Data Communication Basics

Mahalingam Ramkumar
Mississippi State University, MS

February 4, 2015

Fundamental Factors

- Channel bandwidth H
- Signal power S
- Noise power N

Channel Capacity R

- Achievable bit-rate R
- Shannon's Capacity Limit

$$R = H \log_2(1 + S/N) \text{ bps.} \tag{1}$$

S/N in Decibels (dB)

- Signal to Noise Ratio (SNR)
- Usually expressed in dB
- $(S/N)_{dB} = 10 log_{10}(S/N)$
- 20 dB $\rightarrow S/N = 100$;
- 30 dB $\rightarrow S/N = 1000$;

Example

- What is the limit on achievable rate R in a channel with bandwidth 1 MHz and SNR 30 dB?
- 30dB $\rightarrow S/N = 1000$

$$R = 1 \times 10^6 \times \log_2(1 + 1000) \approx 1 \times 10^6 \times 10 = 10Mbps \quad (2)$$

- To double the rate to 20 Mbps, we can
  - Double the bandwidth, or
  - increase SNR to 60 dB (or increase $S/N \approx 1$ *million*)
  - $\log_2(1,000,000) \approx 20$

Nyquist Sampling Theorem

- Any signal "emerging" from a medium is band-width limited
- Bandwidth - H hertz (cycles/sec)
- A band-limited signal (limited to H Hz) can be uniquely reconstructed from discrete samples taken at a rate $2H$ per second

Nyquist Capacity

- Assume each sample can take V possible values
- $\log_2(V)$ bits per sample
- Nyquist capacity $C = 2H \log_2(V)$
- No of distinct recognizable values V is determined by noise.
- If noise is zero, $V \to \infty$, so $C \to \infty$
- Every medium has inherent noise
- Even at absolute zero (0 Kelvin).

Transmission Media

- Two basic types
- Guided
  - Twisted pair
  - Coaxial cables
  - Fiber optics
  - Magnetic media
- Unguided

Twisted Pair

- Telephone cables
- Any wire is an antenna!
- Twisting substantially reduces interference
- Waves from different twists cancel out
- Number of twists per cm
- Categories 3 (16MHz), 5 (100 MHz), 6 (250), and 7(600)
- Ethernet cables

(a) (b)

Coaxial Cables

- Up to 1 GHz
- Twisted pair cables are catching up!
- Better noise immunity than twisted pairs
- Used for TV signals, cable TV ...

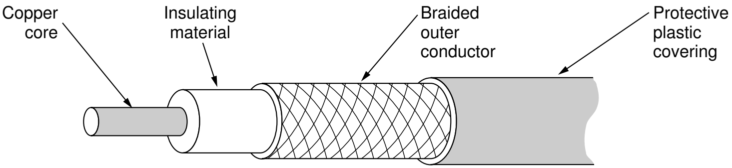

Fiber Optics

- Extremely high bandwidth
- Single mode and multimode
- Single mode fibers can handle upto 50GHz over 100km!

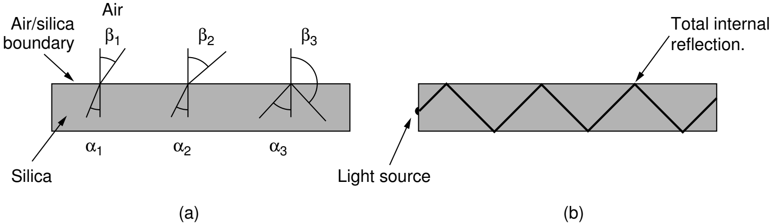

Fiber vs Copper Wire

- Weight vs bandwidth
- Cost
- Security

Electromagnetic Spectrum

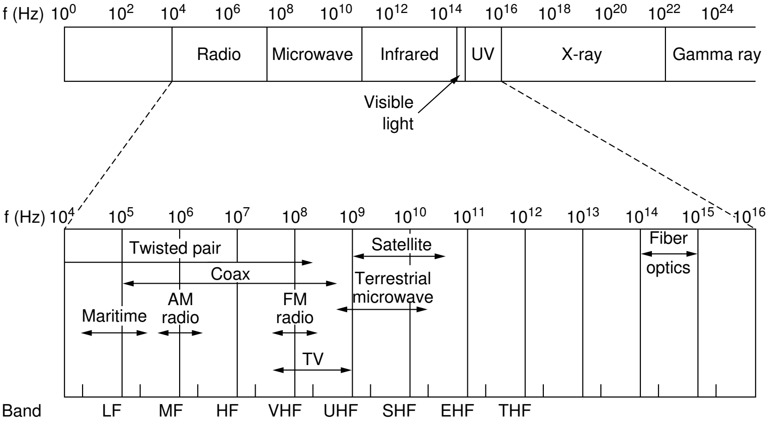

Allocation of Spectrum

ITU-R, FCC

ISM (Industrial, Scientific, Medical) Bands

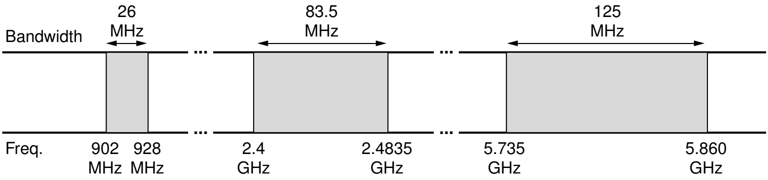

Satellites

- Geo-stationary (about 36,000 km)
- Medium Earth Orbit (18,000 km, 6 hr orbits)
- Low Earth Orbit (less than 1000 km)
- Satellite vs Fiber
    - Remote / hostile areas
    - Point-to-point vs Broadcast
    - Mobile communications
    - "Right of way" for laying cables

Thus far ...

