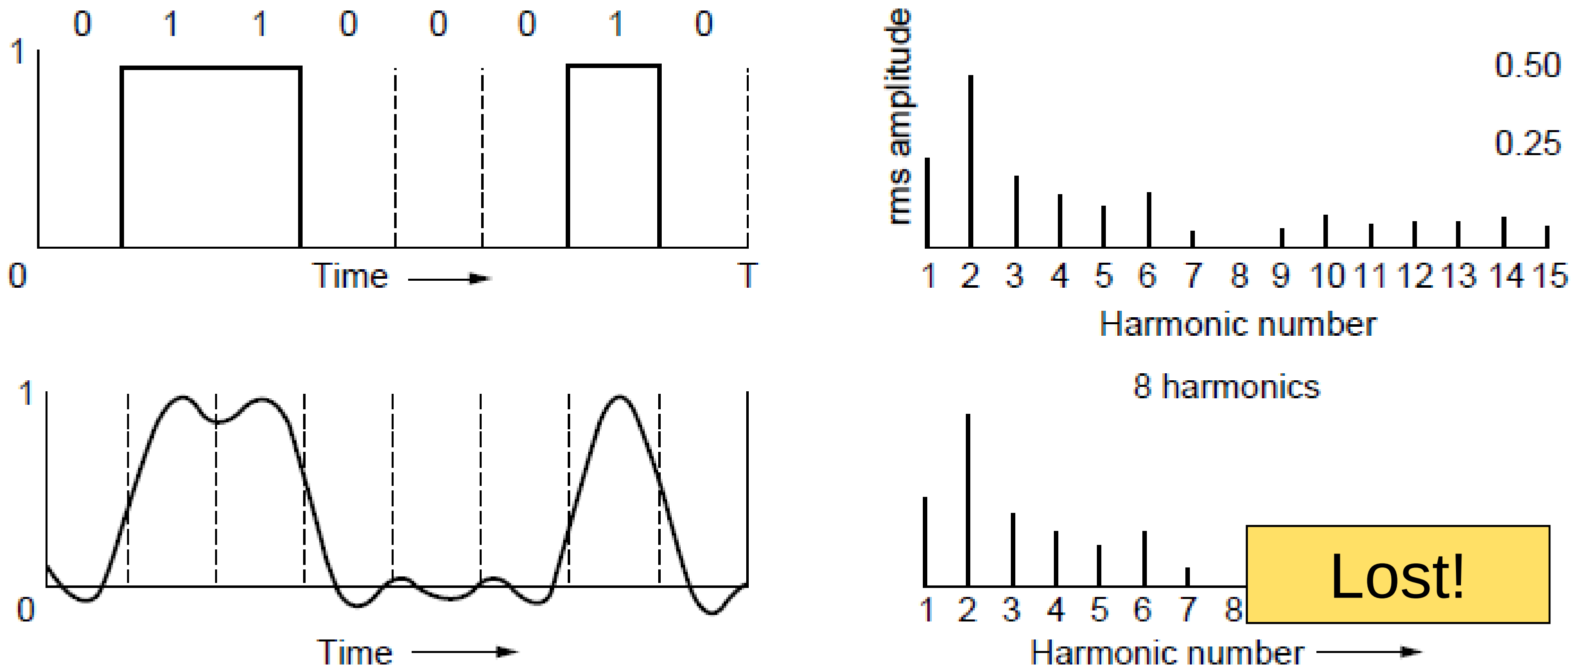

- Less bandwidth permits less rapid signal transitions

Signals over a wire

- If we send a waveform, what do we get out?
  - Signal loss
  - Delay
  - Frequency-dependent attenuation
  - Noise at the receiver

- What do we want to get out?
  - Fidelity versus interpretability

Nyquist Limit (~1924)

- For a noiseless channel with bandwidth B
- Symbols will be distorted, and sending too fast leads to Inter-symbol Interference (ISI)

