CS244a: An Introduction to Computer Networks

Handout 12: Physical Layer

Sending 1's and 0's, Capacity and Clocking

Nick McKeown
Professor of Electrical Engineering and Computer Science, Stanford University
nickm@stanford.edu
http://www.stanford.edu/~nickm

Outline

- Bits at the physical layer
  - Representing 1's and 0's
  - What dictates the data-rate on a link?
- Clock recovery
  - How does a receiver know what data rate the sender used?
  - Elasticity buffers

Signaling bits on a link

Most electrical and optical networks signal bits using two distinct voltage/power levels.

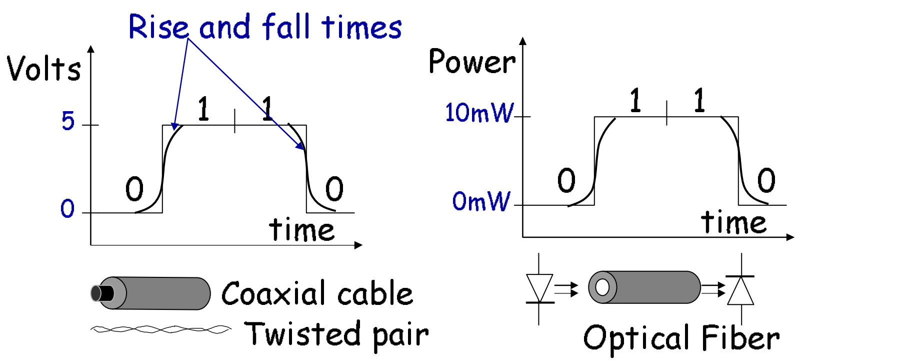

Signaling bits on a link

All links have a maximum bandwidth, and hence a minimum rise and fall time. Intuitively, this limits how close together consecutive bits can be placed, and so limits the maximum capacity or data rate of the link.

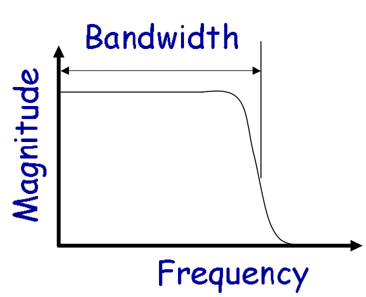

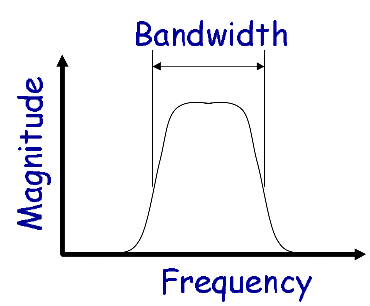

Signaling bits on a link

Multi-level Signaling

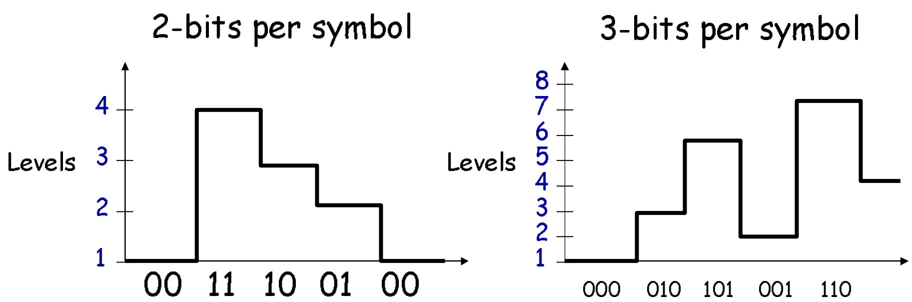

Ultimately, what limits the number of bits I can send per symbol?

Signaling bits on a link

Frequency division multiplexing (FDM)

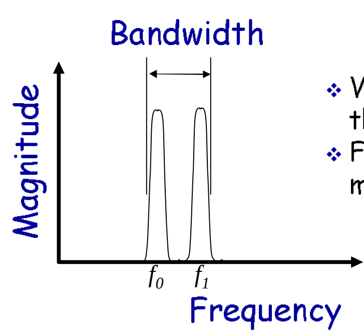

- What limits how close together the frequencies can be?
- For a given bandwidth, how many frequencies can we use?