- Capacity is a function of signal power, noise power and bandwidth of the channel/medium
- Shannon's and Nyquist capacity formulations
- $C = H \log_2(1 + S/N)$ bps (Shannon)
- $C = 2H \log_2 V$ bps (Nyquist): $2H$ samples/sec. Each sample can have one of V different *levels* (number of levels depends upon signal to noise strength at receiver)
- Types of channels and their pros/cons

Modulation

- Modulation: *information* signal + *carrier* signal → *modulated* signal
  - Speech signals have frequencies between 300 and 3000 Hz
  - Signals cannot propagate directly over *all* mediums
  - Speech signals $s(t)$ can be sent directly over a copper cable: but not over air or free-space
  - "Carrier" signals are used to carry the "information" signals
- Modulated signal transmitted over a medium
- Demodulation: extracting information signal from the modulated signal
- Types of modulation: Amplitude (AM), Phase (PM), Frequency (FM)

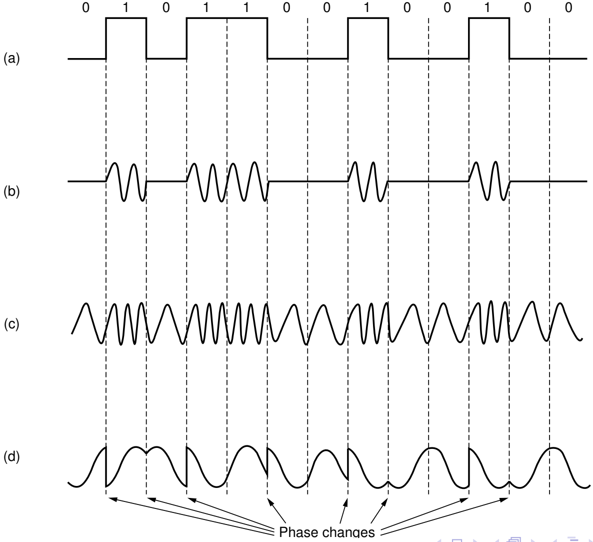
0 1 0 1 1 0 0 1 0 0 1 0 0
(a)
(b)
(c)
(d)
Phase changes

Bandwidth of Modulated Signals

- The same as the bandwidth of the “information” signal
- The bandwidth is now around a “center” frequency — the frequency of the “carrier”
- Many signals can coexist simultaneously in the medium.
- By using different carrier frequencies

Modulation for FAX/Dial-up Internet

- Send bits over telephone cable using Modem (Modulator - Demodulator)
- Cable meant for carrying voice signals between 300-3000 Hz
- Amplitude + phase modulation - series of sinusoidal pulses
- Modulator: Bits are chunked, represented using sinusoidal pulses
- Demodulator: Convert received sinusoidal pulse to bits
- Clipped sinusoid is the basic signal
- Variations achieved by modifying amplitude and phase
- Represented on a *constellation* diagram

QPSK, QAM-16, QAM-64

Quadrature Phase Shift keying (2 bits per sample), Quadrature Amplitude Modulation

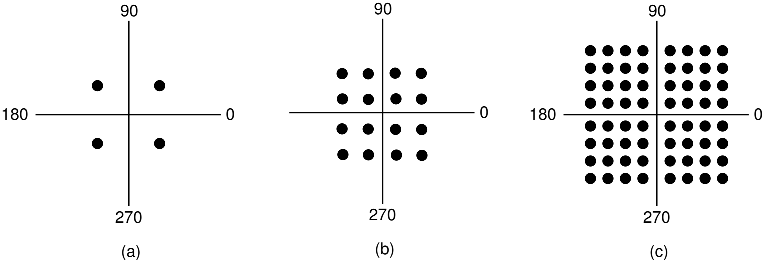

Telephone/FAX Modem

- Each pulse represents multiple bits
- If we use 16 different types of pulses (eg., QAM 16) we can send 4 bits per pulse ($16 = 2^4$)
- Bit rate and Baud rate
  - Baud rate - number of sinusoidal pulses per second
  - Fixed at 2400 for telephone lines
  - Bit rate = Baud rate x number of bits per pulse
  - How do we get 56 K with baud rate of 2400?

Trellis Coded Modulation

Add more bits per sample

Add extra bits to each sample for error correction

V.32 (4+1, 32) - 9600 bps, V.32 bis (6 + 1, 128) 14,400 bps

V.34 - 28,800 bps, V.34 bis - 36,600 bps

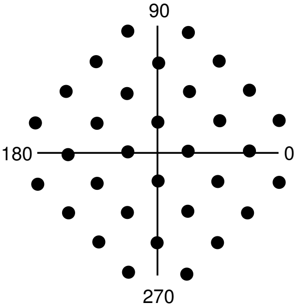

(b)

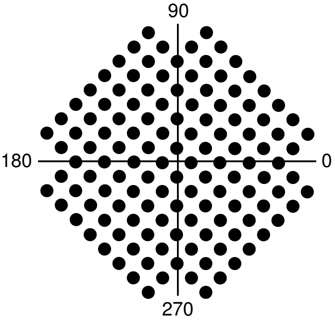

(c)

Modern Modems

- Full duplex - transmissions possible in both directions at the same time
- Half-duplex - Both directions, but not at the same time
- Simplex - Only one direction
- V.90, V.92 - full duplex
- Dedicated uplink and downlink channel
- 56 kbps downlink, 33.6 kbps uplink (V.90)
- 48kbps uplink (V.92) + facility to detect incoming calls while online