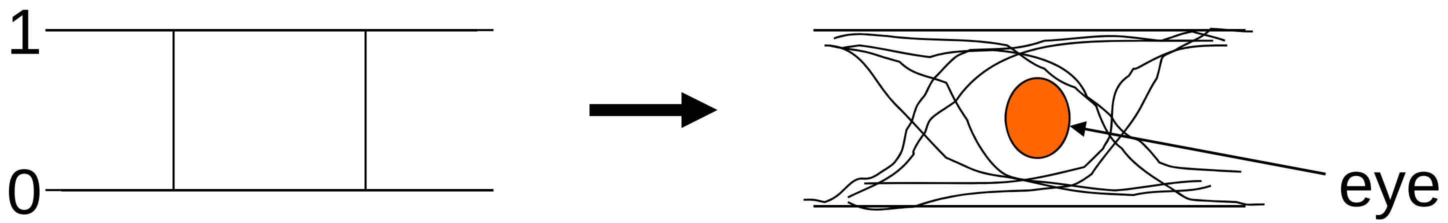

- The maximum rate at which it is possible to send:

  R = 2B symbols/sec

  e.g., 3KHz → 6Ksym/sec

Taking Noise into Account

- Noise limits how many signal levels we can safely distinguish between
  - S = max signal amp., N = max noise amp.

- The number of bits per symbol depends on the number of signal levels
  - E.g, 4 levels implies 2 bits / symbol

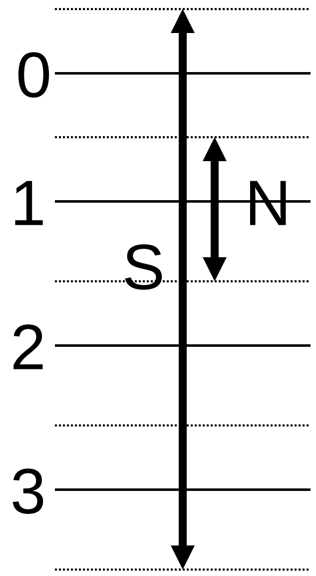

Shannon limit (1948)

$$\text{Capacity} = \text{Bandwidth} \times \log_2 \left(1 + \frac{\text{Signal}}{\text{Noise}}\right)$$

Capacity	Bandwidth	Signal/Noise
Bits/sec	How fast signal can change	How many signal levels can be seen

- Shannon, "A Mathematical Theory of Communication," 1948.
- SNR or Signal-to-Noise ratio defined as deciBels on a log scale
  - SNR = 10 $\log_{10}$ (Signal / Noise), e.g., 30dB = 1000 times

3. Encoding Bits with Signals

- Generate analog waveform (e.g., voltage) from digital data at transmitter and sample to recover at receiver

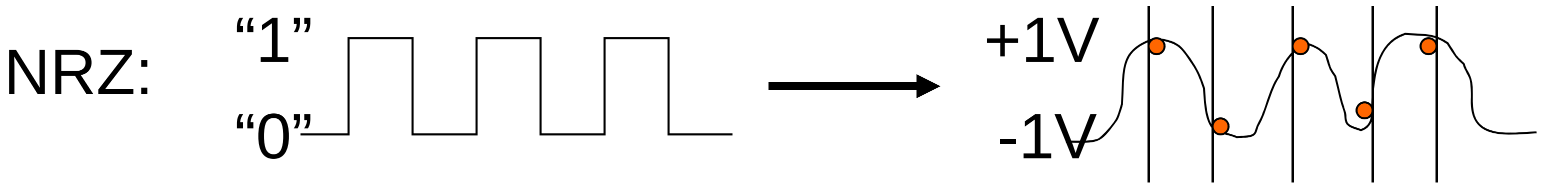

- We send/recover symbols that are mapped to bits
  - May have >2 different symbols, e.g., amplitudes
  - Thus distinguish symbol rate versus bit rate

- This is baseband transmission

NRZ

- Simplest encoding, NRZ (Non-return to zero)
  - Use high/low voltages, e.g., high = 1, low = 0