Coding schemes

- High speed links use a selection of complicated techniques to "squeeze" the maximum data-rate out of the link that they can.
  - Techniques: FDM, phase modulation, multi-level signaling, CDMA, pulse-position modulation (PPM), ...
  - Links: Modems, DSL (Digital Subscriber Line), Fast Ethernet, Gigabit Ethernet, Wireless Ethernet, ...
- Capacity: the maximum data rate of a link
  - Why does the output of a modem sound like "white noise"?
  - What ultimately limits the capacity?

Maximum Capacity/Data Rate

Shannon Capacity:

$$C = B\log_2(1 + S/N)$$

Bandwidth of link (B) — Signal-to-Noise ratio (S/N)

For example:

- Bandwidth of telephone link from telephone to a typical home is approx 3300Hz - 300Hz = 3kHz
- Signal-to-noise ratio is approx 30dB = $10\log_{10}(S/N)$
- Therefore, $C = 3000 * \log_2(1001)$ ~= 30kb/s

Optical fiber has a higher capacity because the bandwidth, B, of a fiber is much greater than for wire; and it is less susceptible to noise, N.

Outline

- Bits at the physical layer
  - Representing 1's and 0's
  - What dictates the data-rate on a link?
- Clock recovery
  - How does a receiver know what data rate the sender used?
  - Elasticity buffers

Encoding for clock recovery

Problem:

- Different hosts use locally-generated clocks of nominally the same frequency, but slightly different.
  E.g. 10MHz +/- 100ppm ("parts per million")[1].
- The receiver needs to "recover" the senders clock from the data stream, for example:

Sender **Receiver**

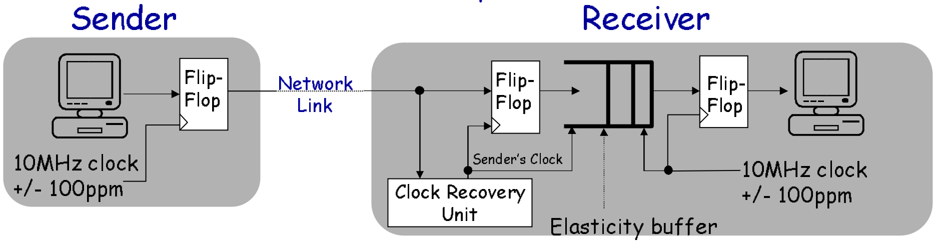

1) One part per million equals 10^{-4}%.

If we don't know the sender's clock

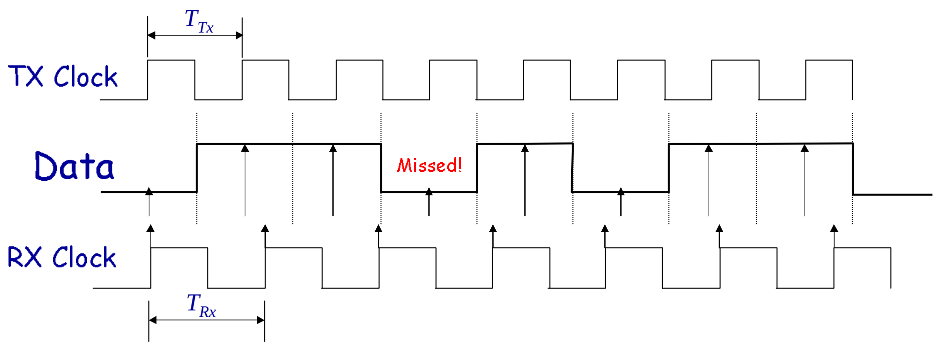

If the RX clock is p ppm slower than the TX clock, then: $T_{Rx} = T_{Tx}(1+10^{-6}p)$.

After $\frac{0.5}{10^{-6}p}$ bit times, the RX clock will miss a bit.

