CS144
An Introduction to Computer Networks

The Link Layer

aka Physical Layer

Nick McKeown
Professor of Electrical Engineering
and Computer Science, Stanford University

Goals for today

1. **Capacity**:
   What determines the maximum data rate of a link?
   - ► How can we get close to the maximum capacity?

2. **Clocks**: Two communicating entities cannot have exactly the same clock or frequency. How can they communicate?

The 4 Layer Internet Model

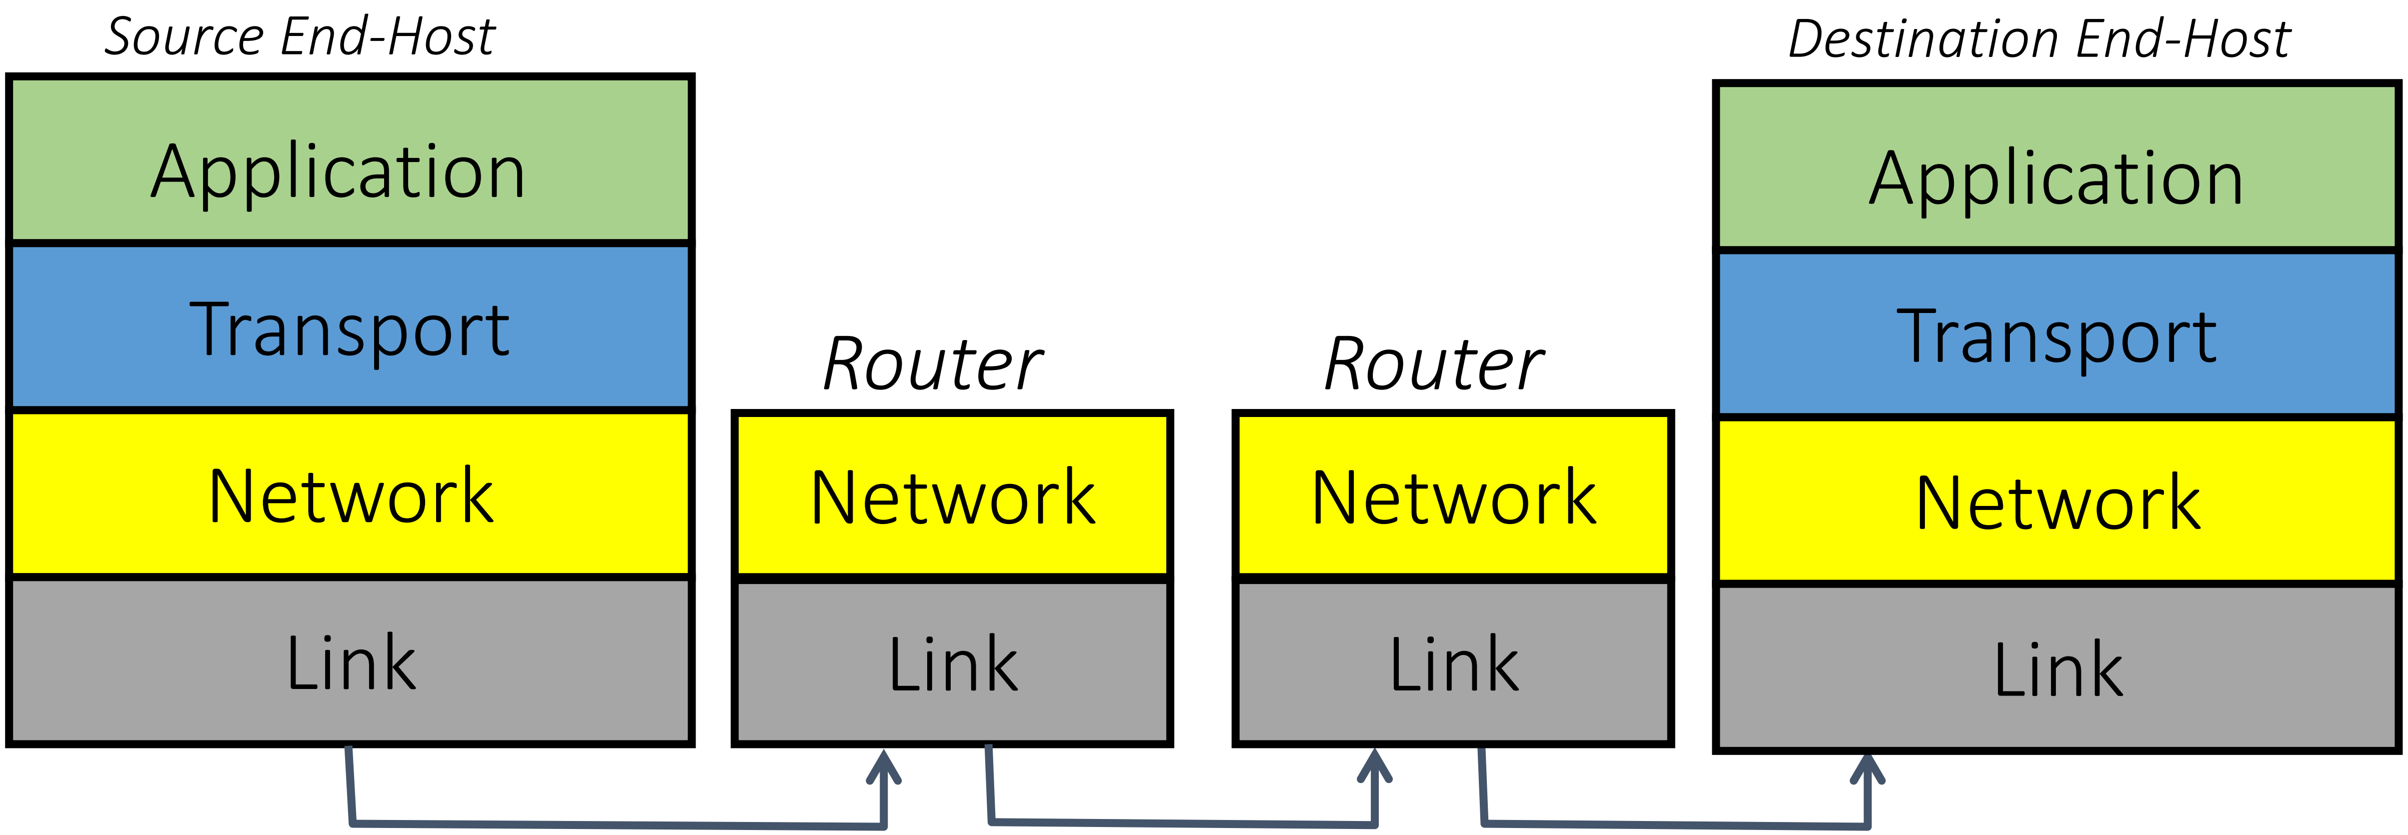

Total time to send a packet across a link: The time from when the first bit is transmitted until the last bit arrives.

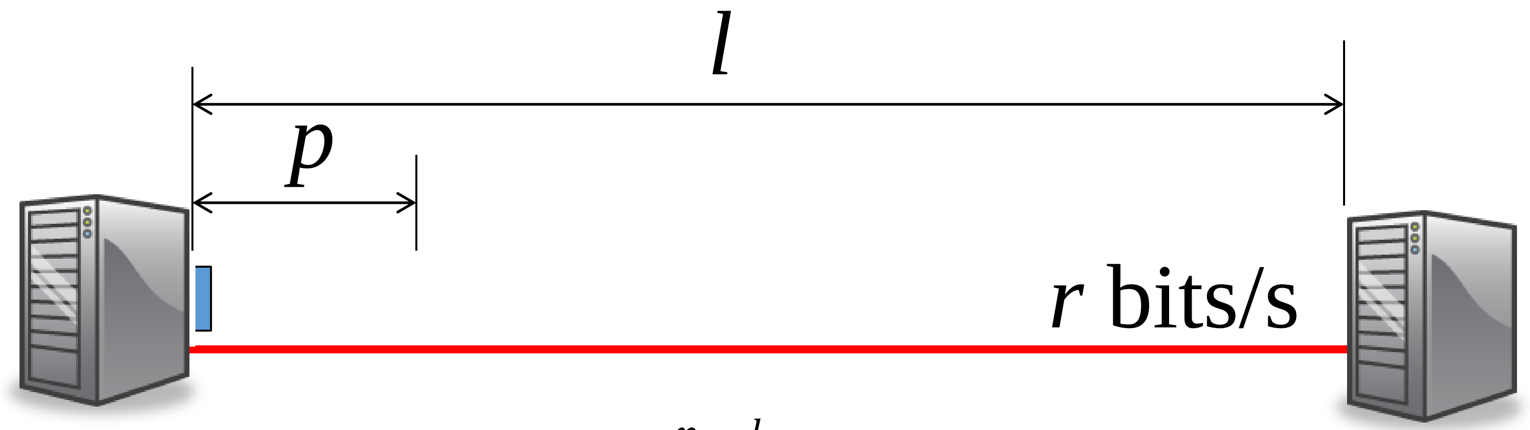

$$t = t_p + t_l = \frac{p}{r} + \frac{l}{c}$$

Example: A 100bit packet takes 10 + 5 = 15μs to be sent at 10Mb/s over a 1km link.


LINKSYS
Secure
Easy
Setup
Power
DMZ
USB
Wireless-G
Ethernet
WLAN 1 2 3 4 Internet
Wireless-G
Broadband Router
With 4-Port Switch
Model WRT54GL

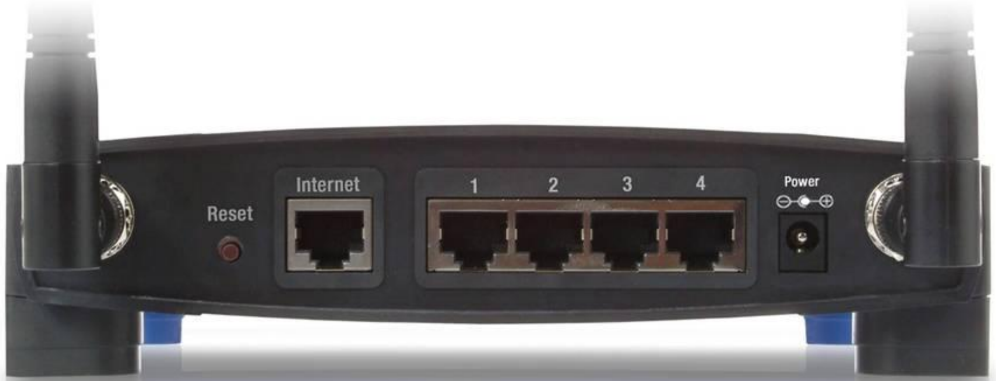
Reset
Internet
1
2
3
4
Power

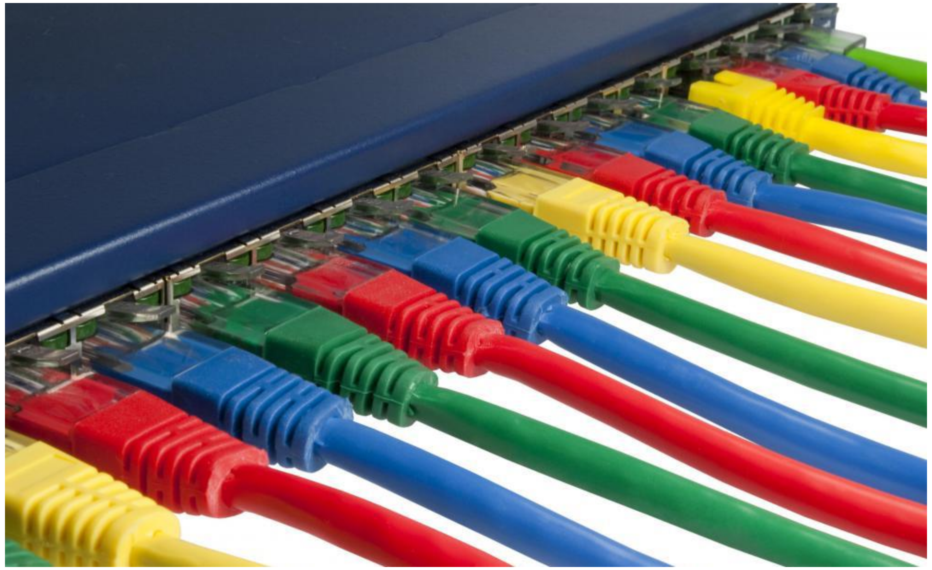

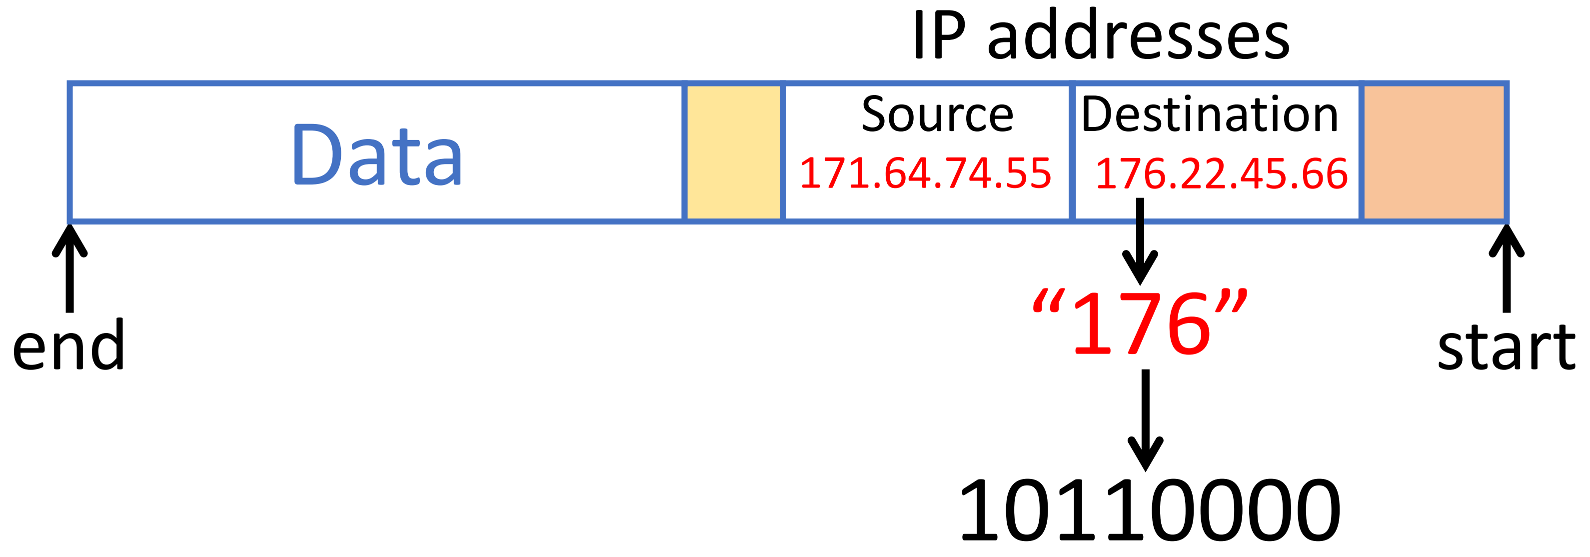
IP addresses
Data
Source
171.64.74.55
Destination
176.22.45.66
end
start
"176"
10110000