Actual Capacity of a Telephone Line

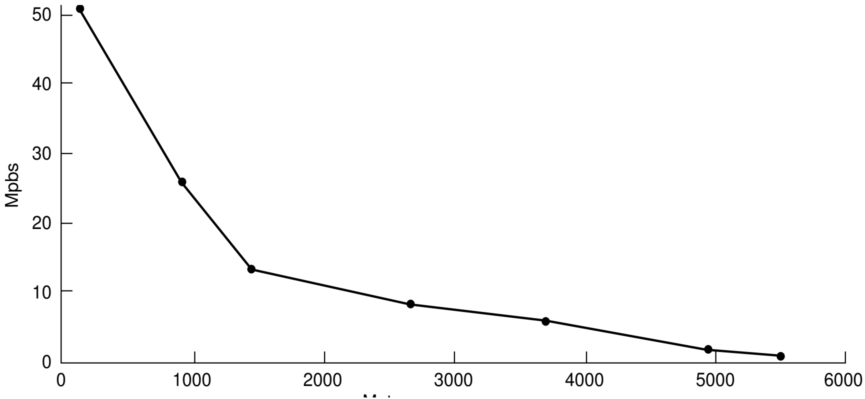

ADSL

- We have been sitting on a gold-mine!
- Telephones filter out higher frequencies to suppress noise
- 1.1 MhZ specturm divided into 256 channels - 4312 Hz each (DMT - Discrete Multitone)
- First channel used for POTS (plain old telephone service)

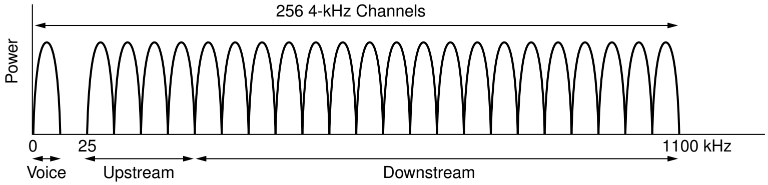

ADSL Equipment

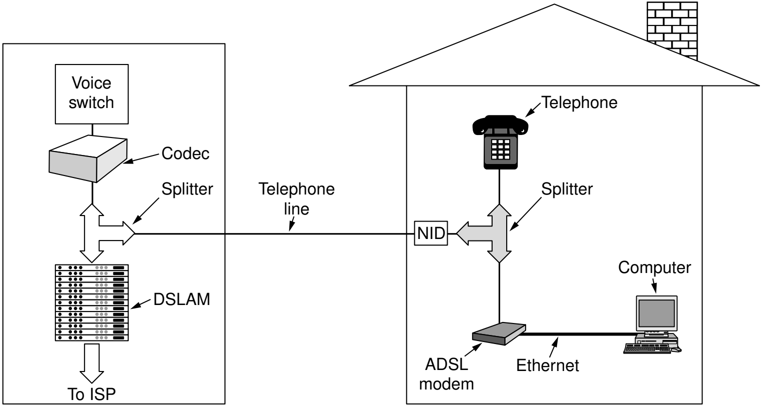

Telephone company end office
Customer premises

Multiplexing and Demultiplexing

- Multiplexing - many to one
  - Many signal mixed together to produce one signal. The multiplexed signal is transmitted over a channel
- Demultiplexing - one to many
  - Multiplexed signal is split into individual components
- Types: FDM, TDM, CDM

Frequency division multiplexing (FDM)

Telephone - 4000Hz per channel (450 + 3100 + 450)

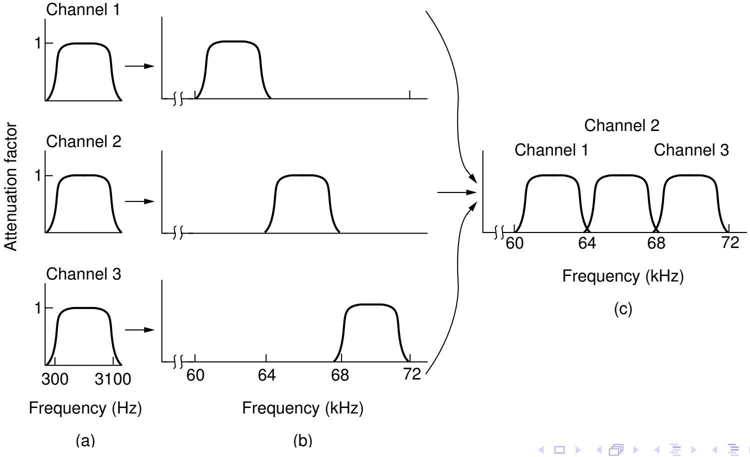

Wavelength Division Multiplexing (WDM)

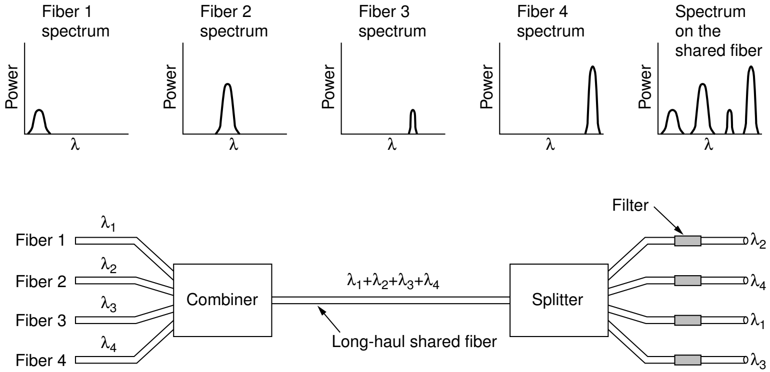

Time Division Multiplexing (TDM)

- *N* signals, each having *M* samples per sec
- Interleaved together in a channel which can carry *NM* samples per sec
- POTS signal (speech or clipped sinusoids) sampled at 8000 samples per sec (sent as 64000 bps - each sample is represented using 8 bits)
- T1 lines - 1.544 Mbps, 24 channels