Asynchronous communication

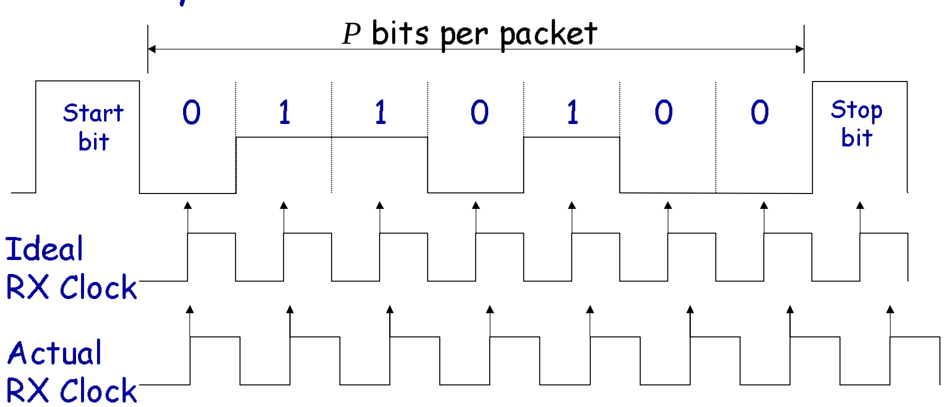

Asynchronous communications sometimes used for links with short packets.

Encoding for clock recovery

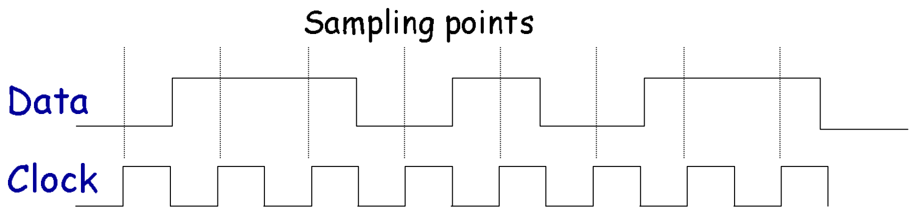

It is more common for the receiver to recover the clock from the received data stream.
If the clock is not sent separately, the data stream must have sufficient transitions so that the receiver can determine the clock.

Encoding for clock recovery

Example #1: Ethernet

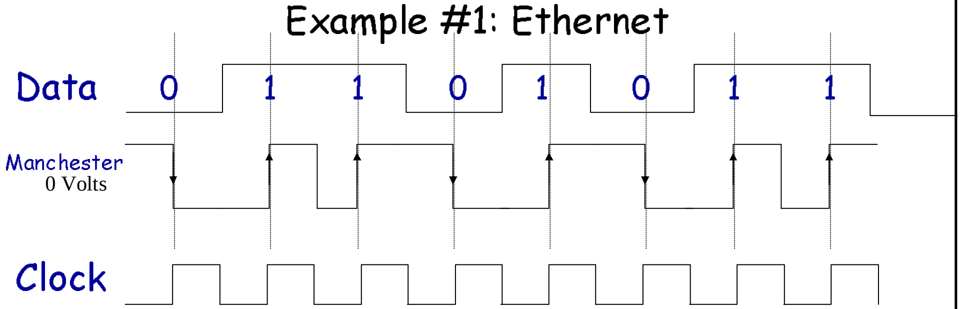

Advantages of Manchester encoding:

- Guarantees one transition per bit period
- Ensures d.c. balance (i.e. equal numbers of hi and lo)

Disadvantages

- Doubles bandwidth needed

The threshold between hi and lo can be set at the long-term average value.