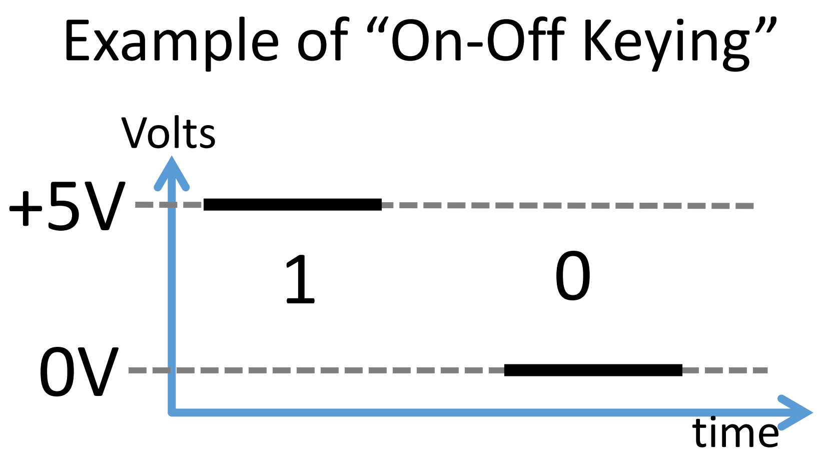
Example of "On-Off Keying"
Volts
+5V
0V
1
0
time

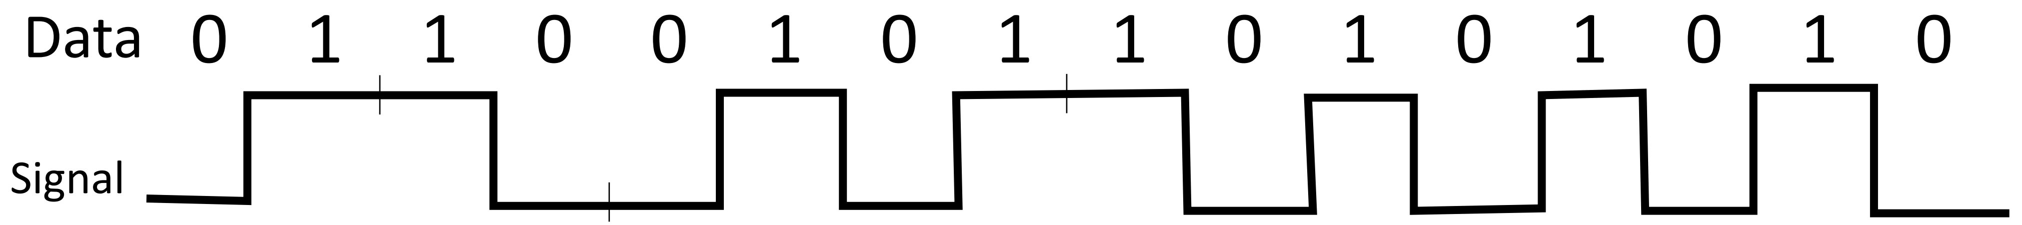
Data
0 1 1 0 0 1 0 1 1 0 1 0 1 0 1 0
Signal

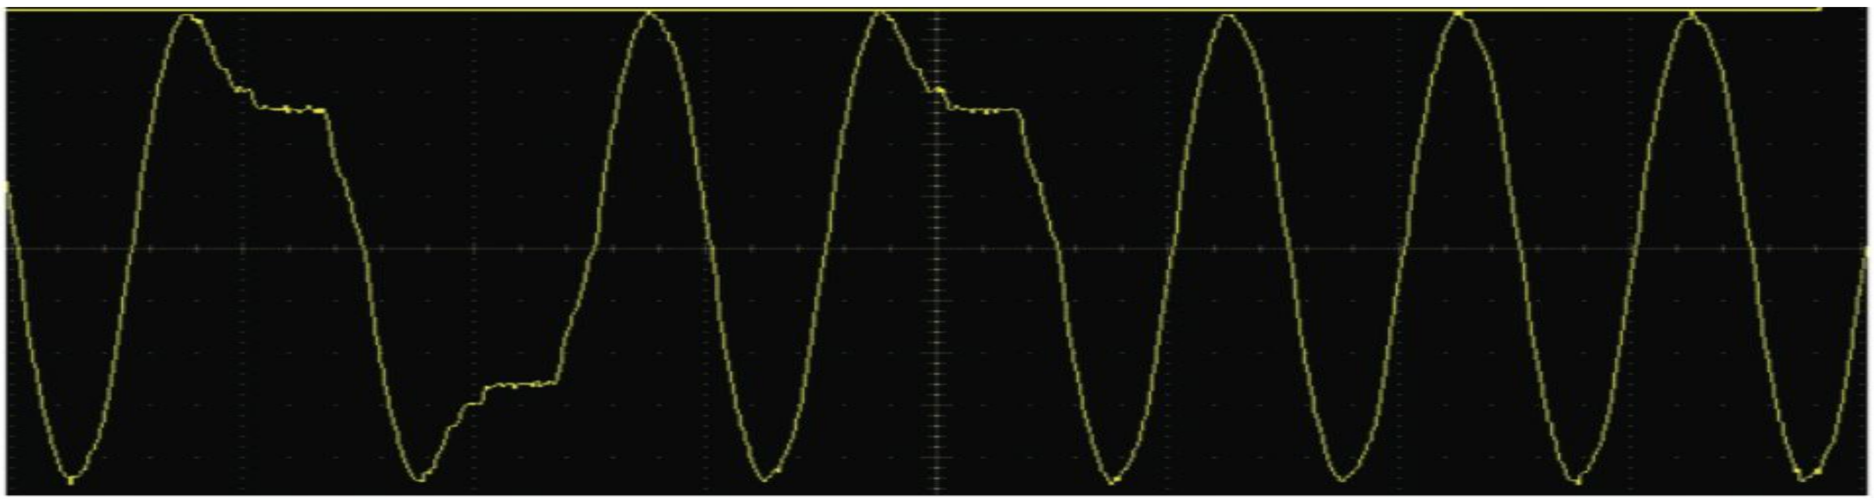

Keith: So what was this about?

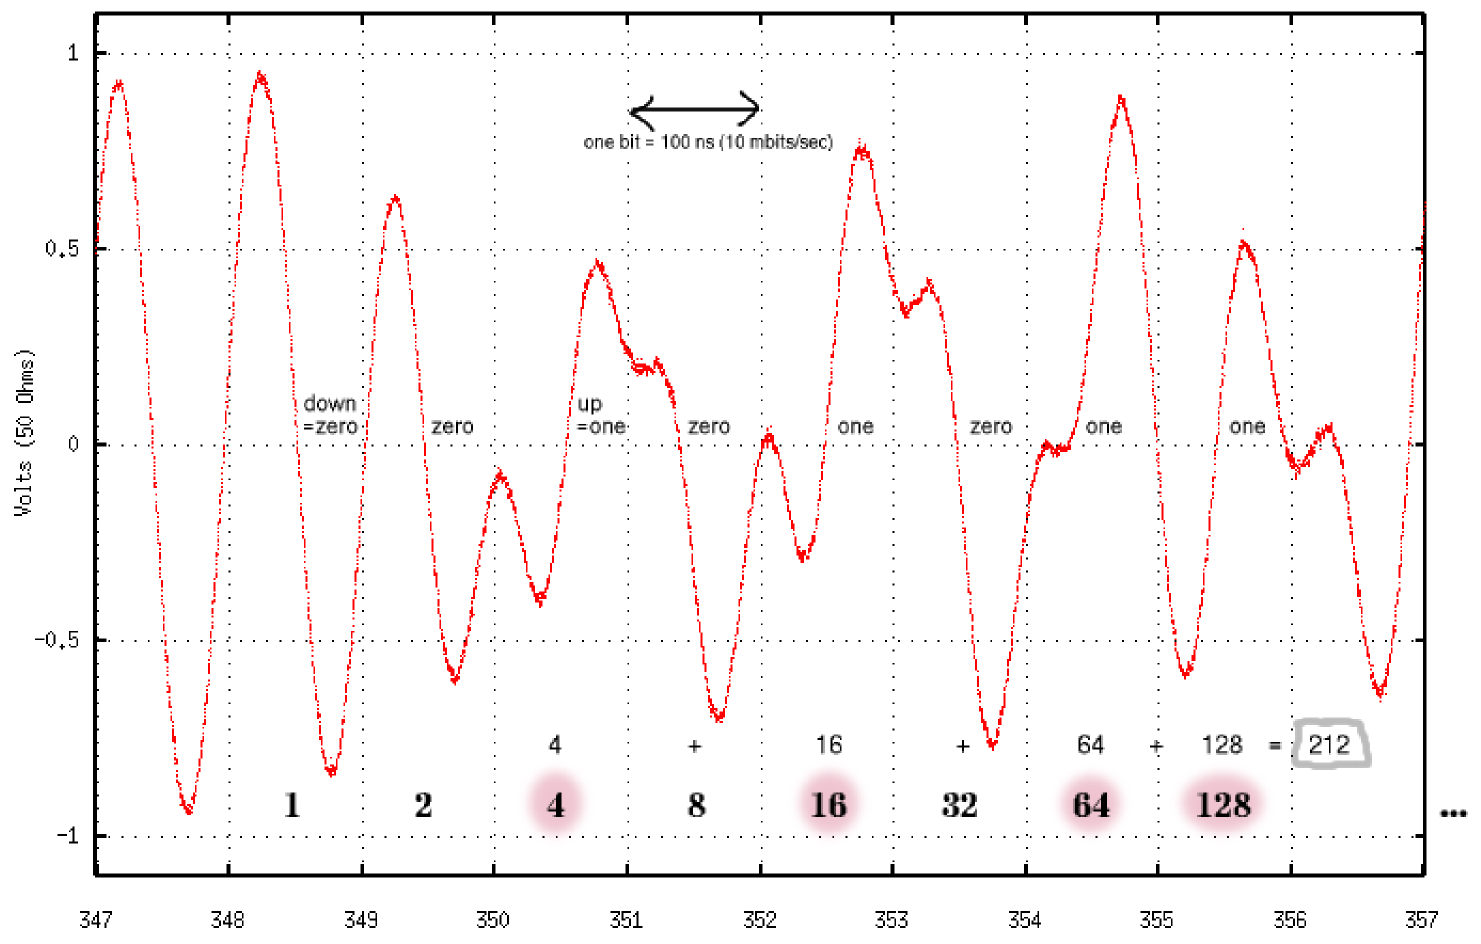

What determines the data rate?

Q: What determines the steepness (i.e. rate) of this change?

Q: How does the rate of change affect the data rate?