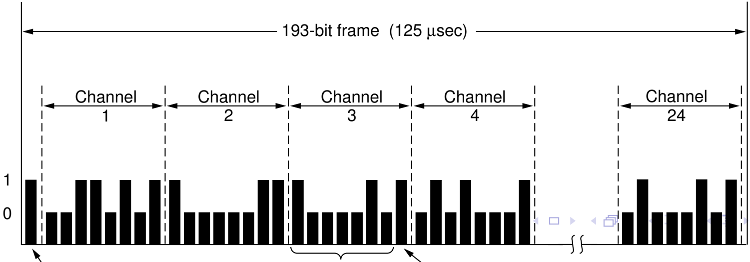

TDM in Telephone Network

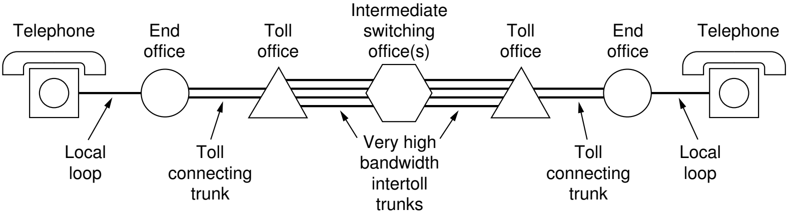

T1, T2, T3, T4

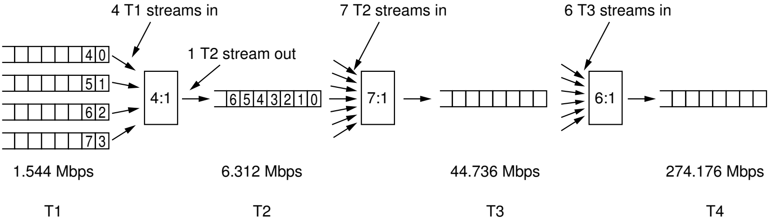
4 T1 streams in
1 T2 stream out
7 T2 streams in
6 T3 streams in
4 0
5 1
6 2
7 3
4:1
6 5 4 3 2 1 0
7:1
6:1
1.544 Mbps
6.312 Mbps
44.736 Mbps
274.176 Mbps
T1
T2
T3
T4

Code Division Multiple Access (CDMA)

- TDM - different channels can overlap in frequencies, but not in time
- FDM - no overlap in frequencies, overlap in time
- CDMA - overlap in both time and frequency
- Separable by orthogonality of *codes*
- Two vectors are orthogonal if their inner-product is zero.

CDMA

- Each bit is converted to a vector of *chips*
- Different users assigned different *orthogonal* chip vectors
- A user assigned vector $\bar{x}$ transmits $\bar{x}$ to send a one or $-\bar{x}$ to send a zero
- $\bar{x}.\bar{x} = 1, \bar{x}.(-\bar{x}) = -1$
- $\bar{x}.\bar{y} = 0$ if $\bar{y}$ is orthogonal to $\bar{x}$.
- In practice difficult to obtain *strict* orthogonality
- $\bar{x}.\bar{y}$ can be made very small compared to 1
- "Tolerable" interference with quasi-orthogonal sequences
- CDMA is also more efficient (closer to Shannons limit)

Why CDMA?

- If n chips per bit, then chip frequency n times higher than bit-frequency
- Energy to sent a bit spread over a larger frequency
- Useful during war-time as transmission virtually indistinguishable for noise in any narrow band.
- CDMA is also more efficient
- Higher frequency and lower S/N takes us closer to Shannons limit
- Capacity increases linearly with bandwidth (only logarithmically with S/N)

Characterizing Data Transmission

- Bit-rate R and bit duration,
- Packet duration τ_p
- Propagation speed c, propagation delay τ_c
- One-way delay
- Round trip time (RTT)

Packet Duration

- Bit duration is the inverse of the bit-rate R
- If bit rate is 1 Mbps, the bit duration is 1 μs
- 1 Kbps $\leftrightarrow$ 1 ms, 1 Gbps $\leftrightarrow$ 1 ns.
- Packet duration is bit-duration times packet size (number of bits in the packet)
- 10 Mbps bit-rate: duration of a 100 byte packet is $\tau_p = 100 \times 8 \times 0.1 = 80\mu s$.
- 1 Gbps bit-rate: duration of a 100 byte packet is $\tau_p = 100 \times 8 = 800ns = 0.8\mu s$.

One-way Delay

- propagation at the speed c (light speed in the medium)
- $c = 3 \times 10^8$ m/s in free-space and air, $c = 2.5 \times 10^8$ m/s in copper.
- Distance traveled in 1 μs is 250 m (25cm in 1 ns)
- Over a 10 Mbps line of length 2500 m, how long does it take to receive a 100 byte packet?
  - propagation delay is $\tau_c = 10\mu s$
  - assume transmission started at time $t = 0$
  - leading edge of the packet sent at $t = 0$; trailing edge sent at time $\tau_p = 80\mu s$
  - leading edge arrives at $t = 10\mu s$; trailing edge arrives at $t = 90\mu s$

One-way Delay

Pictorial Representation

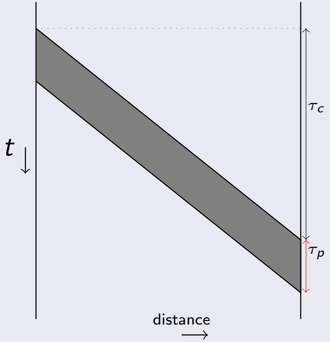

One-way Delay

- One-way delay = Propagation delay τ_c + packet duration τ_p
- y axis is time
- x axis represents distance
- the slope represents speed of propagation

Repeaters

- Used in long transmission lines
- Converts electrical signals to bits, and back to electrical signals, for retransmission
- Repeaters placed strategically to ensure that signal strength received at repeaters is strong enough to eliminate errors
- The conversion process eliminates the effects of noise between repeaters or between repeaters and end-points.
- Repeaters introduce additional propagation delay
- For a line of length 2500 m, with four repeaters (assume each repeater introduces a delay of $4\mu s$) the propagation delay is $10 + 4 \times 4 = 26\mu s$.