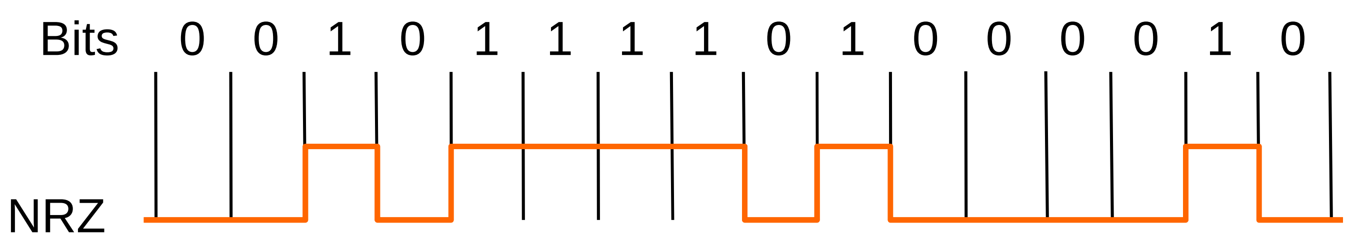

Issue: Clock recovery

- Um, how many zeros was that again?

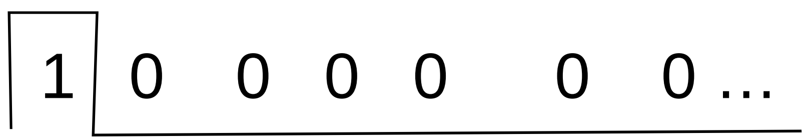

If sender and receiver have exact clocks no problem. But they don't!

- Any brilliant ideas?

Embed clock in signal (Manchester)

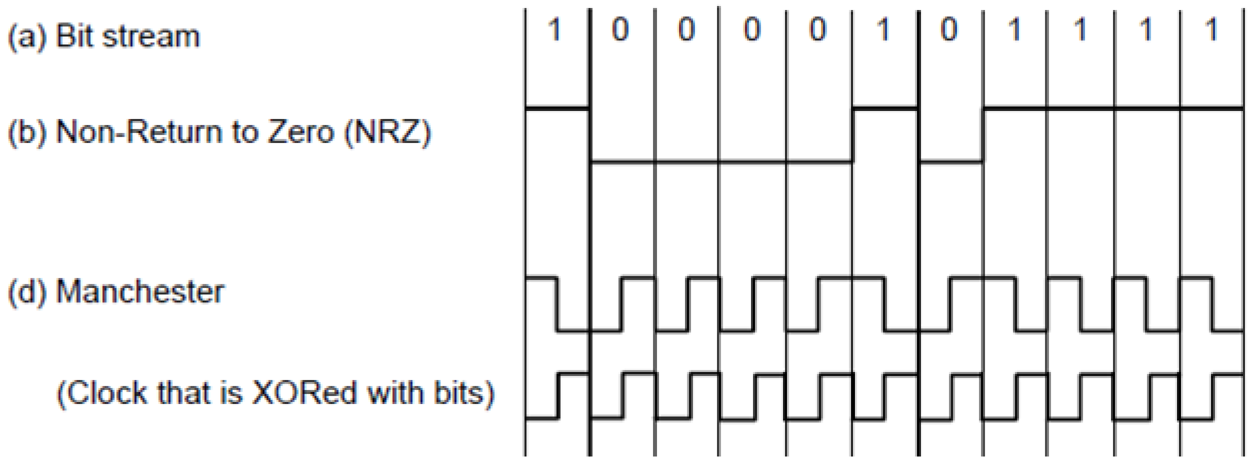

- Signal is XOR of data (NRZ) and clock (transition per bit)
  - Low-to-high is 0; high-to-low is 1
  - Signal rate is twice the bit rate
- Advantage: self-clocking, Disadvantage: BW inefficiency

4B/5B Codes

- We want self-clocking transitions *and* efficiency …
- Solution: map data bits (which may lack transitions) into code bits (which are guaranteed to have them)

- 4B/5B code:
  - 0000 → 11110, 0001 → 01001, … 1111 → 11101
  - Never more than three consecutive 0s back-to-back
  - 80% efficiency, plus use “illegal” codes as markers

- Many more complex codes are available; some use multiple voltage level

Scrambling

- XOR data with known pseudo-random sequence
  - Can generate cheaply with linear feedback shift registers (LFSR)
  - Causes transitions with reasonable probability
  - Also tends to whiten data (better for RF)

- Reverse at receiver by XORing with same sequence

Passband transmission

- For wireless, fiber, need to encode signal by modulating carrier wave … can't propagate at baseband
- Carrier frequency set by assigned bandwidth,
  - e.g., 2.45GHz WiFi
- Modulation: can change carrier
  - Amplitude
  - Phase/frequency