Frequency Spectrum for 10Mb/s Ethernet

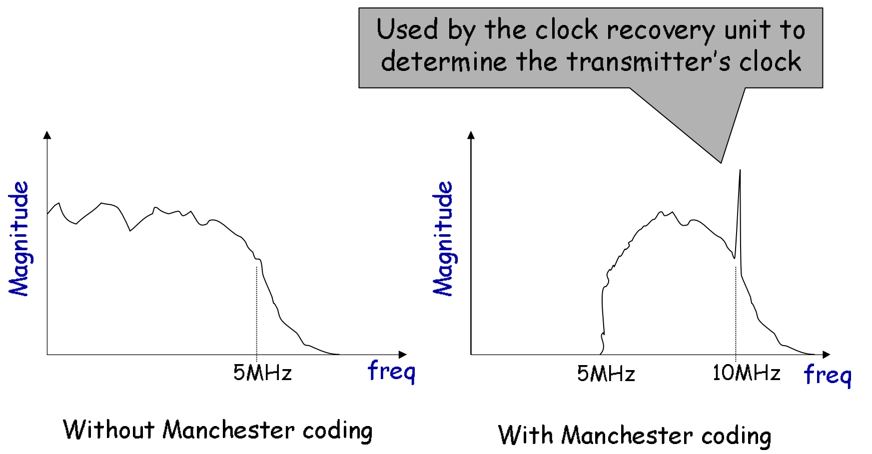

Encoding for clock recovery

Example #2: FDDI

4-bit data	5-bit code
0000	11110
0001	01001
0010	10100
...	...

Advantages of 4b/5b encoding:

- More bandwidth efficient (only 25% overhead).
- Allows extra codes to be used for control information.

Disadvantages

- Fewer transitions can make clock recovery harder.

Outline

- Bits at the physical layer
  - Representing 1's and 0's
  - What dictates the data-rate on a link?
- Clock recovery
  - How does a receiver know what data rate the sender used?
  - Elasticity buffers

The need for an elasticity buffer

Problem:

- The sender's clock may be slower or faster than the receiver's clock. e.g. 10MHz +/- 100ppm ("parts per million").
- Clock tolerance $\equiv \frac{R_{max} - R_{min}}{R}$.
- How big should the FIFO be?

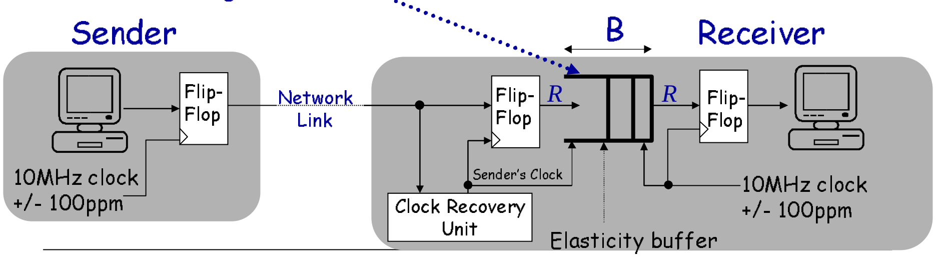

Sizing an elasticity buffer

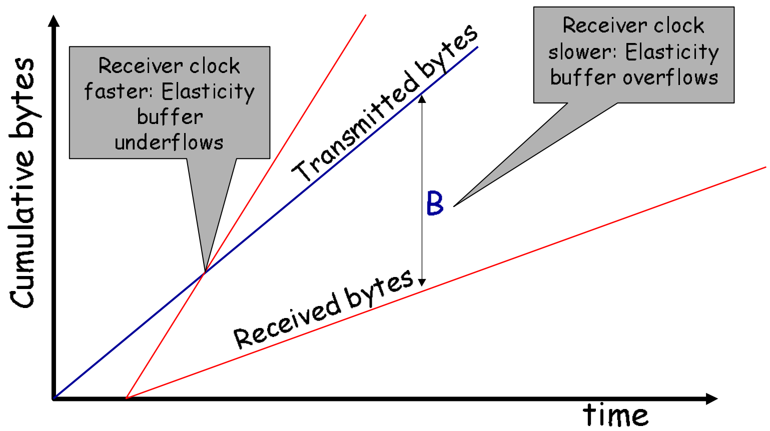

Sizing an elasticity buffer

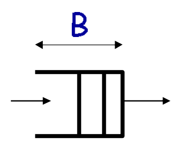

1. At start of new packet, allow buffer to fill to B/2.
2. Size buffer so that it does not overflow or underflow before packet completes.
3. Ensure that the inter-packet gap is long enough to allow buffer to drain before next packet arrives.