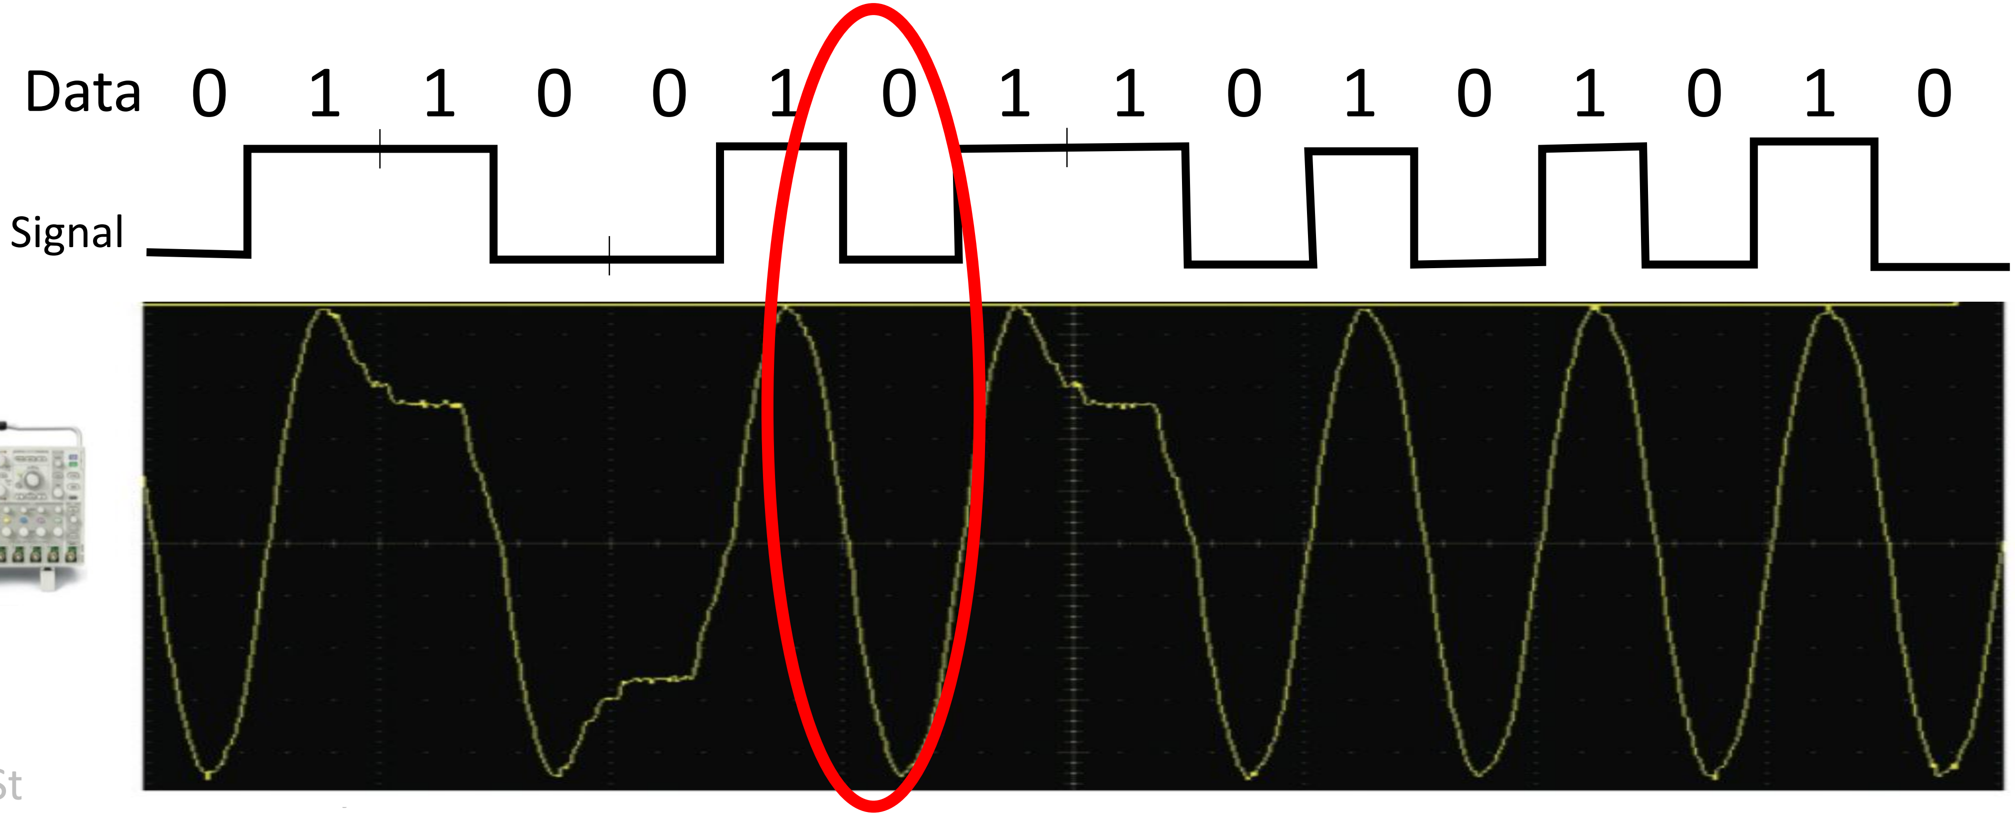

Fiber-optic communication links

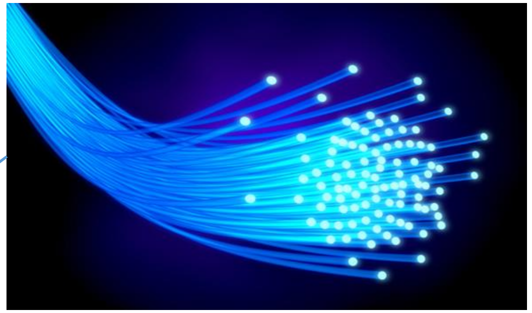

Each fiber is smaller than a human hair

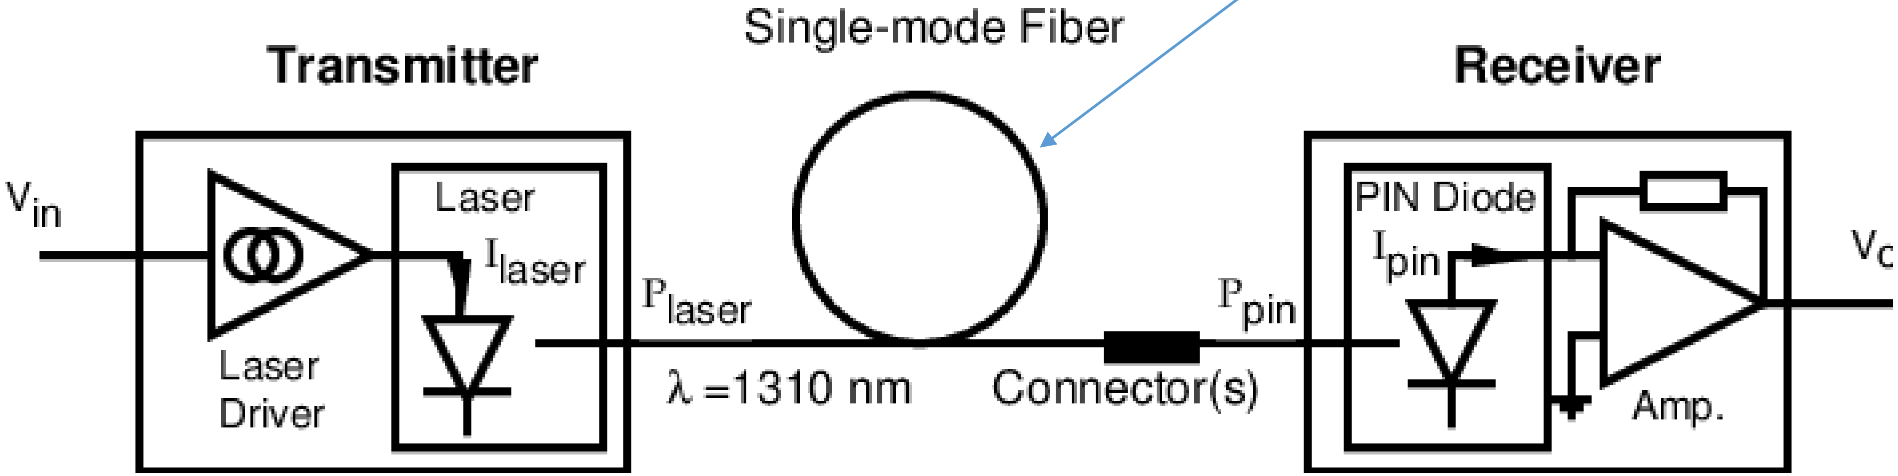

Commercial lasers can be be switched on/off at rates above 100GHz

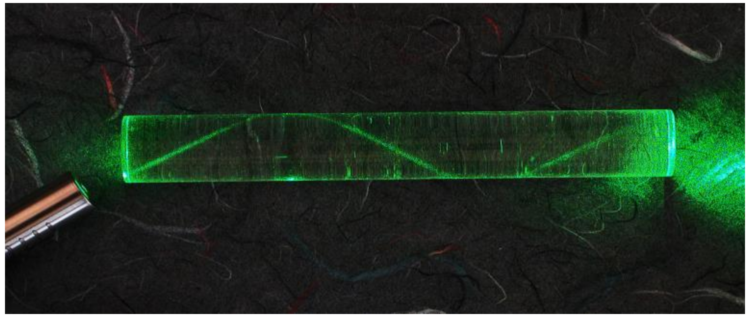

Light travels through the fiber because of total internal reflection

Characteristics of fiber-optical links

High bandwidth:

- Lasers can be switched on/off very fast. (“very steep rate of change”)
- The signal remains the same for a long distance (“low dispersion or spreading”)
- Multiple signals of different colors can be sent over a cable simultaneously (“wavelength division multiplexing” or WDM).

Low signal loss: Typically 0.3dB/km (*)

Electromagnetic interference immunity: An opaque jacket means light cannot enter the cable. Total internal reflection means it is very hard to detect the signal outside the cable.

(*) Power loss (dB) = $10log_{10}\frac{P_1}{P_0}$

For example: If power is reduced from 1W to 1mW, we would write: Power loss (dB) = $10log_{10}\frac{P_1}{P_0}$ = -30dB, or simply a “30dB loss”. 0.3dB/km means about 7% is lost per km.

Fiber-optic communication links

Optical fibers and modules on the front panel of communication equipment

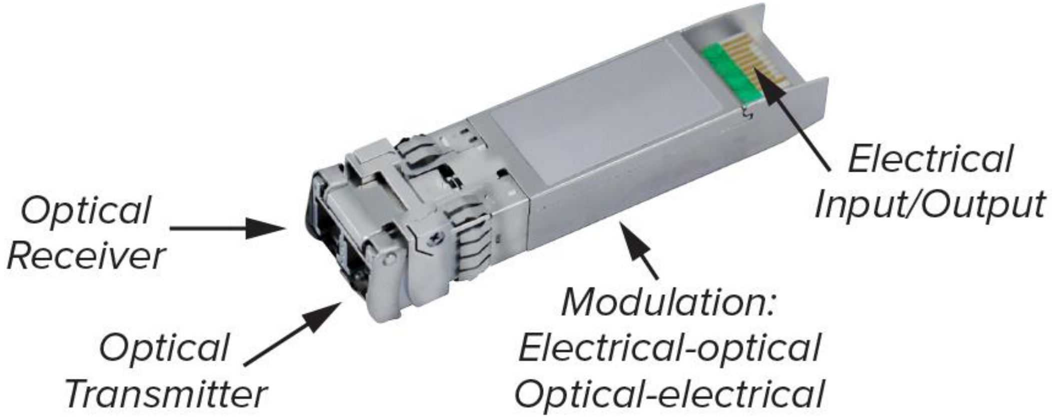

Typically used for very long and/or very fast communications

- Long haul links in the public Internet, e.g. 200km @ 400 Gb/s
- To connect buildings on campus, e.g. 2km @ 100Gb/s
- Between racks of equipment in a data center, e.g. 50m @ 100Gb/s
- World-record NEC (2011) 101.7 Tb/s over 370 WDM channels at 165km!

What determines the maximum data rate of a cable, fiber, wireless link, etc?

Q: What happens if we put the "bits" closer and closer together?

Q: If we can't put them closer together, how can we increase the number bits of information transmitted per second?

Q: What other factors limit the number of bits per second we can transmit?

Q: Are there any other factors other than "Bandwidth" and "Noise" that determine the maximum data rate of a channel?

Claude Shannon

1937: MS Thesis proposed using Boolean algebra for digital circuit design.

1948: "A Mathematical Theory of Communication" led to the field of Information Theory and Shannon Capacity