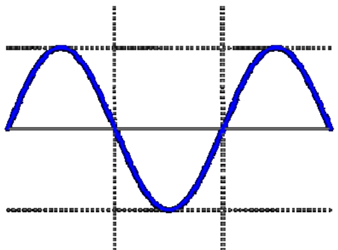

Modulation examples

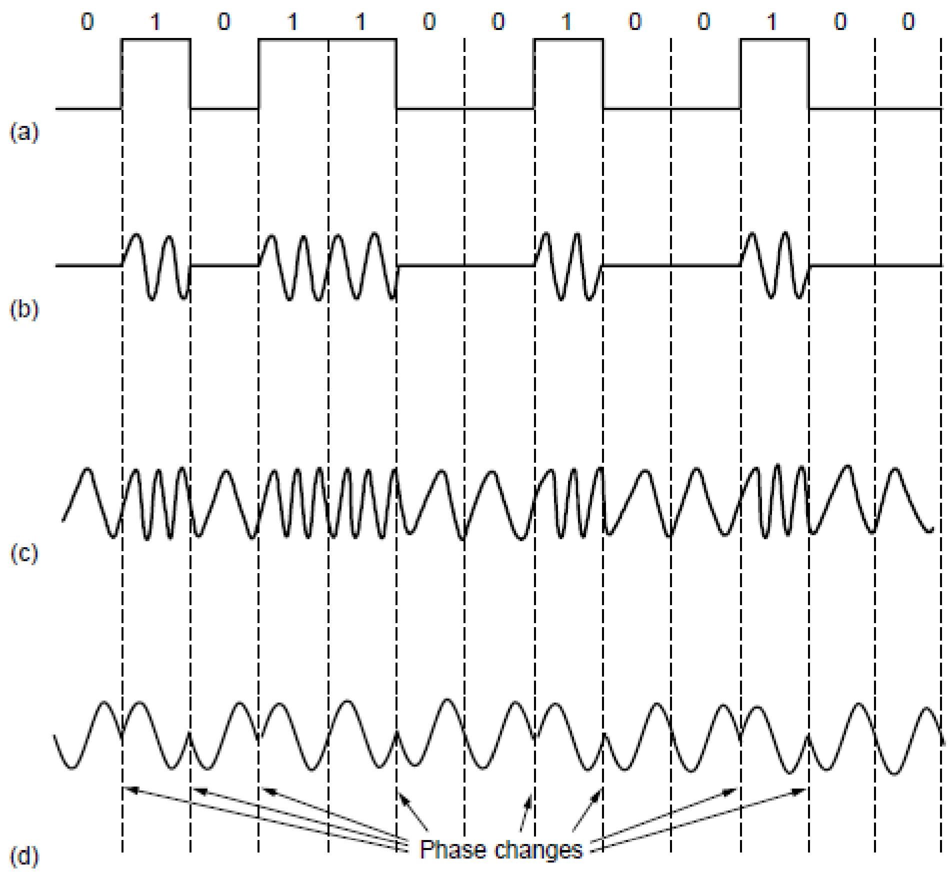

(a) A binary signal

(b) Amplitude shift keying

(c) Frequency shift keying

(d) Phase shift keying

Constellations

- Express modulation as a "constellation"
  - Points are amplitude/phase modulations for valid symbols

- Many names for schemes:
  - BPSK, QPSK … QAM

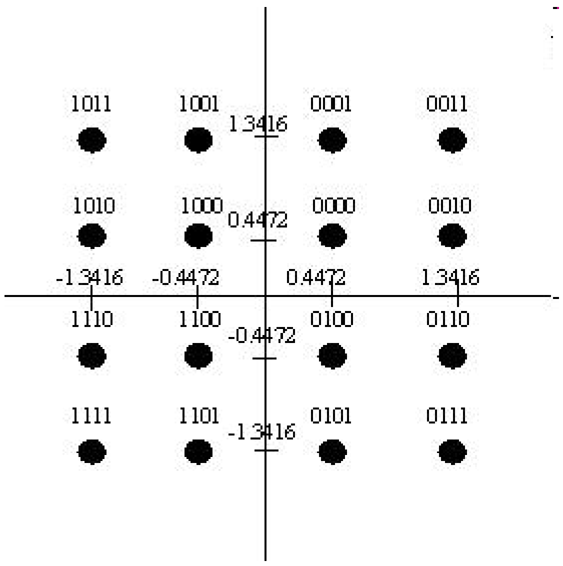

QAM 16 constellation in 3G (HSPDA)

