MASSACHUSETTS INSTITUTE OF TECHNOLOGY
DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

6.02 Introduction to EECS II
Spring 2011

Quiz 3

Name SOLUTIONS	Score

☐ 10a Devavrat Shah 24-402

☐ 11a Devavrat Shah 24-402

☐ 12n Fabian Lim 38-166

☐ 1p John Sun 38-166

☐ 2p John Sun 38-166

Please write your answers legibly in the spaces provided. You can use the backs of the pages if you need extra room for your answer or scratch work. Make sure we can find your answer!
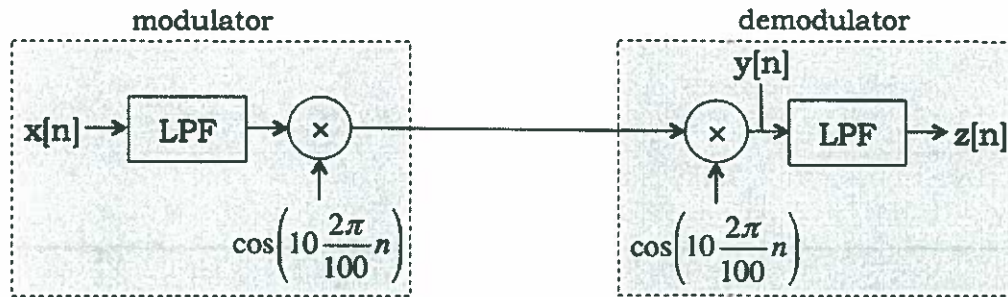
You can use a calculator and one 8.5" x 11" cribsheet.

Partial credit will only be given in cases where you show your work and (very briefly) explain your approach.

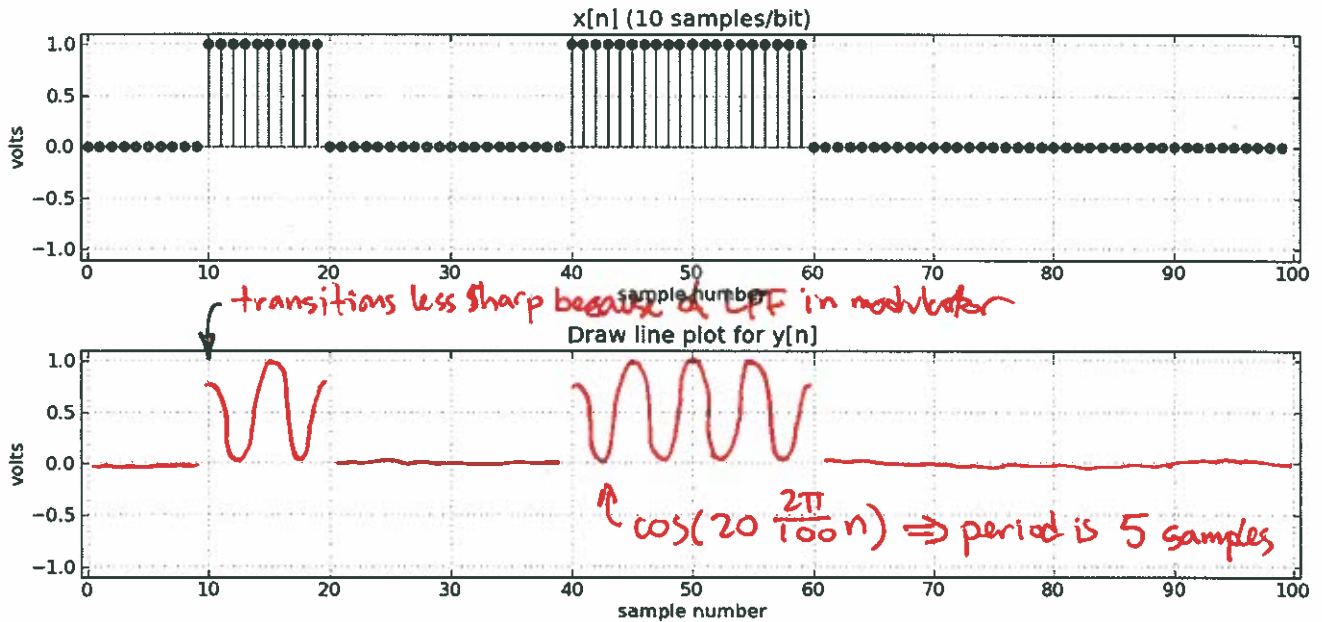| Prob. #1 (31 pts)	Prob. #2 (9 pts)	Prob. #3 (27 pts)	Prob. #4 (33 pts)

Problem 1. Modulation/Demodulation (31 points)

Consider the following standard modulation/demodulation scheme where $N=100$.

modulator   **demodulator**

$x[n] \rightarrow \boxed{\text{LPF}} \rightarrow \times \longrightarrow$ $y[n]$ $\times \rightarrow \boxed{\text{LPF}} \rightarrow z[n]$

$$\cos\left(10\frac{2\pi}{100}n\right) \qquad\qquad \cos\left(10\frac{2\pi}{100}n\right)$$

(A) (5 points) The figure below shows a plot of the input, *x[n]*. Please draw the approximate <u>time-domain waveform</u> for *y[n]*, the signal that is the input to the low-pass filter in the demodulator. Don't bother drawing dots for each sample, just use a line plot to indicate the important timing characteristics of the waveform.