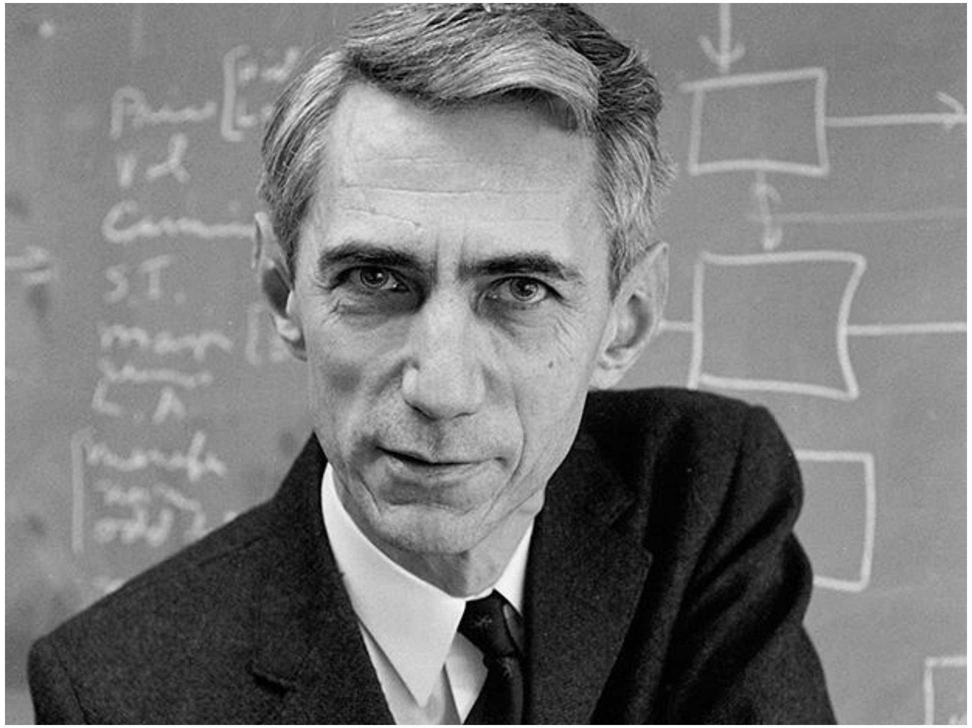

Claude Shannon (1916 – 2001)
Mathematician, Electrical Engineer

Shannon's Juggling Machines https://www.youtube.com/watch?v=dyc5bgpY86c

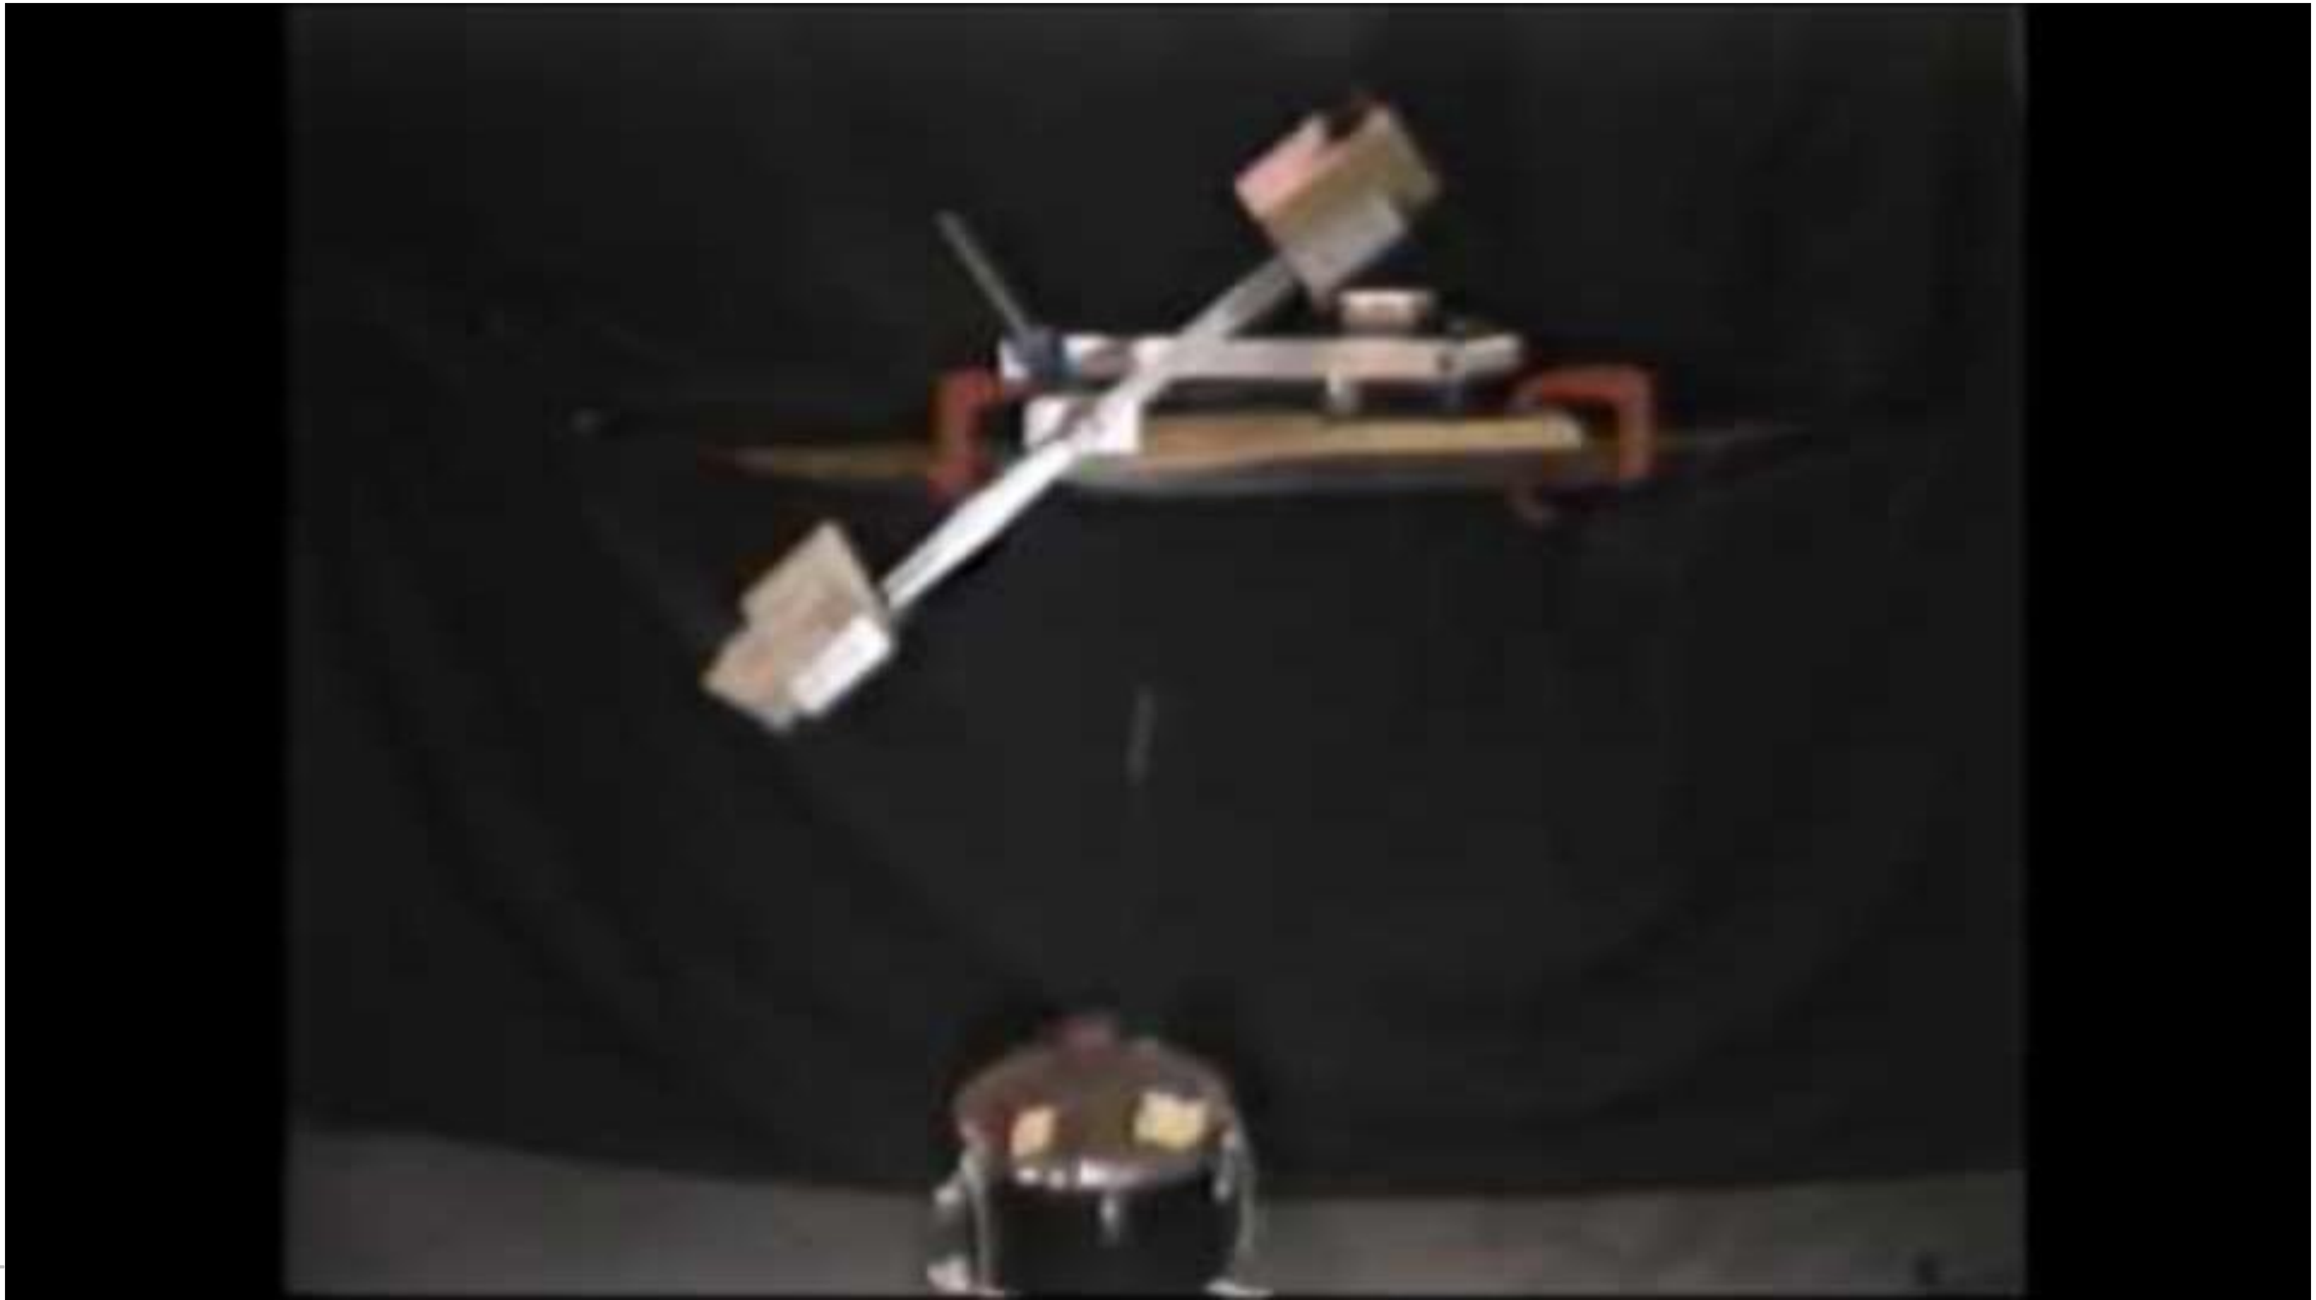

Shannon Capacity

- Shannon capacity is the maximum error-free rate we can transmit through a channel (e.g. wire, fiber, air, …).
- The maximum "data rate".
- Under some mild assumptions:

$$\text{Shannon Capacity} = \text{B}\ log_2\left(1 + \frac{S}{N}\right)$$

- In other words, it depends only on Bandwidth and Signal-to-Noise ratio!
- EE376A: Information Theory. Wow.

Shannon Capacity

But: It doesn't tell us how to achieve the capacity.

Hence: Huge numbers of researchers and engineers have spent decade trying to approach it, in different contexts:

- WiFi
- Cellular telephones
- Ethernet
- Optical Fibers
- ADSL Broadband access
- Modems
- …

Analog signals

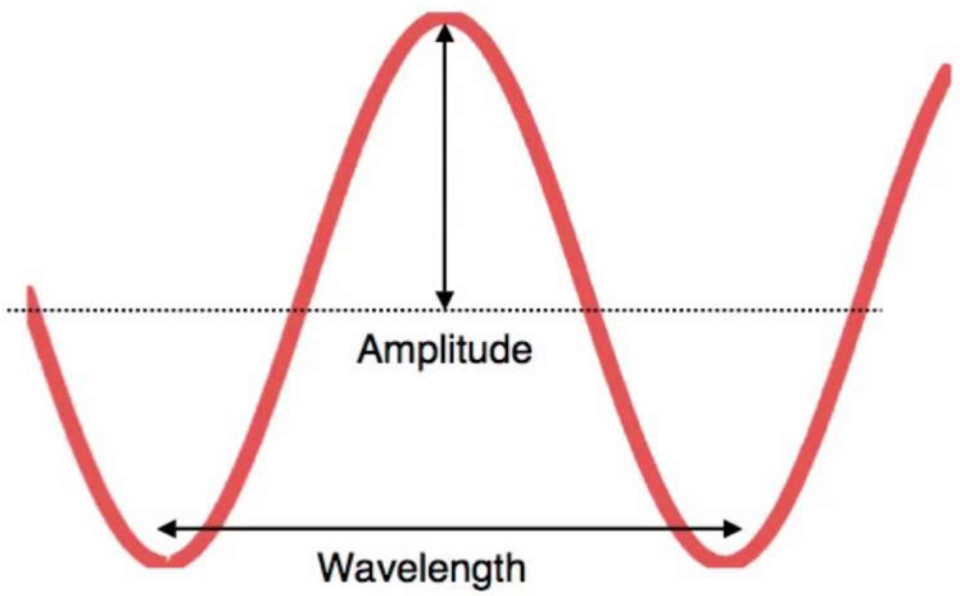

Frequency = 1/wavelength
Bandwidth: size of frequency range
Phase: location of peak within the wavelength

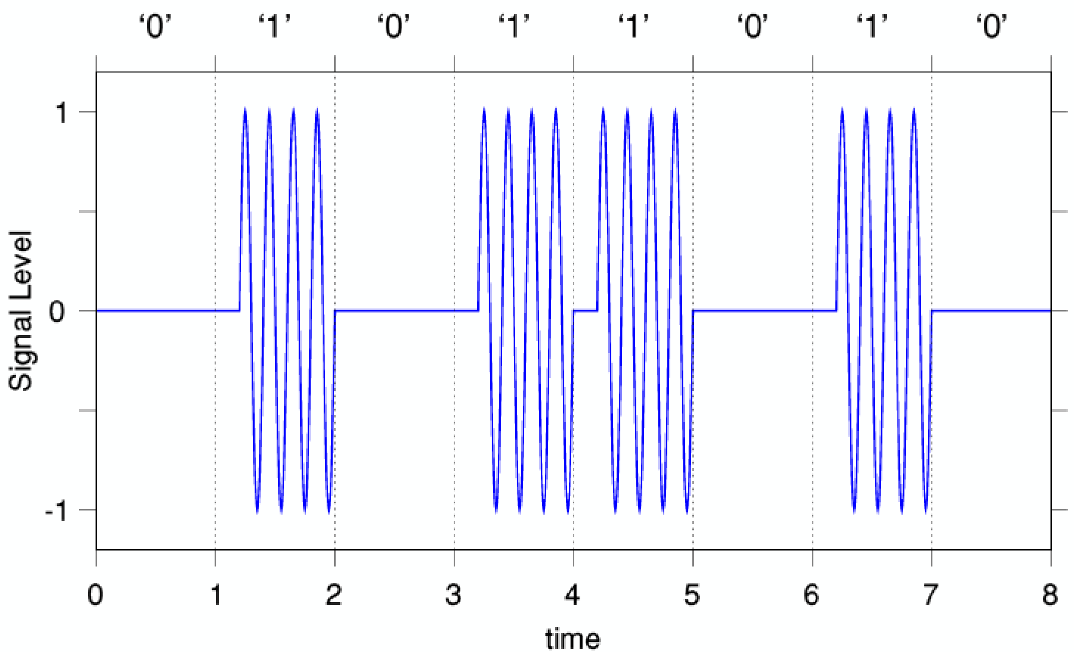

Figure 19.1 Simple on/off binary keying.

On-Off Keying (OOK)

- One frequency
- 2 amplitudes

Sending 0s and 1s

Frequency Shift Keying (FSK)

Amplitude Shift Keying (ASK)

Phase Shift Keying (PSK)

- For the same frequency + amplitude, vary the phase
- No variation in power (amplitude) or wavelength (frequency)

Phase in Analog signals

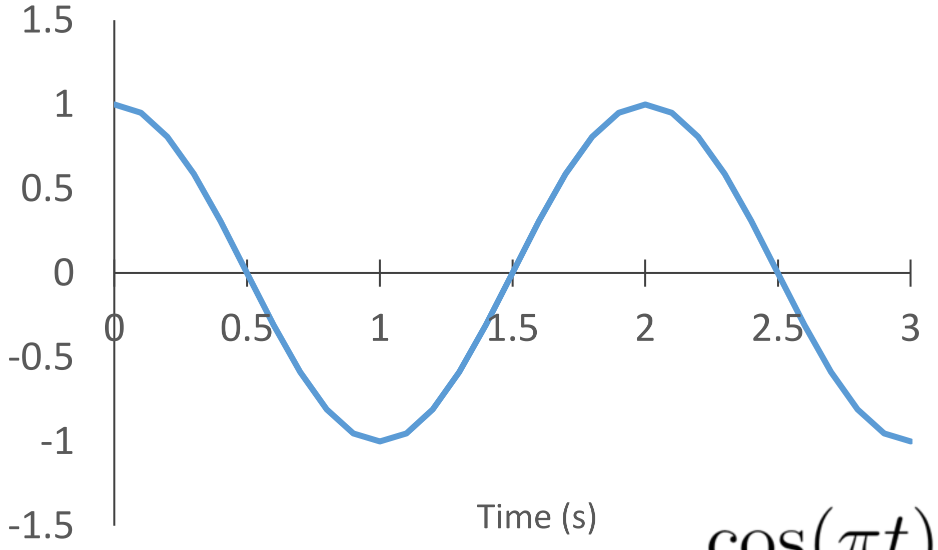

$\cos(\pi t)$

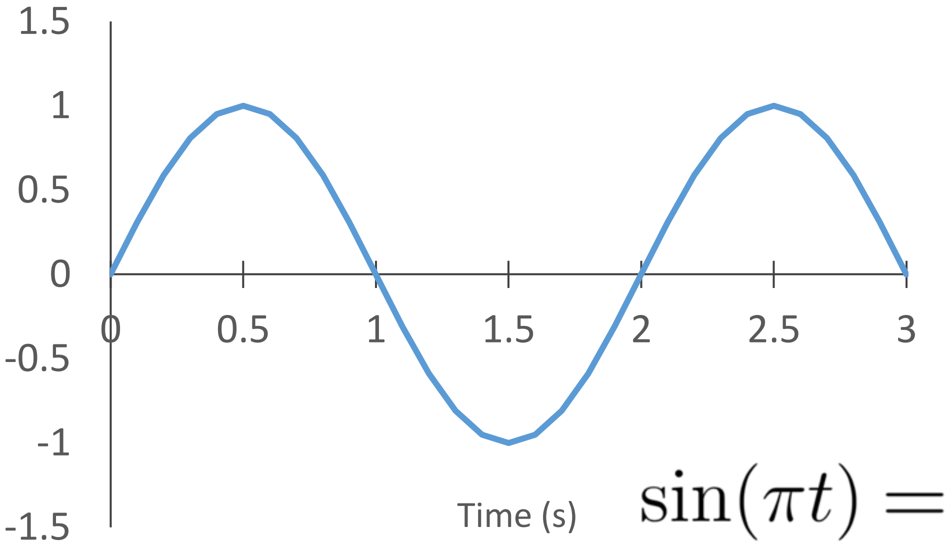

(-90°)

$$\sin(\pi t) = \cos(\pi t \underbrace{- \frac{\pi}{2}})$$

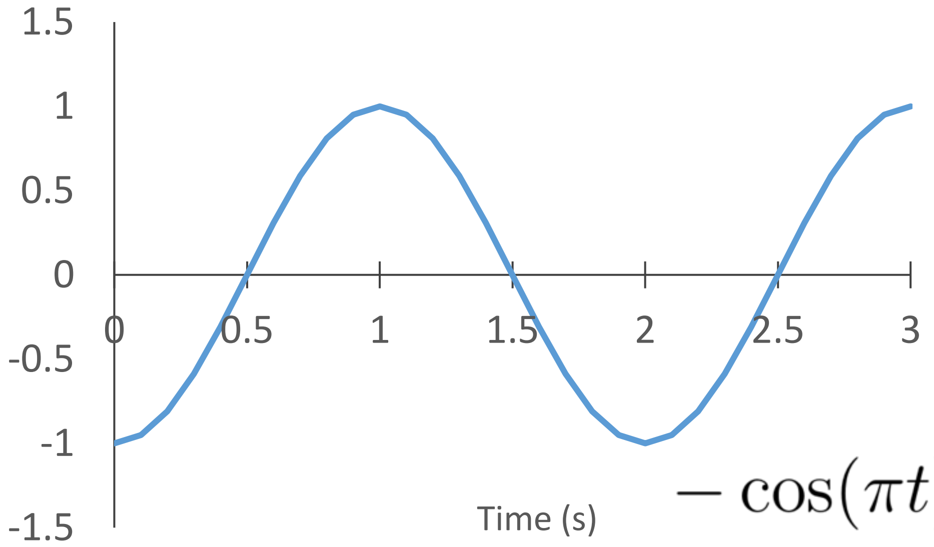

(-180°)