BER versus SNR

- Need higher SNR for more complex modulations to keep a low bit error rate

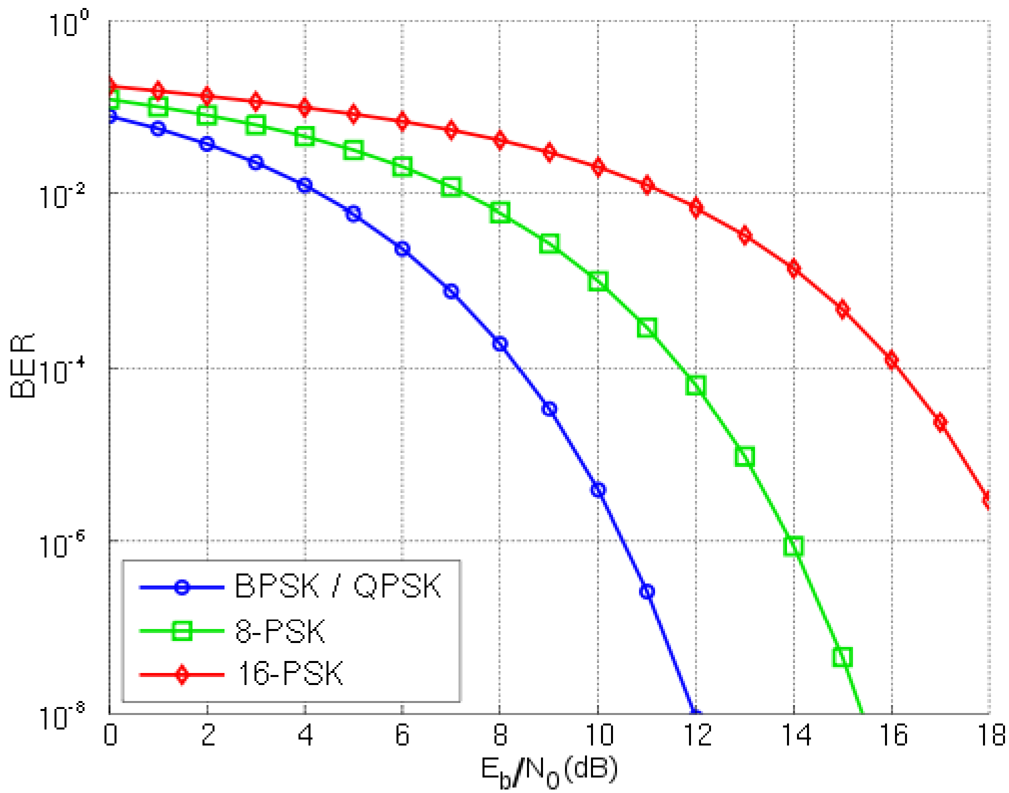

4. Abstract model of a link

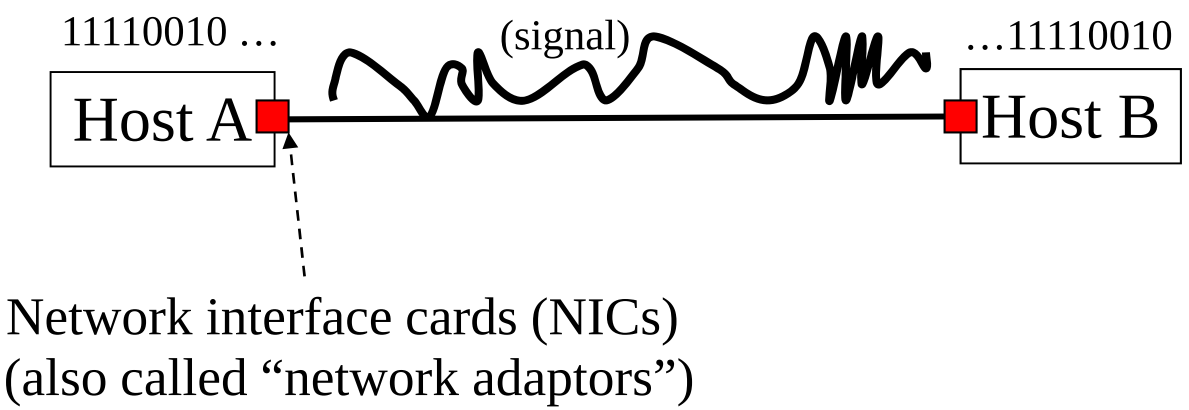

What really happens

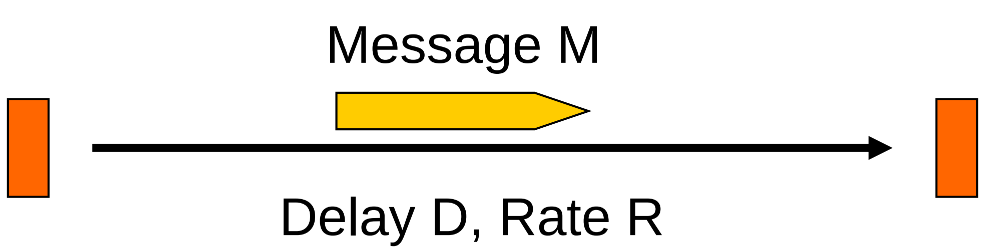

Abstract link for our purposes

Model

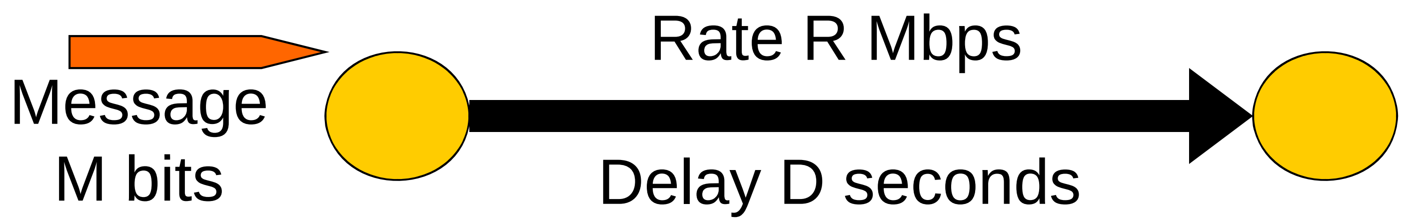

- Typically all we will need (but not so good for wireless!)

- Other parameters that are important:
  - Whether the media is broadcast or not
  - The kind and frequency of errors (bit error rate, BER)

Message Latency

- How long does it take to send a message?

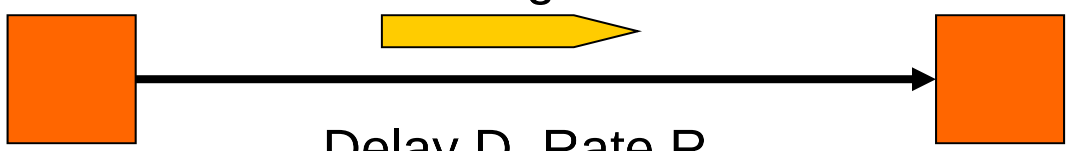

- Two terms:
  - Propagation delay = distance / speed of signal in media
    - How quickly a message travels over the wire
    - 2/3c for copper wire
  - Transmission delay = message (bits) / rate (bps)
    - How quickly you can inject the message onto the wire
- Later we will see queuing delay …

One-way Latency

Dialup with a modem:

- D = 10ms, R = 56Kbps, M = 1024 bytes
- Latency = 10ms + (1024 x 8)/(56 x 1024) sec = 153ms!

Cross-country with T3 (45Mbps) line:

- D = 50ms, R = 45Mbps, M = 1024 bytes
- Latency = 50ms + (1024 x 8) / (45 x 1024*1024) sec = 50ms!