Preventing overflow

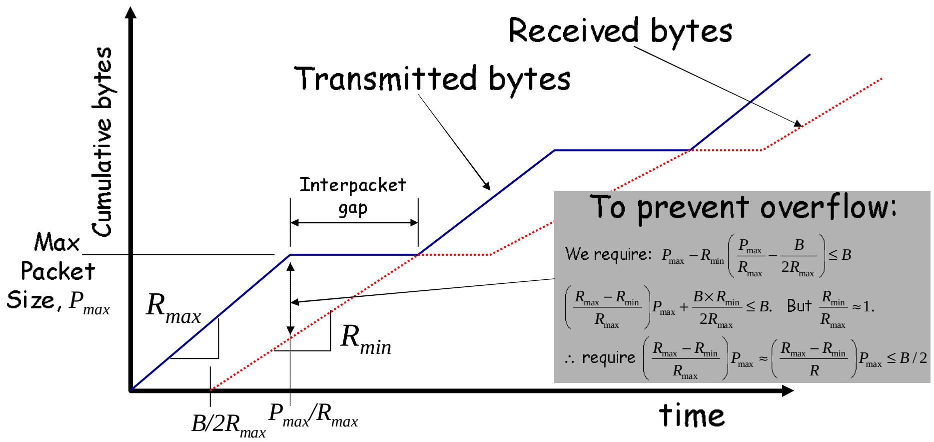

To prevent overflow:

We require: $P_{\max} - R_{\min}\left(\frac{P_{\max}}{R_{\max}} - \frac{B}{2R_{\max}}\right) \le B$

$\left(\frac{R_{\max} - R_{\min}}{R_{\max}}\right)P_{\max} + \frac{B \times R_{\min}}{2R_{\max}} \le B.$ But $\frac{R_{\min}}{R_{\max}} \approx 1.$

$\therefore$ require $\left(\frac{R_{\max} - R_{\min}}{R_{\max}}\right)P_{\max} \approx \left(\frac{R_{\max} - R_{\min}}{R}\right)P_{\max} \le B/2$

Preventing underflow

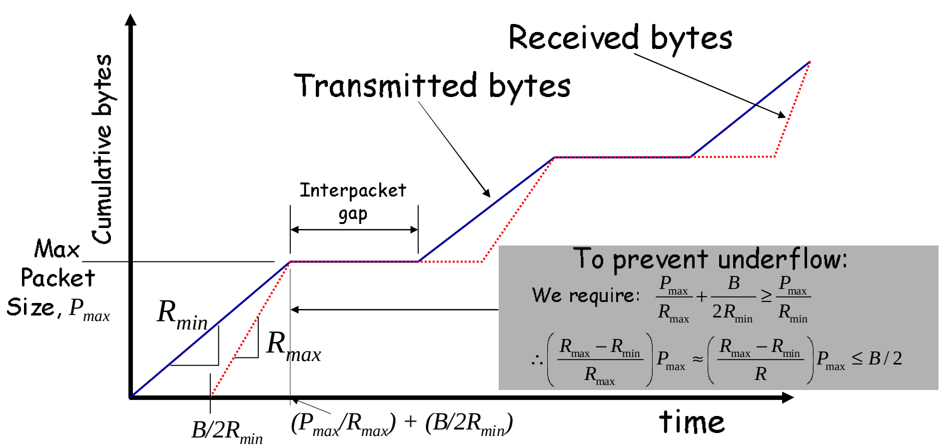

To prevent underflow:

We require: $\frac{P_{\max}}{R_{\max}} + \frac{B}{2R_{\min}} \ge \frac{P_{\max}}{R_{\min}}$

$\therefore \left(\frac{R_{\max} - R_{\min}}{R_{\max}}\right)P_{\max} \approx \left(\frac{R_{\max} - R_{\min}}{R}\right)P_{\max} \le B/2$

Sizing an elasticity buffer

Example: FDDI

Maximum packet size 4500bytes

Clock tolerance +/- 50ppm

$$\left(\frac{R_{\max} - R_{\min}}{R}\right) = 100 \times 10^{-6}$$

$$\therefore B \geq 2(4500 \times 8 \times 100 \times 10^{-6}) = 7bits$$

Therefore,

1. Buffer larger than 7 bits
2. Wait for at least 3.5 bits before draining buffer
3. Inter-packet gap at least 3.5bits