$$-\cos(\pi t) = \cos(\pi t \underbrace{- \pi})$$

- Same frequency
- Same amplitude
- Different phase

Phase in Analog signals

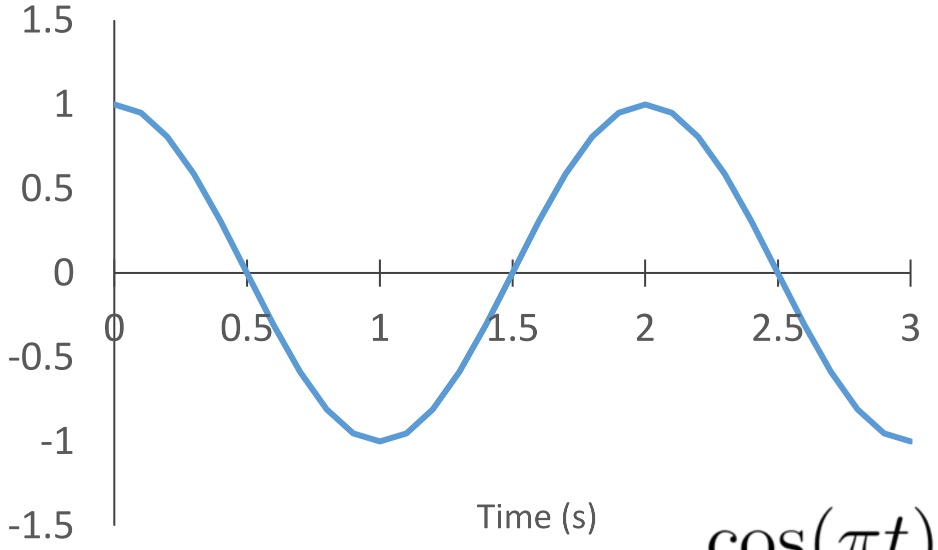

$$\cos(\pi t)$$

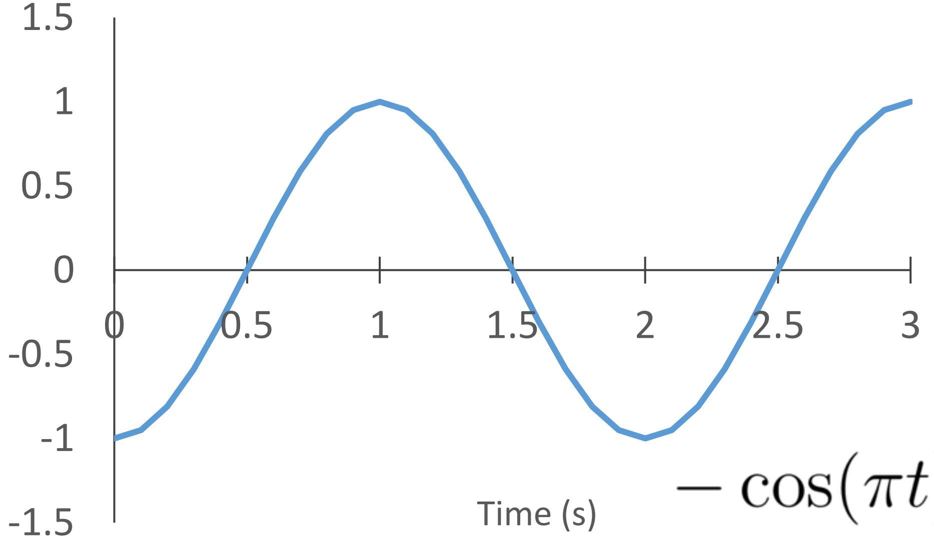

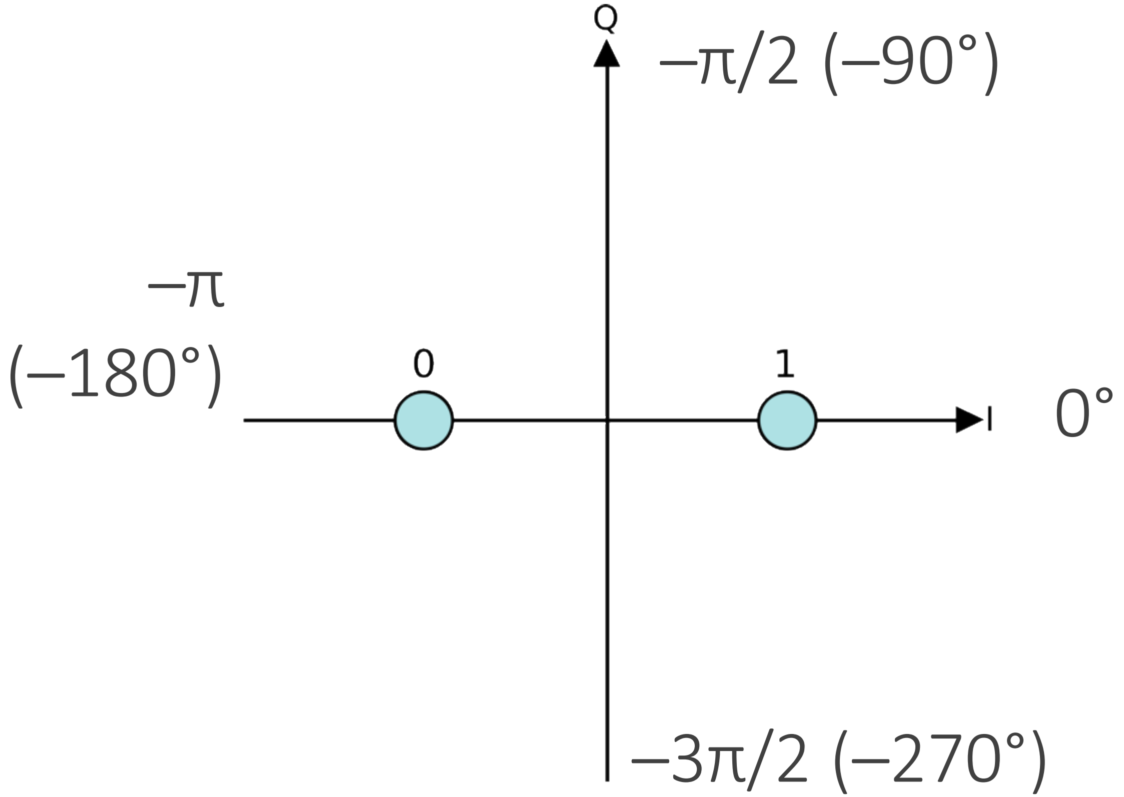

(-180°)

$$-\cos(\pi t) = \cos(\pi t - \pi)$$

- Same frequency
- Same amplitude
- Different phase

I/Q constellations

For the same frequency:

- What I/Q constellation (amplitude, phase) should I select?
- How should I assign a symbol (amplitude, phase) = to bits?

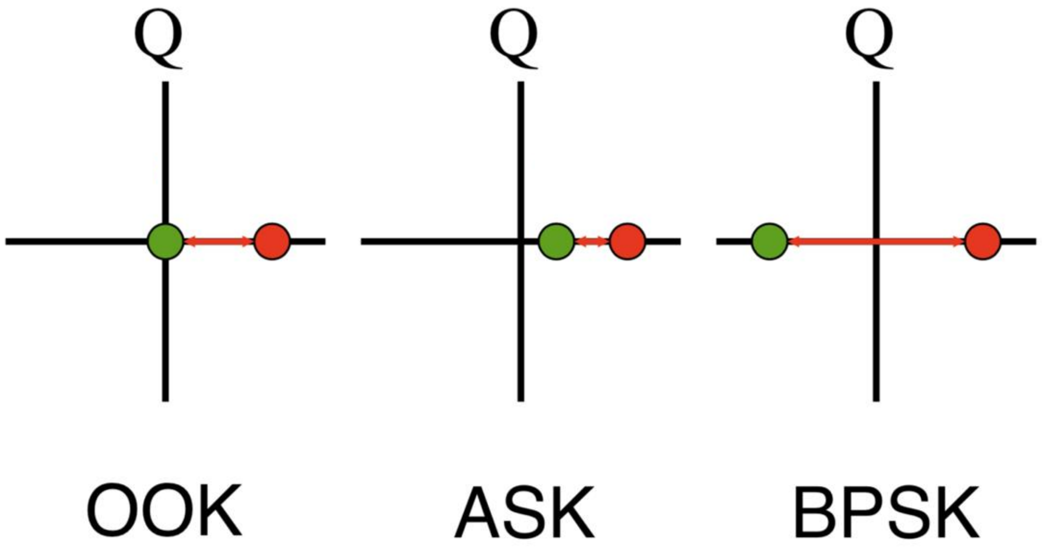

Quadrature Phase Shift Keying (QPSK)

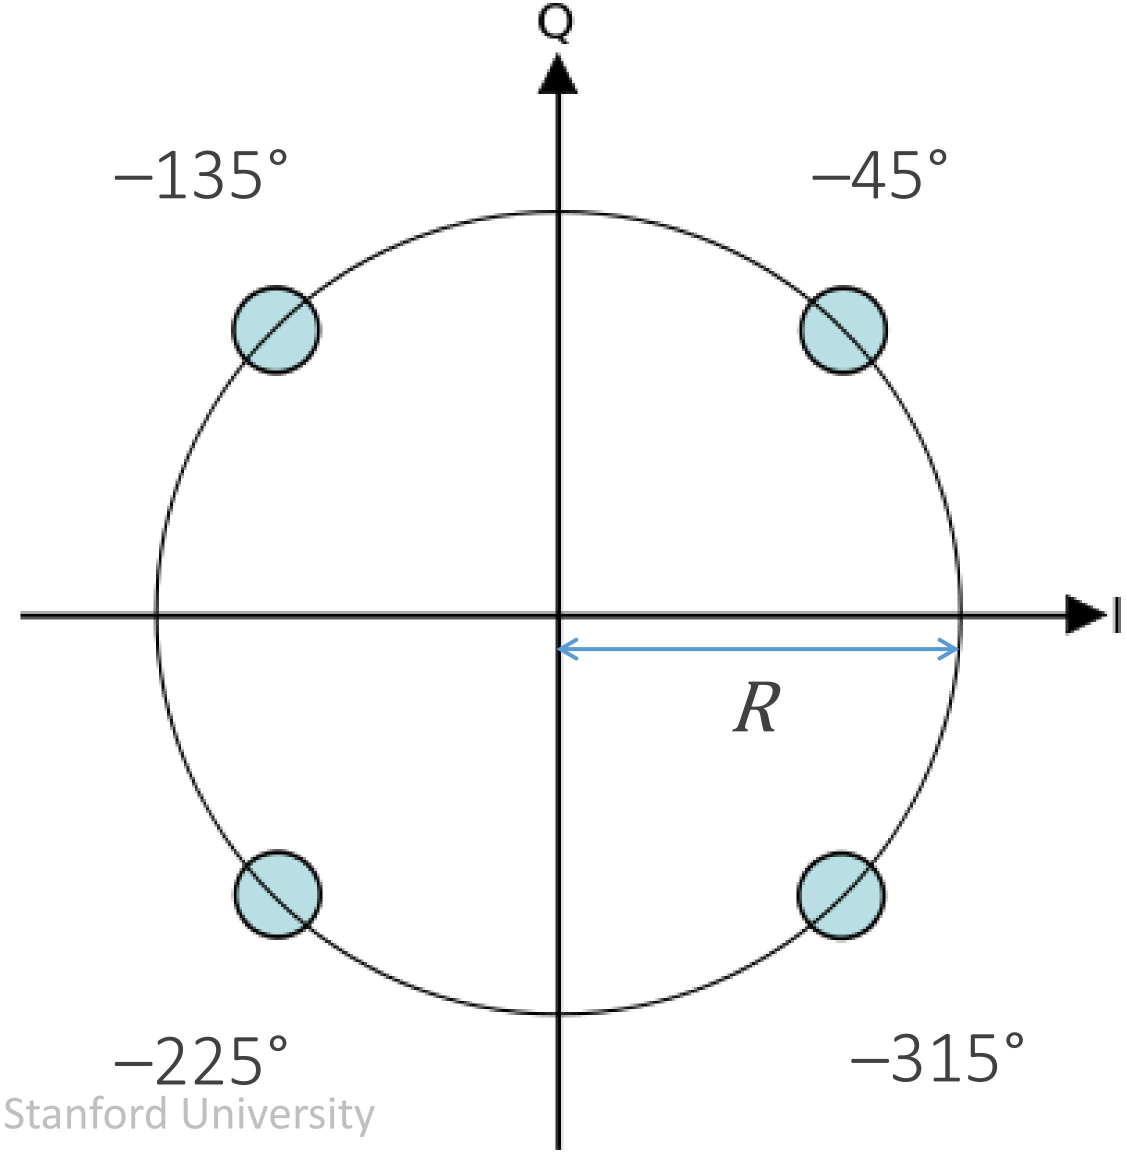

1. For each symbol:
   - What is the amplitude?
   - What is the phase?
2. Represent each symbol as a bit (or bits).

Quadrature Amplitude Modulation (16-QAM)