- Either a slow link or long wire makes for large latency

Bandwidth-delay product: Messages occupy space on the wire

- Consider a 1b/s network, suppose latency is 16 seconds.
  - How many bits can the network "store"?
  - This is the <u>bandwidth-delay</u> product
  - Measure of "data in flight."
  - 1b/s * 16s = 16b

- Tells us how much data can be sent before a receiver sees any of it.
  - Twice BD tells us how much data we could send before hearing back from the receiver something related to the first bit sent.

BD Example

BD = 50ms * 45Mbps = 2.25 * 10^6 = 280KB

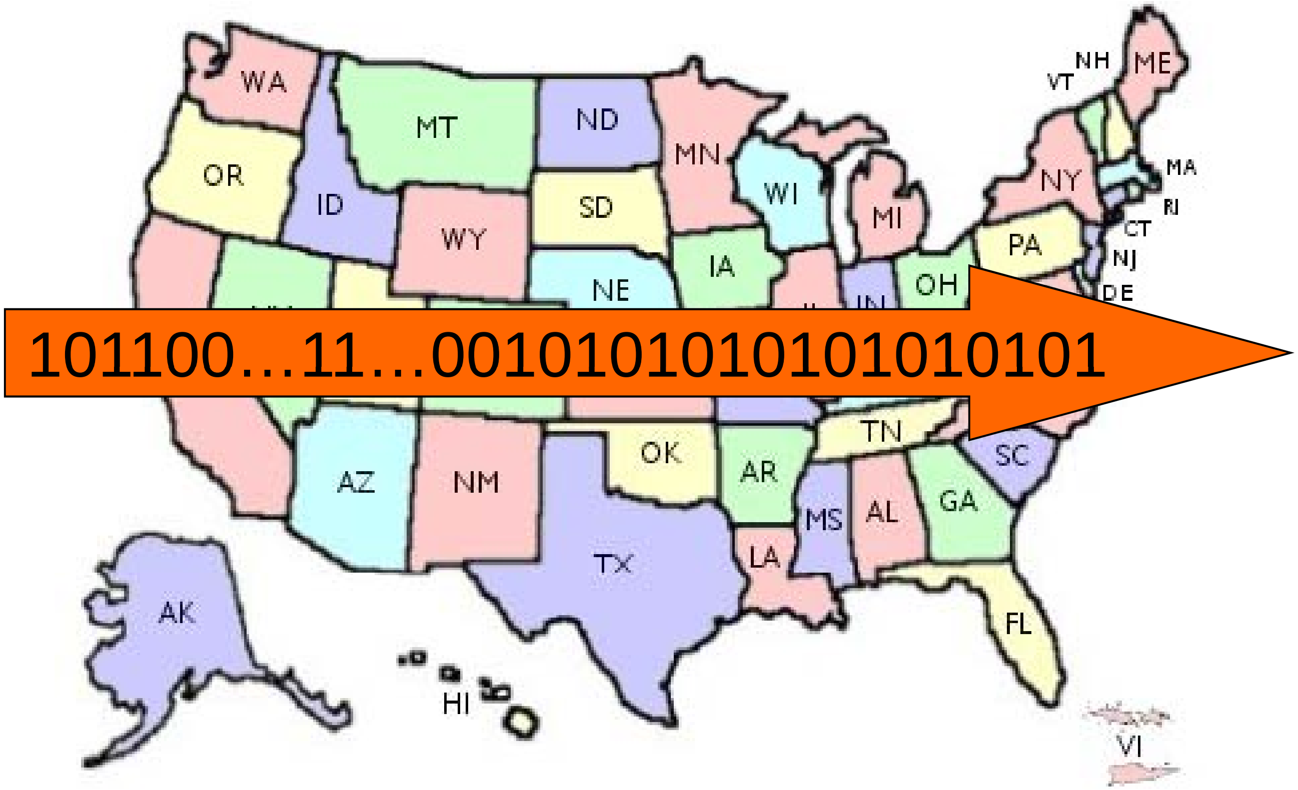

Wireless versus Wired links

- Wireless complications:
  - Broadcast channel has interference effects
  - Link capacity varies lots over time, e.g., as endpoints move
  - Which "wireless links" are up even varies over time

- Endpoint moves → SNR changes due to RF effects → rate must go down if SNR falls to keep low BER; or rate wants to go up if SNR rises to use spectrum efficiently

- Wired is about engineering the right link properties
- Wireless is about adapting to the channel capacity