Repeater Delay

Repeater Delay

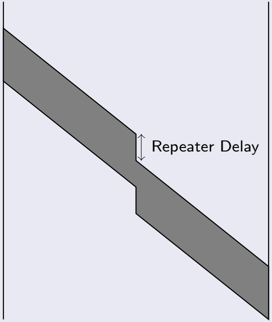

Repeaters

- packet duration does *not* affect repeater delay
- repeaters do not have to receive the entire packet before starting the conversion
- Why not repeaters for analog communications?

Processing Delay

- It takes a finite amount of time for the receiver to "process" the received packet (say, τ_r)
- Data-link frames typically have a cyclic redundancy check (CRC) that will be verified (if inconsistent drop the corrupted packet)
- Routers will need time to look-up routing tables to determine the best next hop.
- There may also be *queuing* delay at routers
- Routers may have several incoming and outgoing interfaces
- A queue in an incoming interface may be stuck if the outgoing interface for the packet at the head of the queue is busy.

Round Trip Time

- Typically, every packet is acknowledged.
- Round trip time (RTT) = Packet duration τ_p + propagation delay τ_c + processing delay τ_r + duration of acknowledgement τ'_p + ack propagation delay τ_c + ack processing delay τ'_r.
- A decent approximation is RTT $= 2 \times (\tau_p + \tau_c + \tau_r)$
- What is the RTT over a 10 Mbps line of length 2500 m, if the packet size (both sent and ACK) is 100 bytes, and the processing delay is 5 μs?
- For the same length, packet size, and processing delay, what is the RTT if the bit-rate is 1 Gbps?

Multi-hop Propagation

One way propagation delay for a channel with n hops is $n(\tau_c + \tau_p + \tau_r)$

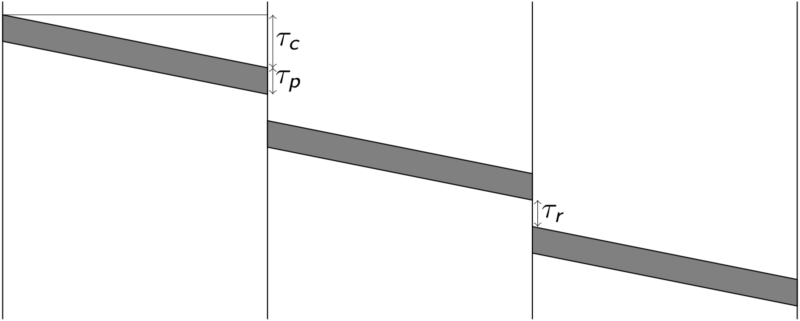

A Simpler Representation of Delays

Representation

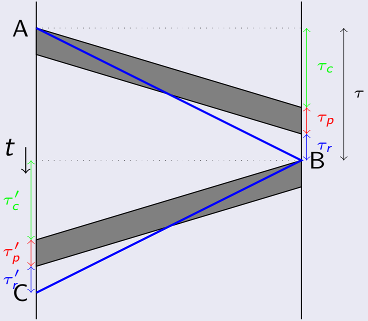

Simpler Representation

- Forward path AB, reverse path BC
- $\tau = \tau_c + \tau_p + \tau_r$ total forward path delay ($\tau' = \tau'_c + \tau'_p + \tau'_r$ total reverse path delay)
- In the simpler representation only the lines AB and BC will be depicted
- Also often rotated by 90 degrees (time in X-axis)

Switching

- Circuit Switching
- Cut-through Switching
- Store and Forward Switching
- Circuit switching was used in Telephone networks (not any longer)

Circuit Switching

- Physical path established over multiple links
- Circuit established at the physical layer level
- Delay for multi-hop propagation is the sum of propagation delay over each link
- Was used in Telephone networks (not any longer)

Cut-through Switching

- Begin forwarding a packet to next hop even before the entire packet has been received
- Leading bytes in the packet will indicate destination
- Useful when look-up for destination can be accomplished very fast
- Delay for multi-hop propagation similar to the situation when repeaters are used.

Store and Forward or Packet Switching

- Entire packet received before the packet can be sent to next hop
- Packet duration influences the overall delay between end-points.
- Two types: Virtual Circuits and Datagram Switching

Datagram Switching

Datagram Subnet

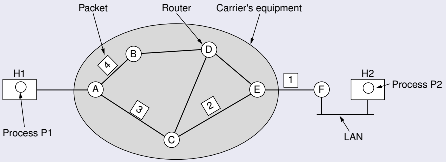

A's table (initially)

Dest.	Line
A	–
B	B
C	C
D	B
E	C
F	C

A's table (later)

Dest.	Line
A	–
B	B
C	C
D	B
E	B
F	B

C's table

Dest.	Line
A	A
B	A
C	–
D	D
E	E
F	E

E's table

Dest.	Line
A	C
B	D
C	C
D	D
E	–
F	F

Datagram Switching

- Routers maintain a routing table indicating best next hop for each destination
- Routing tables are dynamic
- Every packet routed individually
- Paths may change dynamically
- Packets not guaranteed to be received in the same order.

Virtual Circuit Switching

VC Subnet

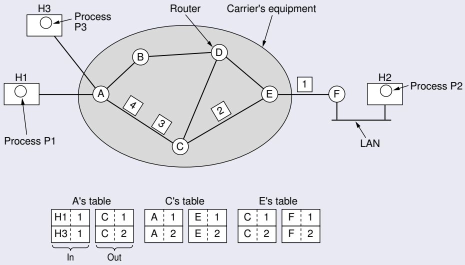

Virtual Circuit

- A path is established before even the first packet can be sent
- Path accepted by all routers in the path
- Packets marked with a path identifier
- Path identifier helps the router determine how to forward the next packet
- Size of routing tables in a router depends on number of active paths through the router.