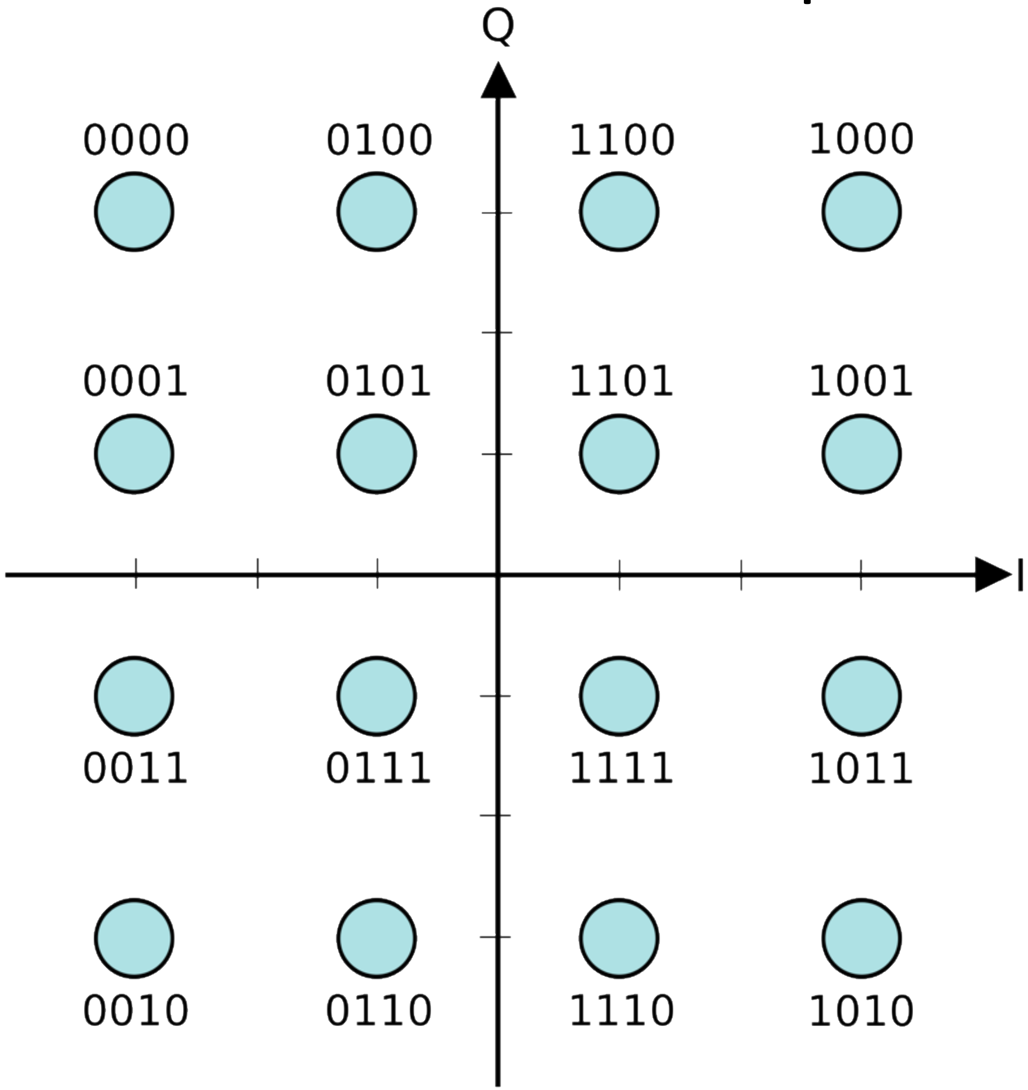

1. How many symbols?
2. How many amplitude variations?
3. How many phase variations?
4. How many bits per symbol?

Example 32 bit word transmission using 16-QAM

1100 1001 0100 1110 1100 0110 1100 1111
0 1 2 3 4 5 6 7

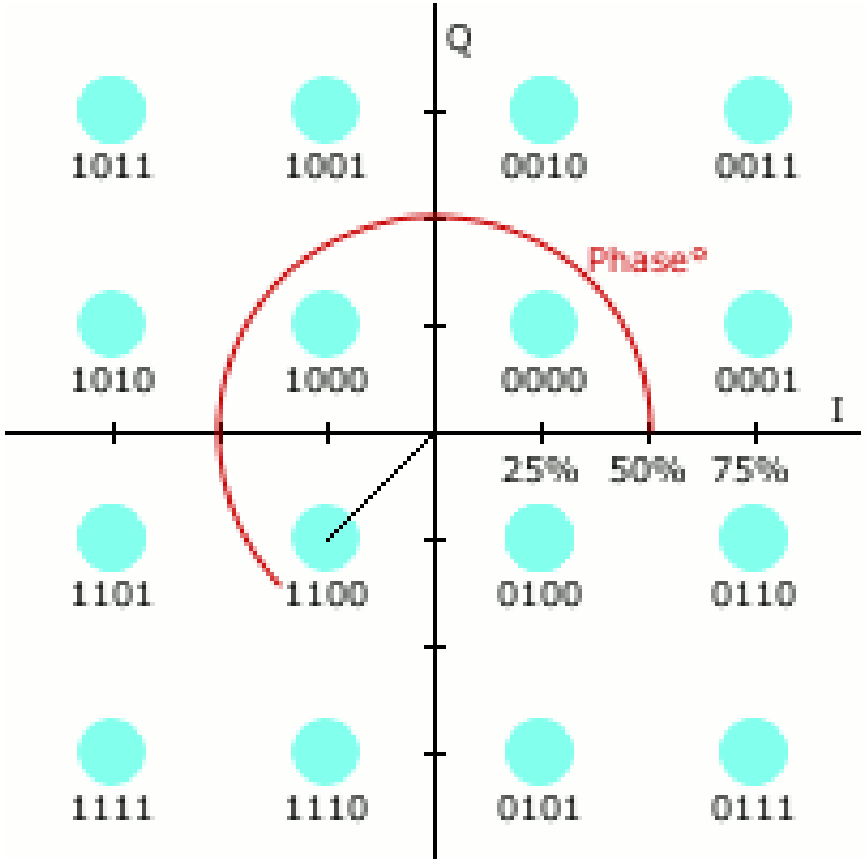

Amp	Phase	Data
25%	225°	1100

Examples today

ASK/OOK: Wired Ethernet
FSK: Bluetooth
BPSK: 802.11 abgn
QPSK: 802.11 abgn, LTE
16-QAM: 802.11abgn, LTE
64-QAM: 802.11 abgn, LTE, 5G
256-QAM: 5G
1024-QAM: Home powerline communication
32768-QAM: ADSL (digital data over long telephone cables)

Clocks

If we don't know the sender's (TX) clock

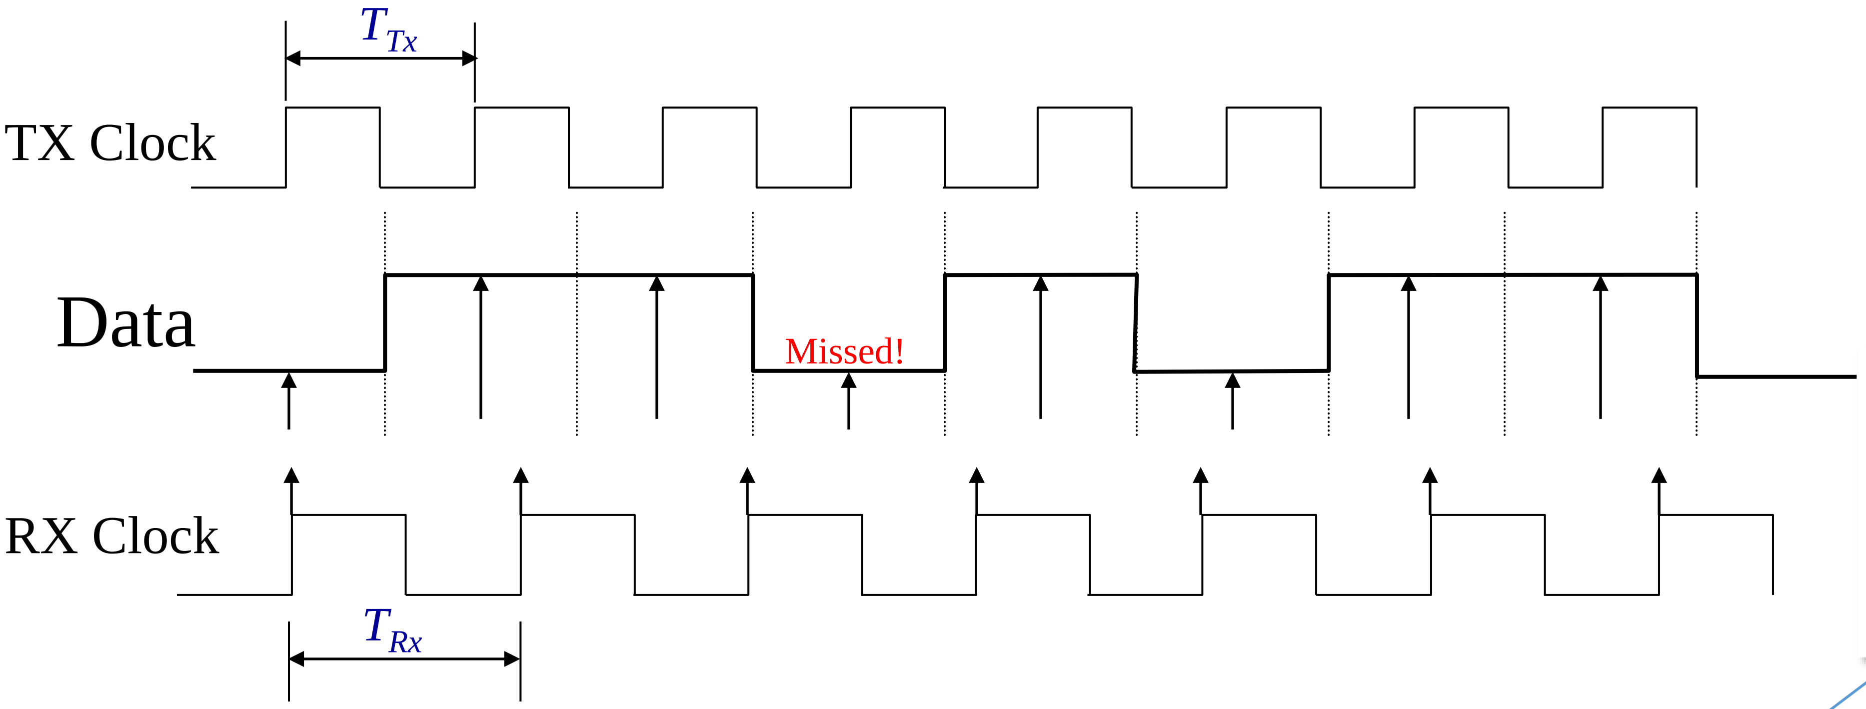

Usually expressed in parts per million (ppm) rather than percentage

If the RX clock is p% slower than the TX clock, then: $T_{RX} = T_{TX}\left(1 + \frac{p}{100}\right)$

After $\frac{0.5}{10^{-2}p}$ bit times, the RX clock will miss a bit completely.

Synchronous communication on network links

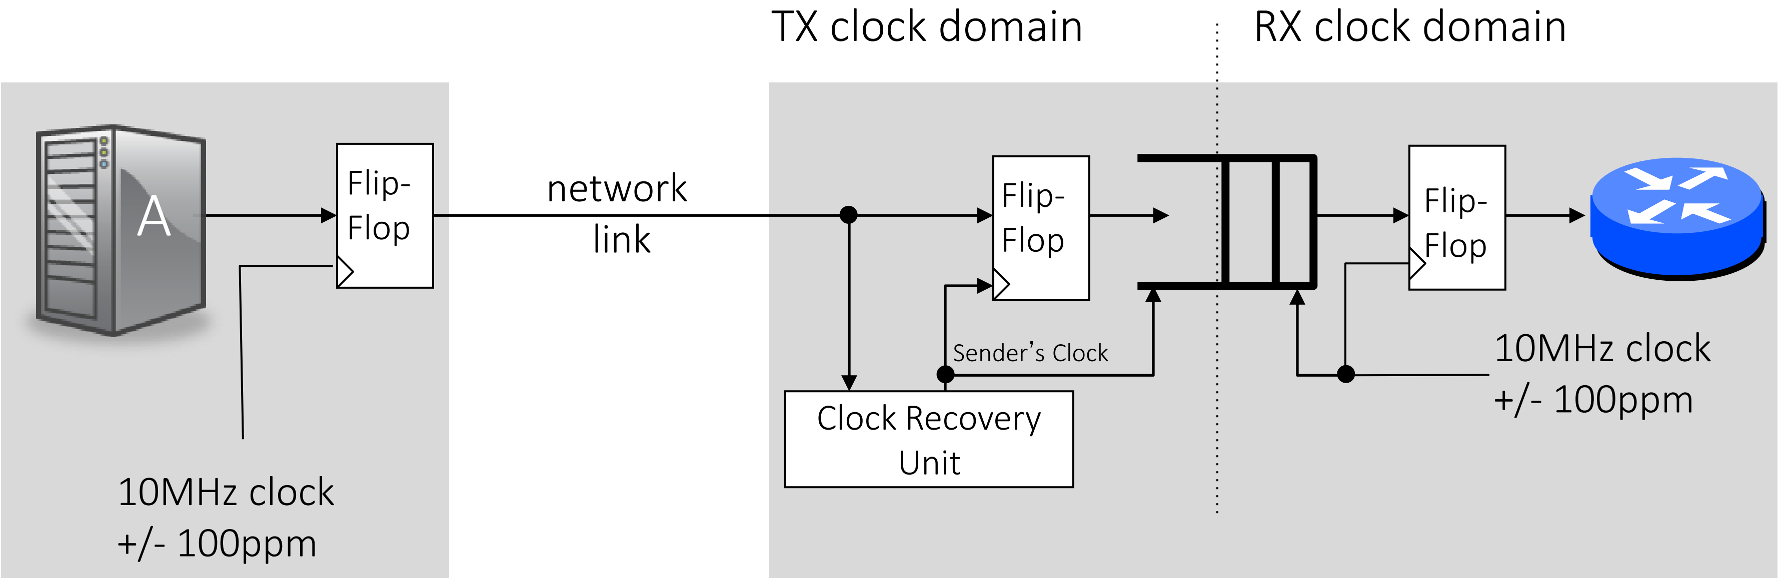

Encoding for clock recovery

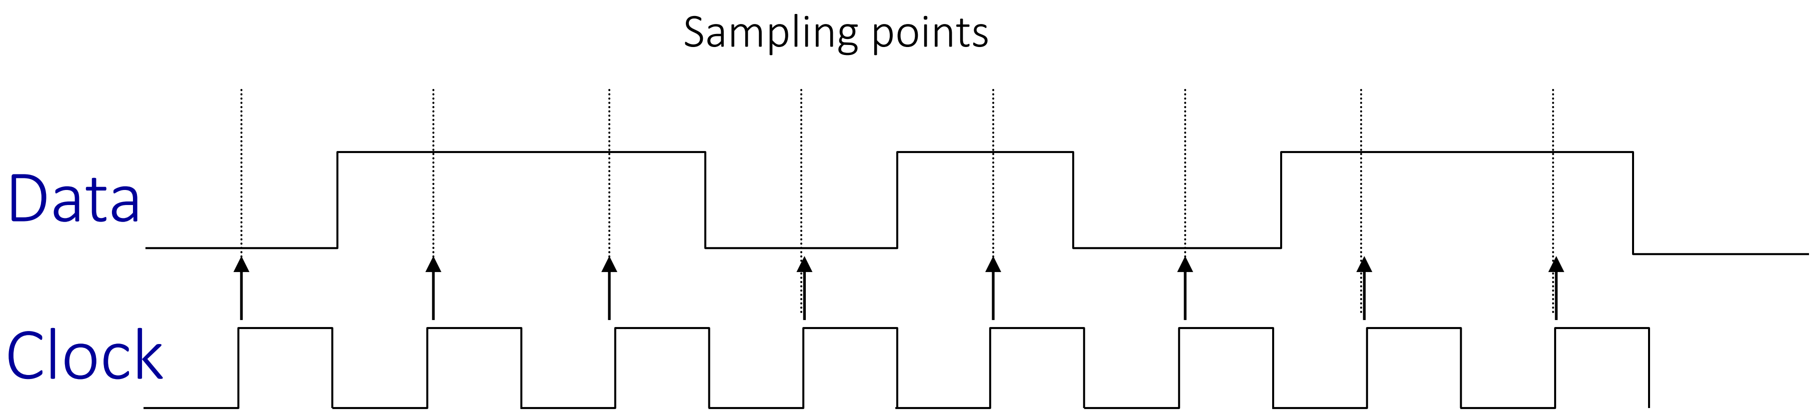

If the clock is not sent separately, the data stream must have enough transitions for the receiver to determine when to sample the arriving data.

Encoding for clock recovery

Example #1: 10Mb/s Ethernet

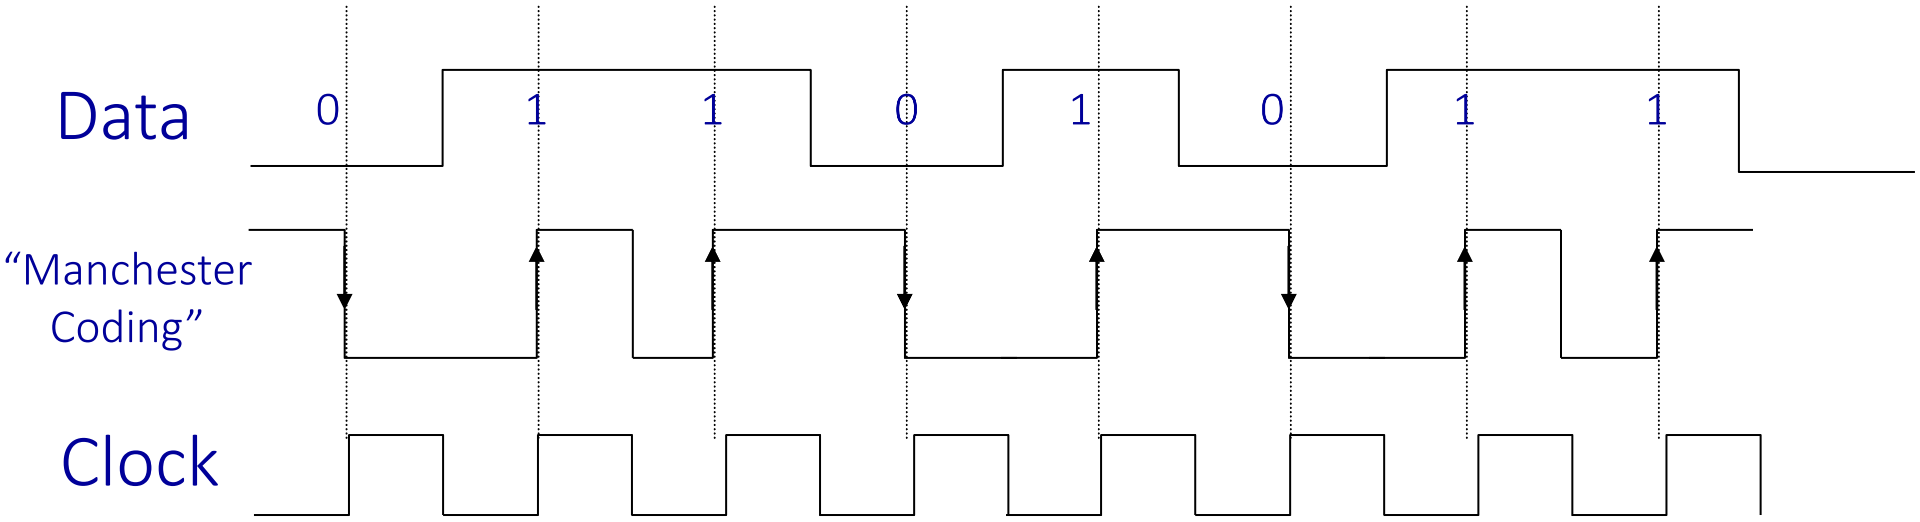

Advantages of Manchester encoding:

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

Disadvantages

- Doubles bandwidth needed in the worst case.

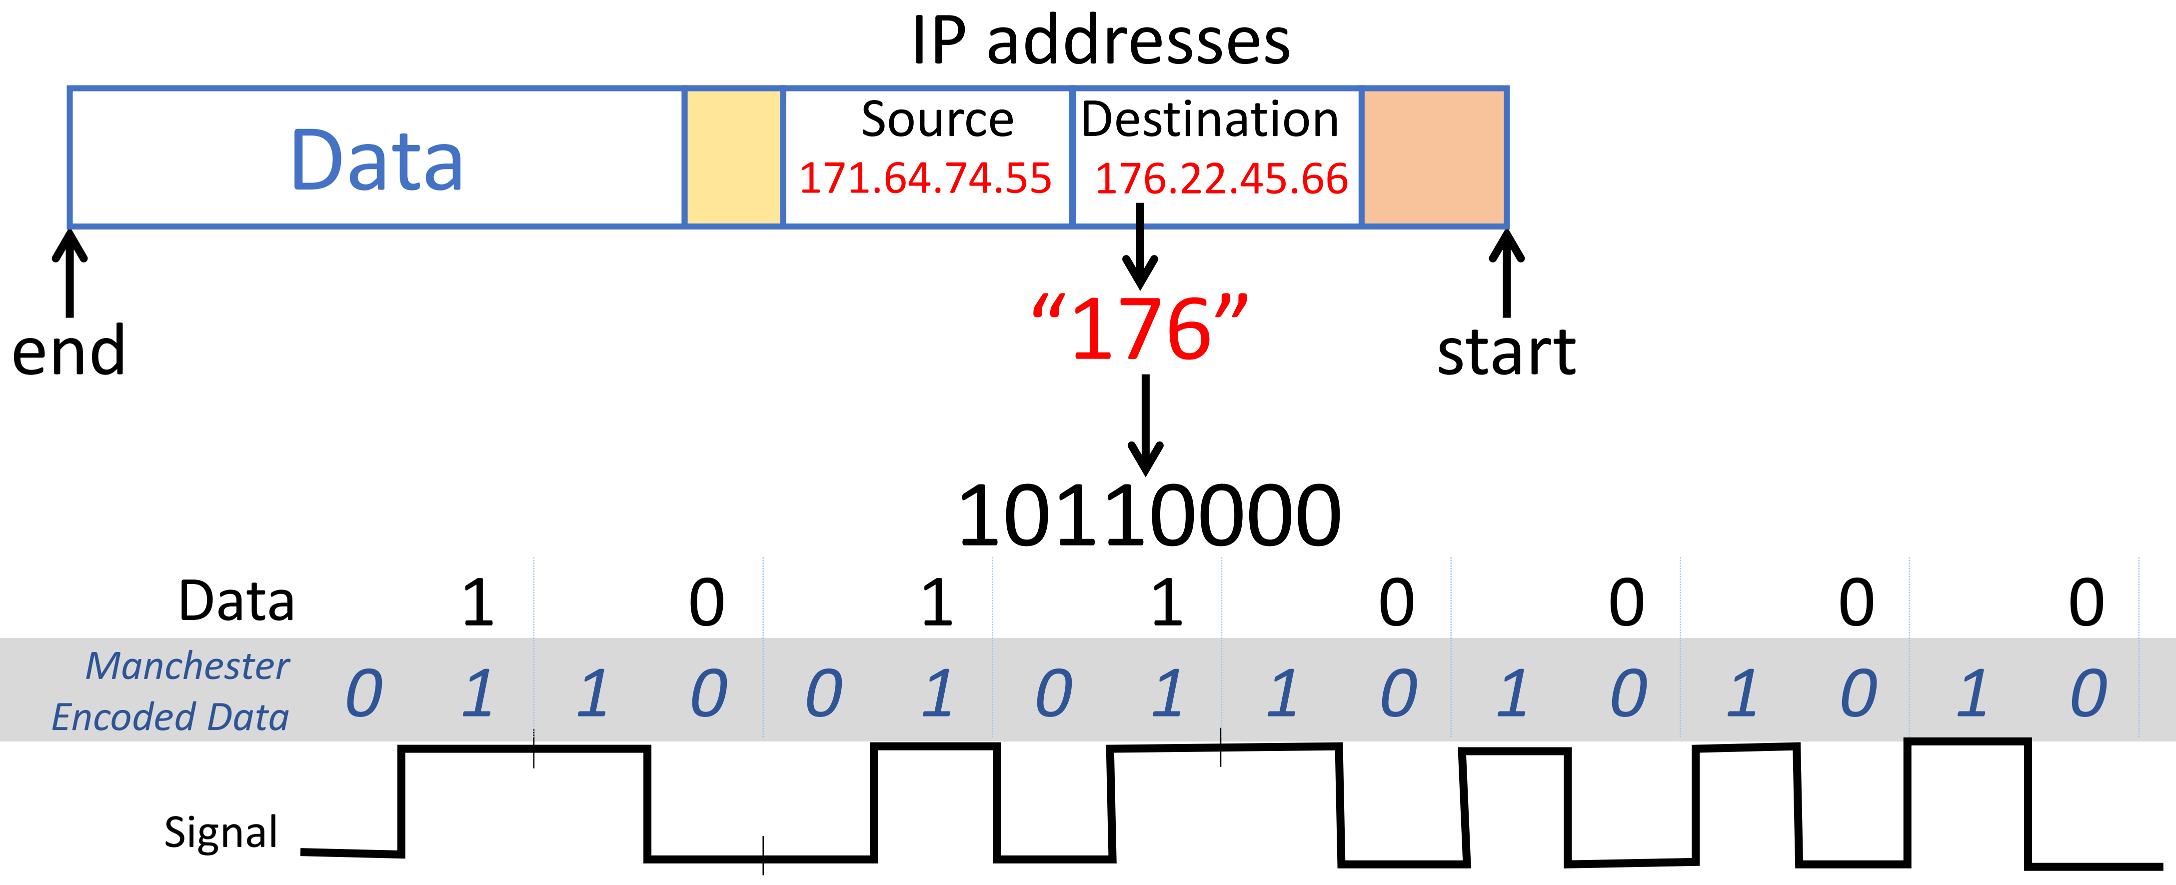
IP addresses
Data
Source
171.64.74.55
Destination
176.22.45.66
end
start
"176"
10110000
Data
1 0 1 1 0 0 0 0
Manchester Encoded Data
0 1 1 0 0 1 0 1 1 0 1 0 1 0 1 0
Signal

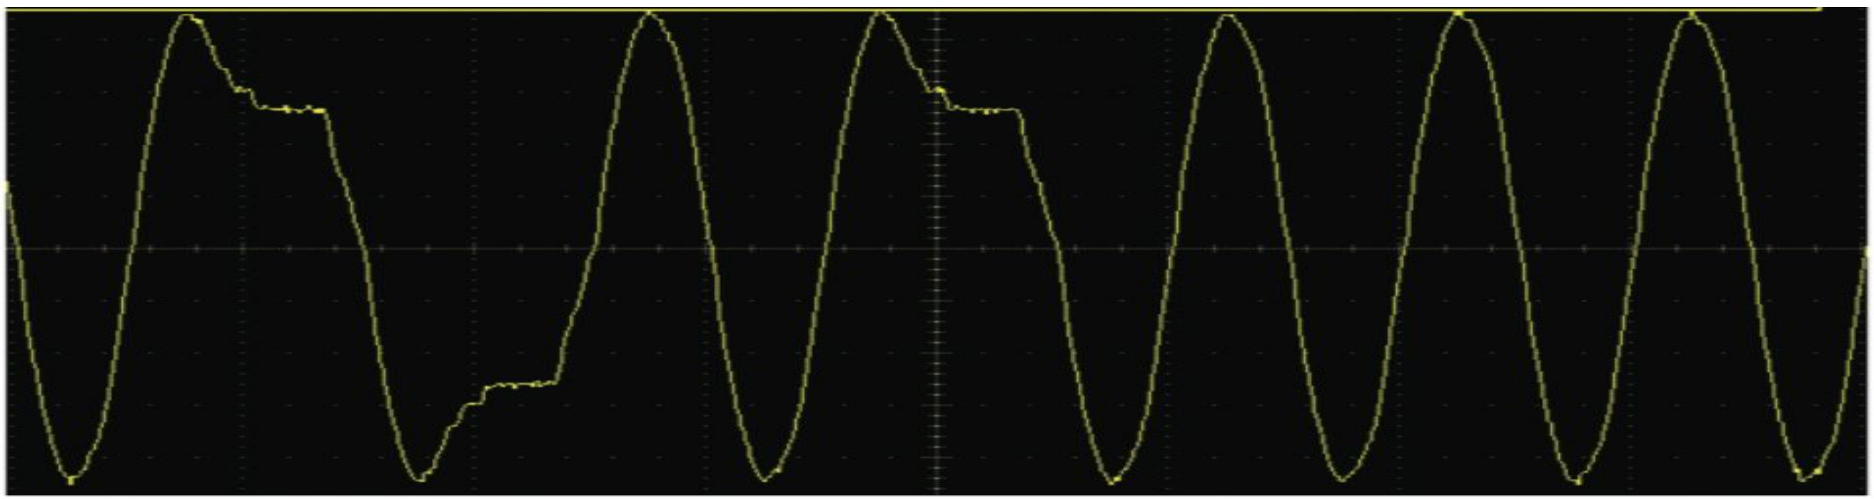

Encoding for clock recovery

Example #2: 4b/5b encoding

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 makes clock recovery a little harder.