Datagram vs VC

Issue	Datagram subnet	Virtual-circuit subnet
Circuit setup	Not needed	Required
Addressing	Each packet contains the full source and destination address	Each packet contains a short VC number
State information	Routers do not hold state information about connections	Each VC requires router table space per connection
Routing	Each packet is routed independently	Route chosen when VC is set up; all packets follow it
Effect of router failures	None, except for packets lost during the crash	All VCs that passed through the failed router are terminated
Quality of service	Difficult	Easy if enough resources can be allocated in advance for each VC
Congestion control	Difficult	Easy if enough resources can be allocated in advance for each VC

Connectionless vs Connection-Oriented

- Connectionless: Snail mail, telegram, etc.
- Establishing connections makes the job of end points easier.
- Circuit switching established a physical layer connection
- Virtual circuits are established at the network layer level
- In datagram-routing connection can be established at a higher level (transport layer).

Latency Example

- Audio channel, 100 kbps (bit duration is $10\mu s$)
- 10 hops, each 500 Km. $\tau_c = 20$ ms (total for all 10 hops)
- Assume low processing delay of $1\mu s$ at each hop (total $\tau_r = 10\mu$ s)
- If packet size is 1000 bytes, packet duration $\tau_p = 80ms$
- Total delay $20 + 0.01 + 10 * 80 \approx 820ms$
- What if packet size is 100 bytes?
- Total delay $20 + 0.01 + 10 * 8 \approx 100ms$
- What if channel rate is 100 Mbps? delay due to packet duration becomes negligible.

To realize low latency channels

- Circuit switching — only propagation delay.
- Cut-through switching
- *Store and forward* (packet switching) vs cut-through switching — error checking at every hop
- For circuit and cut-through switching packet duration is irrelevant.
- For packet switching use smaller packet sizes in low-bit rate channels
- In high bit-rate links packet duration can be negligible (compared to propagation delay).
- With ever increasing bit-rates the advantages of Circuit Switching and Cut-through Switching become less relevant.

Reliable Data Transfer

- Need to transmit a sequence of packets over a link.
- This is a requirement for data-link and transport layers
- Link can be real (data link layer) or virtual (transport layer)
- Packets can be corrupted
- Over a virtual link packets can also arrive out of order, be duplicated, dropped, $\cdots$.

ARQ (Automatic Repeat Request) Protocols

- Packets need to be acknowledged
- Timeout for retransmission
- But acknowledgments can get lost too!

Lost Acknowledgements

In (B) ACK was lost. So sender retransmits packet P1. Can both sender and receiver agree that $P1'$ is a retransmission of $P1$ and not $P2$?

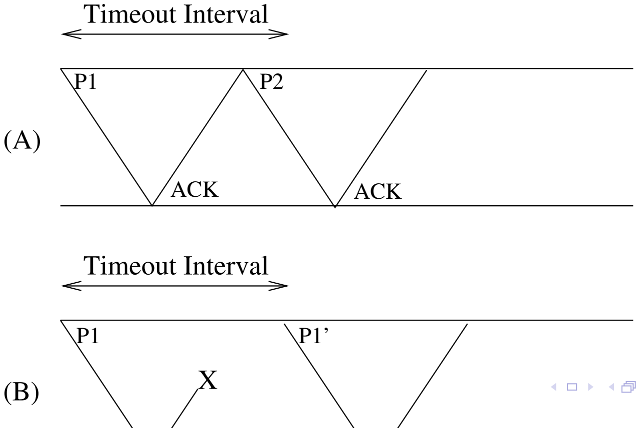

Packet Numbers!

- Packets need to be numbered
- Obviously numbers will be “recycled”
- How many unique sequence numbers do we need?

Ambiguous Acknowledgements

The ACK was P1 was delayed. Sender retransmits P1. Soon after, ACK for P1 is received and sender sends P2 which gets lost.

How does the sender know that the second acknowledgement is for *P*1 or *P*2?

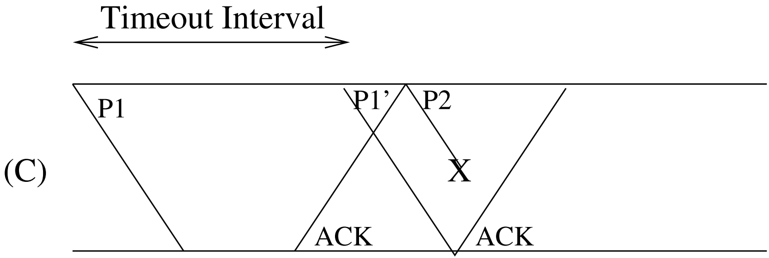

Numbering

- Acknowledgements need to be numbered too!
- ABP protocol (alternating bit protocol)
- One bit packets numbers - 0 or 1
- One bit ACK numbers
- Also called Stop and Wait Protocol (SWP)
- Is ABP/SWP unambiguous?

Alternating Bit Protocol (ABP)

(D)

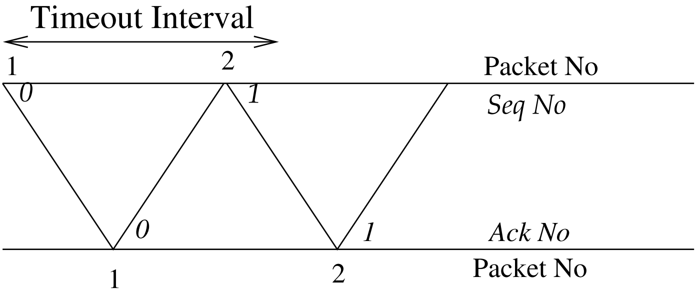

ABP is Unambiguous!