Elasticity Buffer

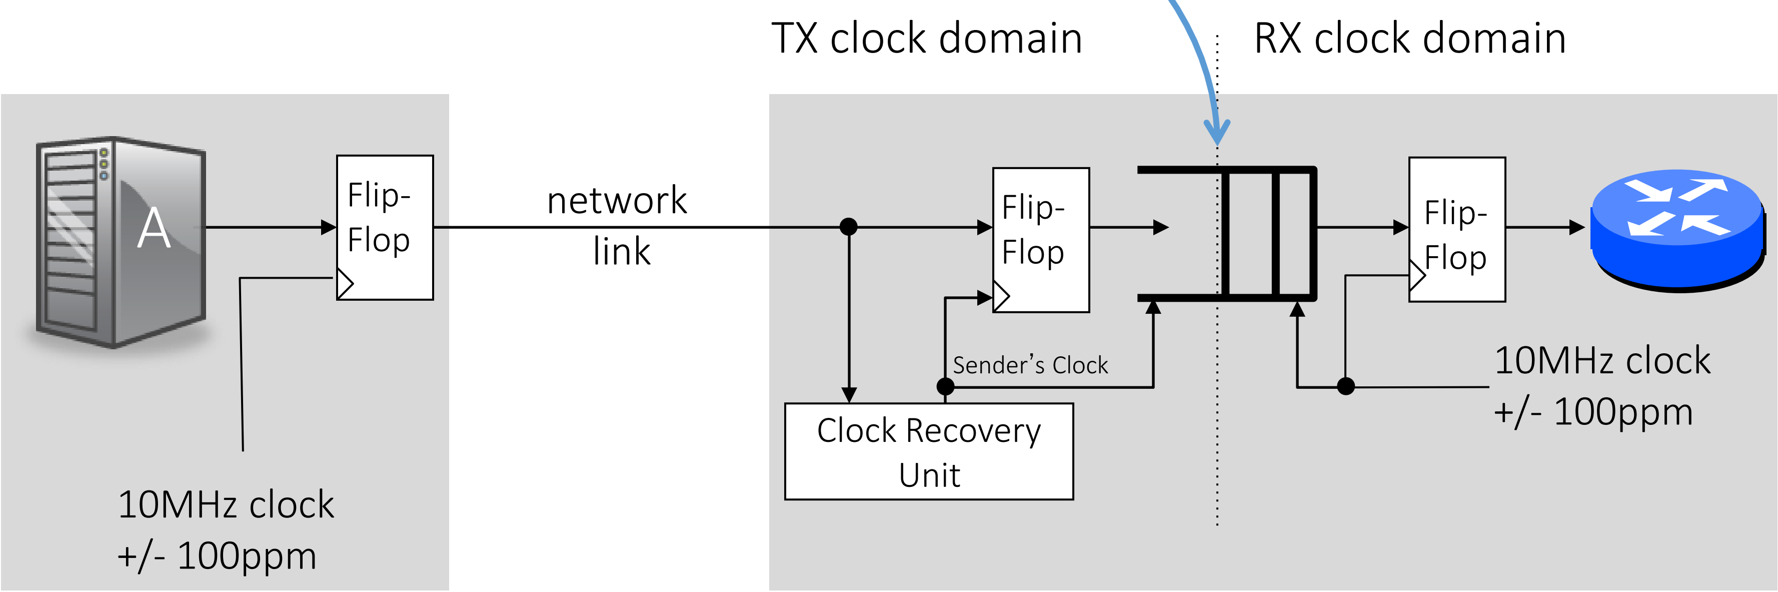

Sizing an elasticity buffer

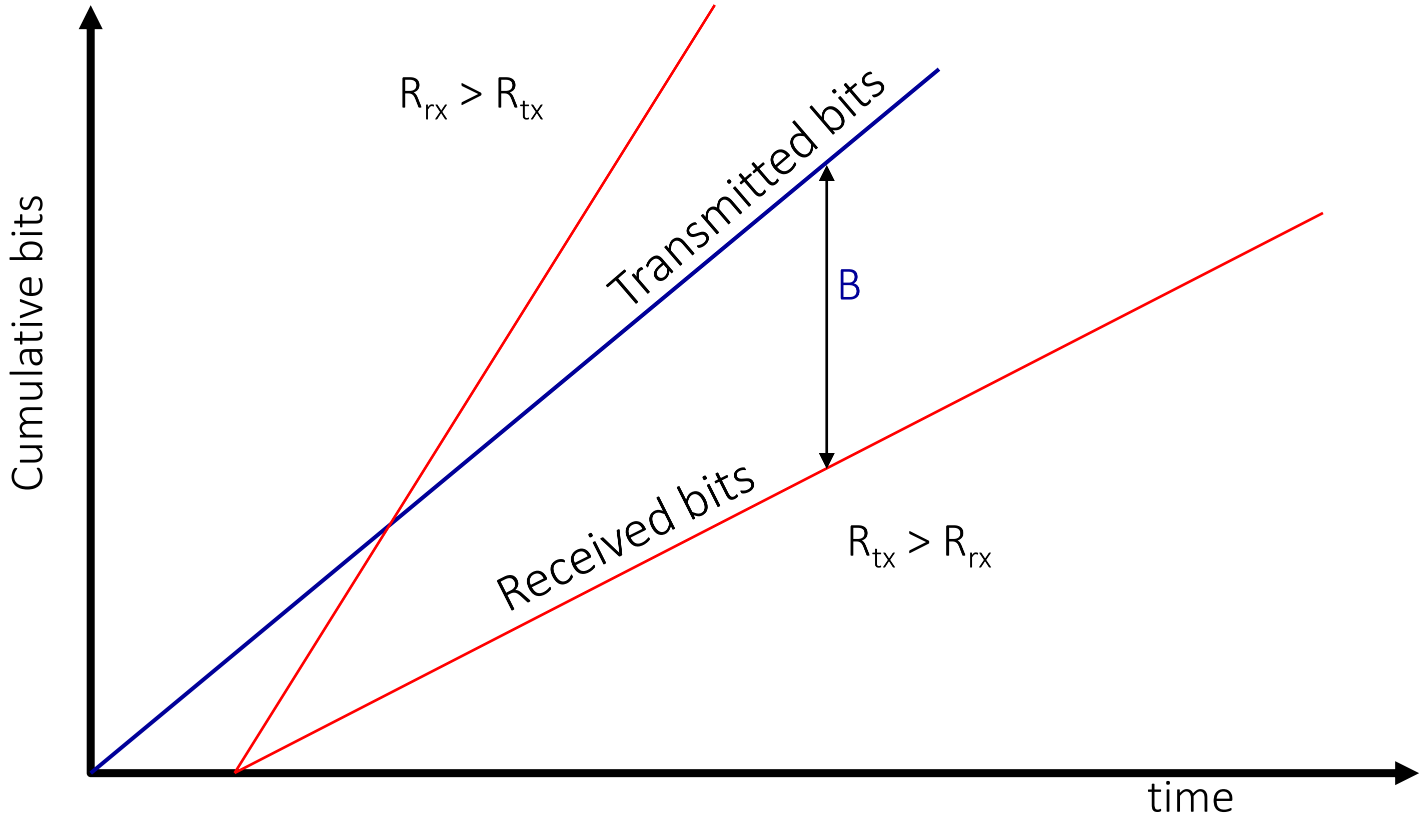

Sizing an elasticity buffer

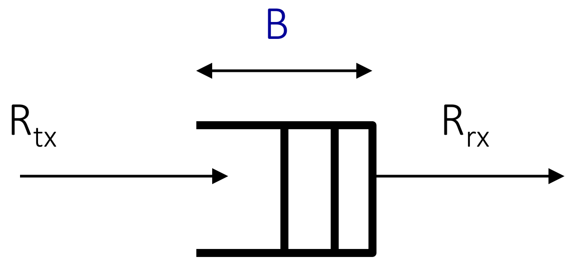

1. Hold buffer nominally at B/2.
   - At start of new packet, allow buffer to fill to B/2.
   - Or, make sure buffer drains to B/2 before new packet.
2. Size buffer so that it does not overflow or underflow before packet completes.
3. ($R_{tx} > R_{rx}$): Given inter packet gap, size B/2 for no overflow.
4. ($R_{rx} > R_{tx}$): Given max length packet, pick B/2 for no underflow.

Preventing overflow

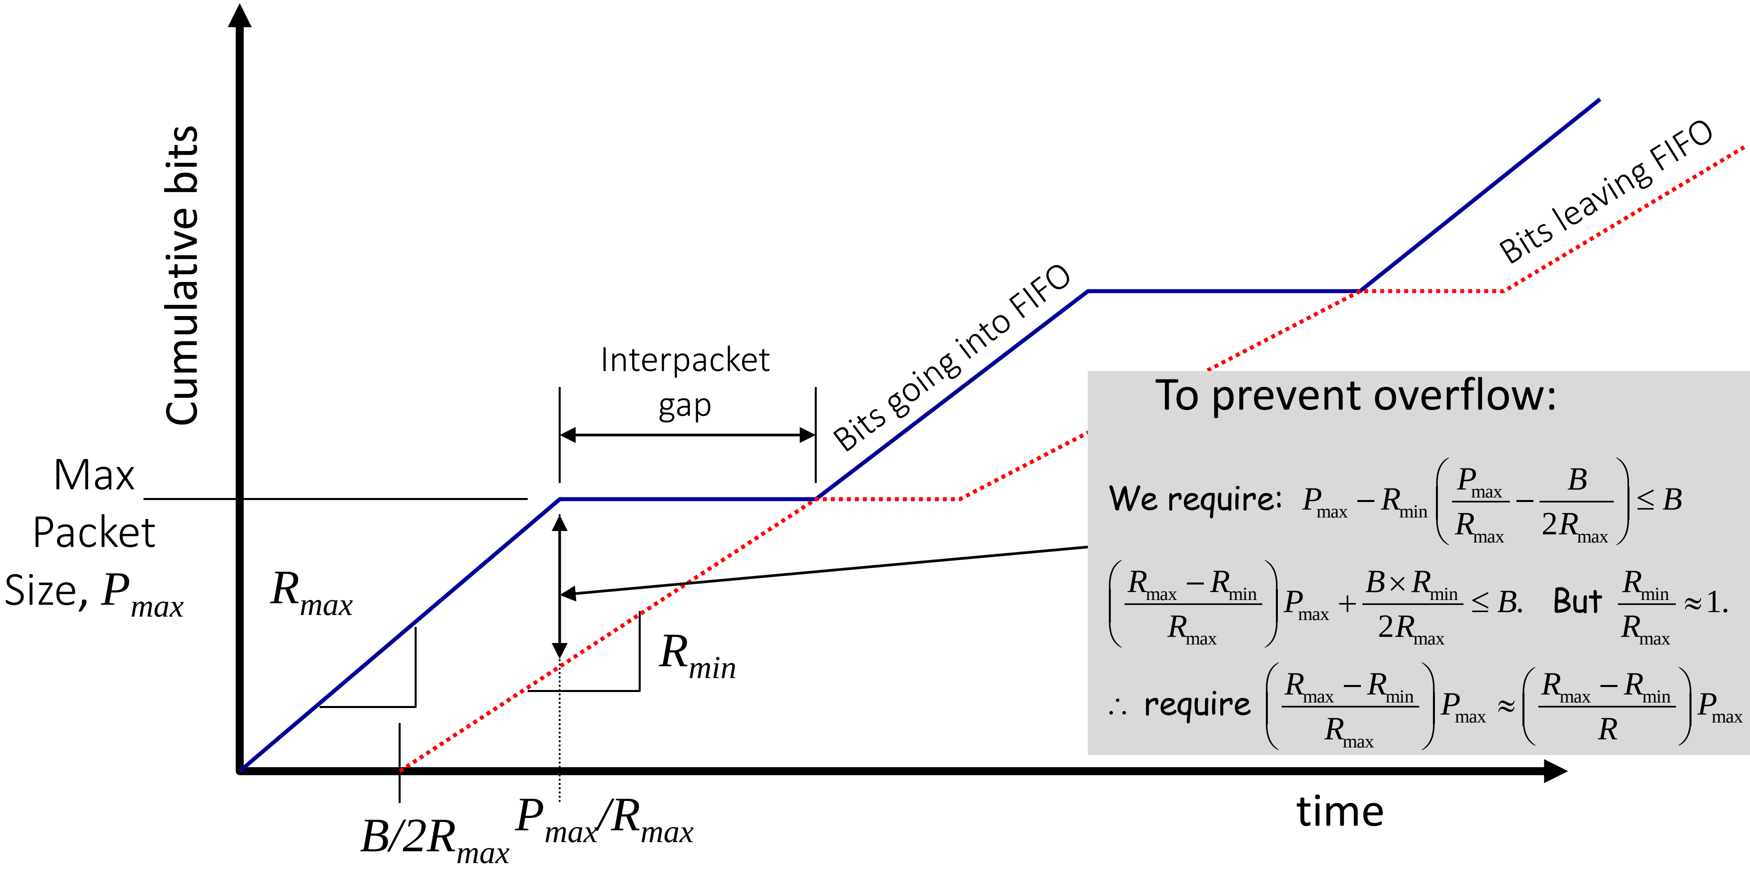

To prevent overflow:

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

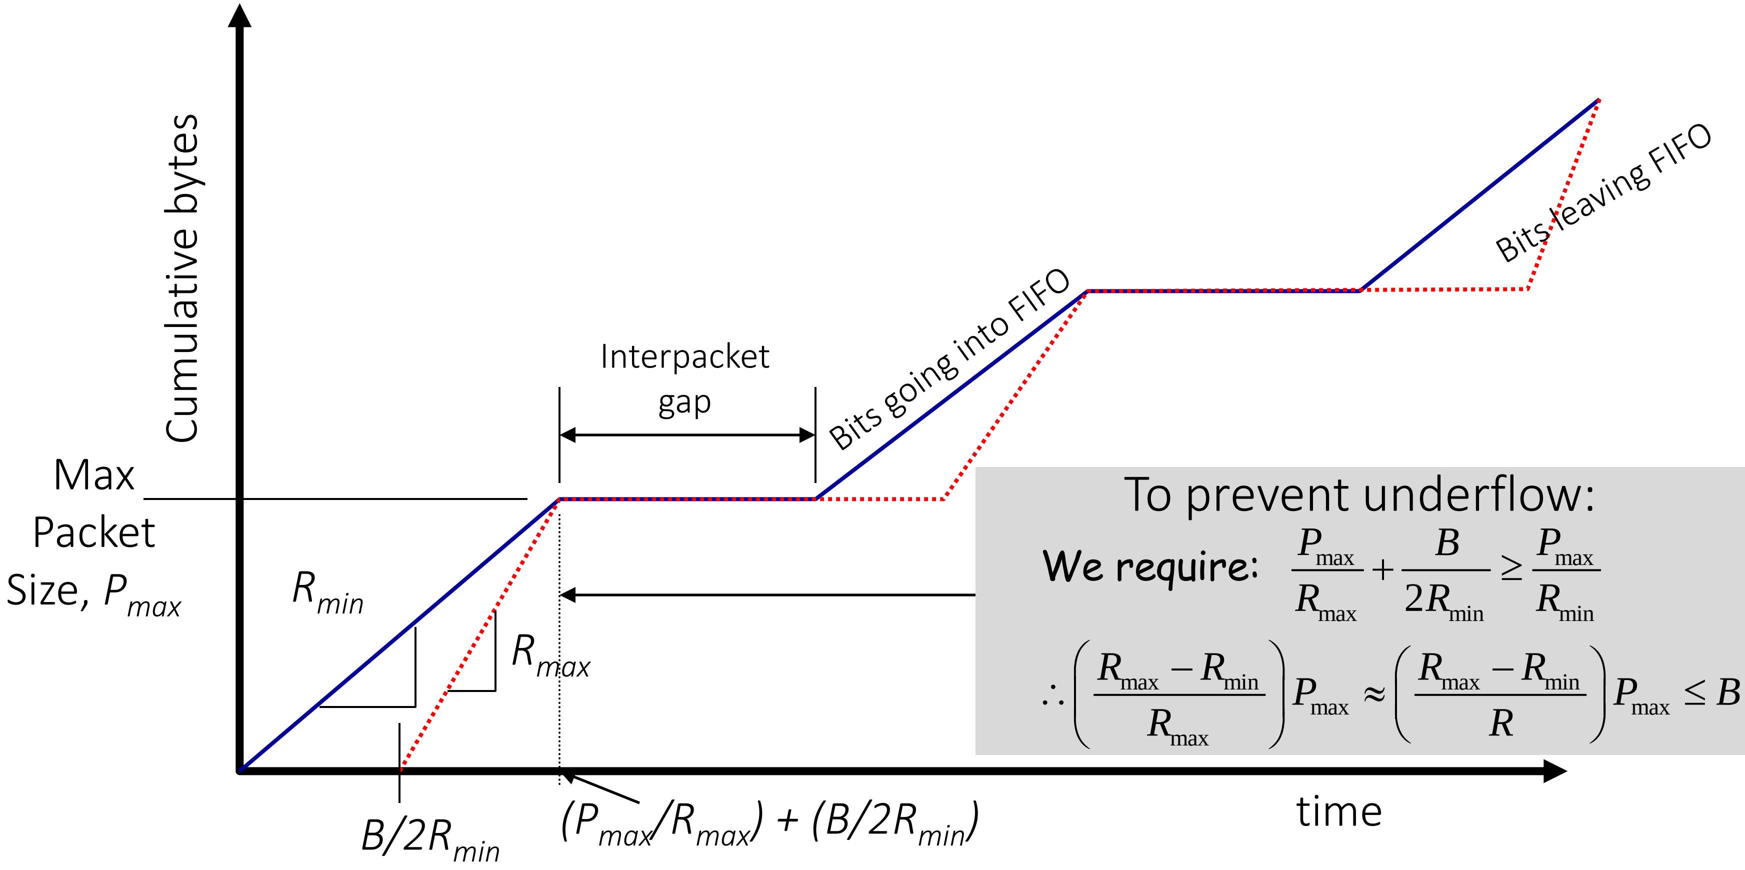

To prevent underflow:

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

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

Sizing an elasticity buffer

Example

Maximum packet size 4500bytes

Clock tolerance +/- 100ppm

$$\left(\frac{R_{max} - R_{min}}{R}\right) = 200 \times 10^{-6}$$

$$\therefore B \geq 2(4500 \times 8 \times 200 \times 10^{-6}) = 14 bits$$

Therefore,

1. Elasticity buffer needs to be at least 14 bits
2. Wait for at least 7 bits before draining buffer
3. Inter-packet gap at least 7 bits