(E)

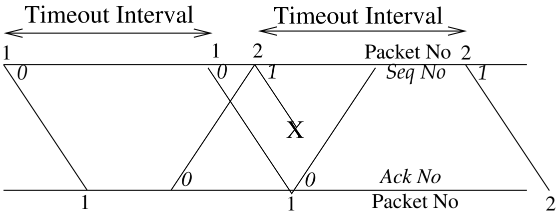

(F)

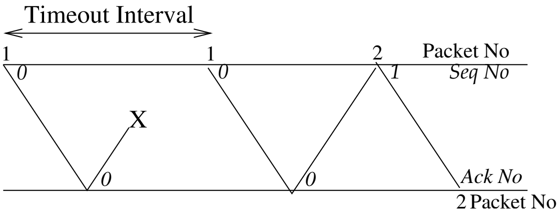

ABP Sender Rules

- If the last packet sent was numbered 0,
  - allowed to send the next packet with number 1 only after ACK for 0 is received
  - (else retransmit 0)
- If the last packet sent was numbered 1,
  - allowed to send the next packet with number 0 only after ACK for 1 is received
  - (else retransmit 1)

ABP Receiver Rules

- If last packet received was numbered 0
  - If 1 received ACK and keep (new packet)
  - If 0 received ACK and drop (retransmission of previous packet)
- If last packet received was numbered 1
  - If 0 received ACK and keep (new packet)
  - If 1 received ACK and drop (retransmission of previous packet)

Effective Data Rate

- 100 Mbps link $A \leftrightarrow B$ ($R = 100x10^6$), $L = 2000m$, packet size $F = 100$ bits, $\tau_r = 1\mu$s
- $P = \tau_p = \frac{100}{100000000} = 1\mu$s, $\tau_c = \frac{2000}{2.5\times10^8} = 8\mu$s
- What is the effective data rate between A and B when ABP is used?
- RTT $= 2\tau_p + 2\tau_c + 2\tau_r = 2 + 16 + 2 = 20.$
- In each 20 μs interval transmission occurs only for 1 μs
- Effective data rate is 5 Mbps (P/RTT x R)

Improving Throughput

- Sender receives the ACK for the first packet after 20 μs
- Sender does nothing for 19 μs
- What if sender is allowed to send 20 packets *before* receiving the ACK for the first packet?
- Only after the ACK for the first packet is received the sender can send packet 21
- Only after the ACK for the second packet is received can the sender send packet 22, and so on
- Window size of 20 packets
- More generally, Window size $\geq$ RTT/packet_duration

Pipelining Example, $W \geq RTT/\tau_p = 13$

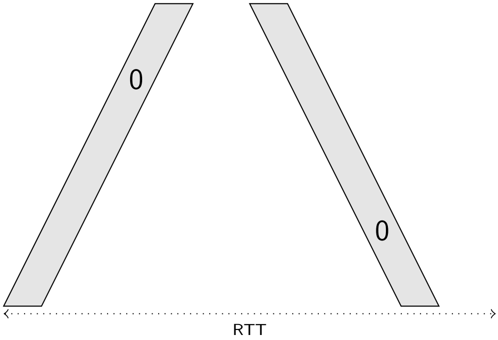

Pipelining Example, $W \geq RTT/\tau_p = 13$

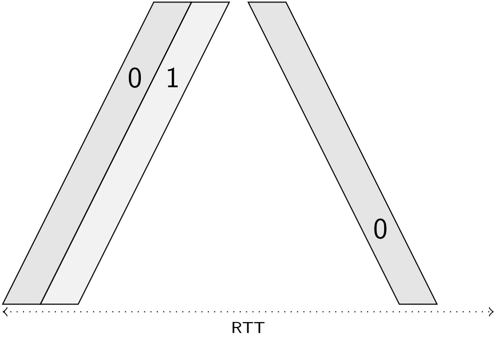

Pipelining Example, $W \geq RTT/\tau_p = 13$

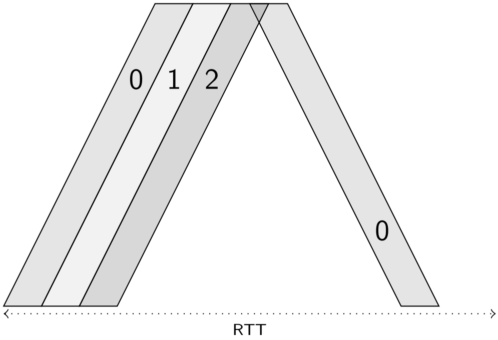

Pipelining Example, $W \geq RTT/\tau_p = 13$

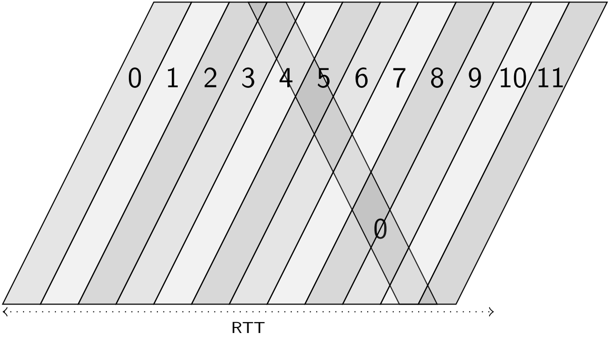

Pipelining Example, $W \geq RTT/\tau_p = 13$

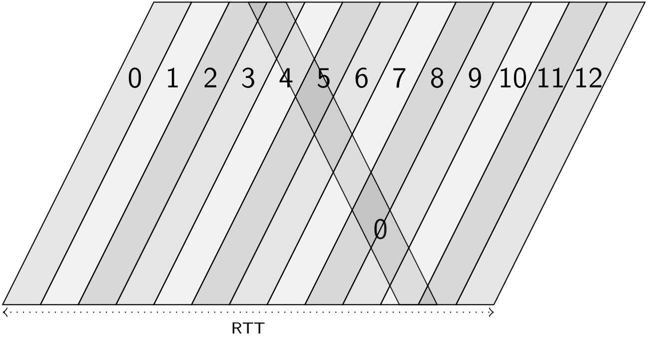

13 packets (0 to 12) sent in one RTT

ACK for 0 processed; ready to send 13

Pipelining Example, $W \geq RTT/\tau_p = 13$

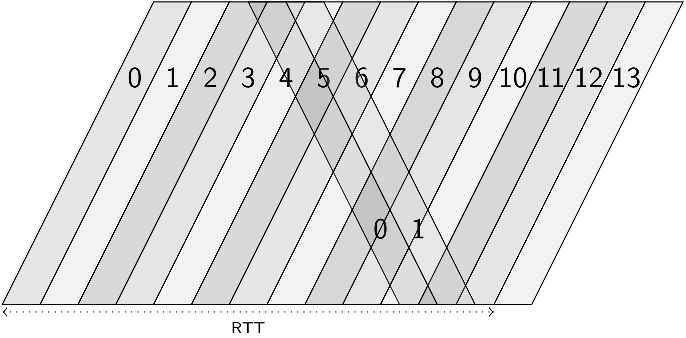

13 packets (0 to 12) sent in one RTT

ACK for 0 processed; ready to send 13

ACK for 1 rcd; ready to send 14

Pipelining Example, $W \geq RTT/\tau_p = 13$

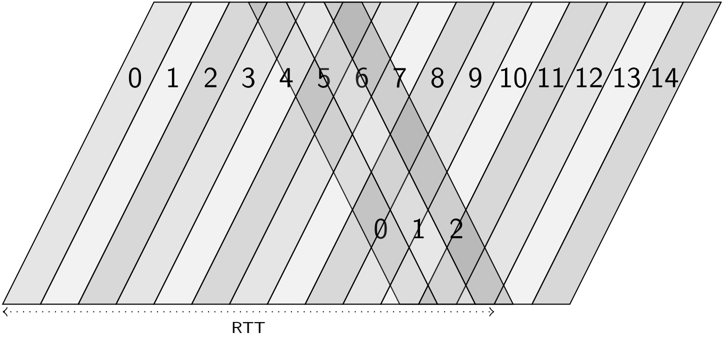

13 packets (0 to 12) sent in one RTT

ACK for 0 processed; ready to send 13

ACK for 1 rcd; ready to send 14

Pipelining

- The secret to higher throughput
- Assumes nodes can do “multiple” things at the same time.
- ABP needs only half-duplex channels, for pipelining we implicitly assume full duplex channels (ACKs and packets cross each other)
- Nodes may need to buffer packets when things do not go well
- What is the buffer size needed?
- How do we number packets and acknowledgements ?

Selective Repeat Protocol

- Similar to ABP, with window size W packets
- Or ABP is SRP with window size 1 packet
- Buffer size of W packets
- Numbering packets?

SRP Protocol

Sender Window (size 5)

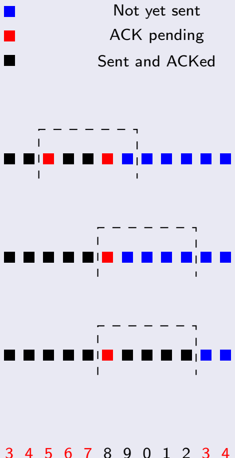

SRP Sender Rules

- Legal to send any packet within the window
- Head of the queue blocked by an outstanding ACK.
- Sender window can not advance until ACK is received for the blocking packet.
- Normally next unsent packet (blue) is sent.
- On time-out for the blocking red packet resend the packet.
- Sender may have to buffer up to W packets.
- $2W$ unique packet numbers — 0 to $2W - 1$.

SRP Window

Rx. Window

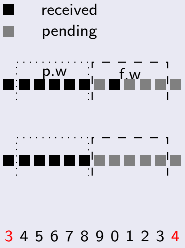

Receiver Rules

- Past Window and Future Window; each of size W.
- Future Window blocked by earliest pending packet.
- Next packet received can be from either window
- If packet from past window ACK and drop packet
- If packet from future window ACK and store packet
- Past Window is the *lower bound* on sender window position. reception of 8 $\rightarrow$ ACK for 3 was already received by sender.
- Future window is the *upper bound* on sender window position

SRP Window

Rx. Window

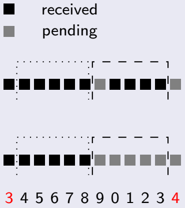

Receiver Rules

- Up to W packets may need to be buffered
- $2W$ possibilities for next packet. Using $2W$ distinct numbers eliminates ambiguities.
- If next received packet is numbered $4, 5, 6, 7$ or 8 ACK and discard.
- If next received packet is numbered $9, 0, 1, 2$ or 3 ACK and store.

Negative Acknowledgments (NACK)

- Assume ACK for packet number 2 not received
- Normally sender would wait for a timeout period
- What if an ACK for 3 (while 2 is missing) is interpreted as a NACK for 2?
- ACK for 3 is generally received well before time-out for 2.

Go Back N - GBN Protocol

- Window size W
- Sender logic is the same
- Receiver does not buffer packets (no window)
- Out of sequence packets are simply discarded
- Periodically sends acknowledgement for the last received packet number
- Example, if Rx gets packets 1,2,3,4,5,7,8,9 the ACK for 5 is sent. 7,8,9 not ACKed.
- Sender retransmits 6,7,8,9...

Windowing Protocols in DL vs TL

- Window size need not always be expressed in packets
- If expressed in packets we assume equal duration packets
- Window size is "number of packet durations in one RTT"
- It can be in bits too "number of bits that can be sent in one RTT"
- Or bytes "number of bytes that can be sent in one RTT"
- In transport layer window size is expressed in bytes
- When windowing protocols are used in data link layer, the link parameters like τ_p, τ_c, τ_r, RTT etc., are known at design time.
- When windowing protocols are used in TL how well do we know the "virtual link" parameters?