Draw line plot for y[n]

x[n] (10 samples/bit)

transitions less sharp because of LPF in modulator

Draw line plot for y[n]

$$\cos\left(20\frac{2\pi}{100}n\right) \Rightarrow \text{period is 5 samples}$$

(B) (5 Points) Building on the scheme shown in the previous question, suppose there are multiple modulators and demodulators all connected to a single shared channel, with each modulator given a different modulation frequency. If the low-pass filter in each <u>modulator</u> is eliminated, briefly describe what the effect will be on signal *z[n]*, the output of a demodulator tuned to the frequency of a particular transmitter.

**Brief description**

w/ LPF:
$|a_k|$ for xmit

transmission is band-limited

removing LPF!
$|a_k|$

transmission covers entire spectrum, high-freqs show up as noise when added to other transmissions

6.02 Spring 2011   - 2 of 10 -   Quiz 3

(C) (5 Points) The plot on the left below shows a_k, the Fourier series coefficients of the signal at the output of a transmitter where N=36. If the channel introduces a 3-sample delay, please plot the Fourier series coefficients of the signal entering the receiver.
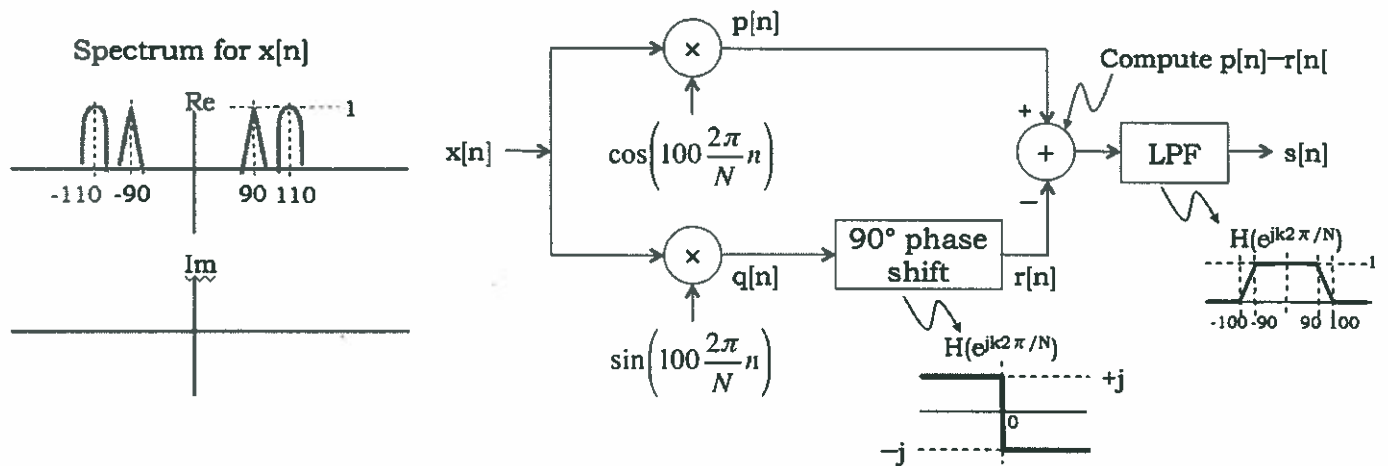
Plot Fourier series coefficients
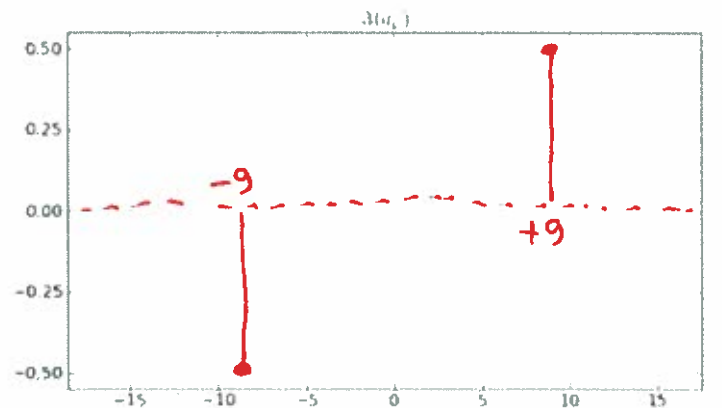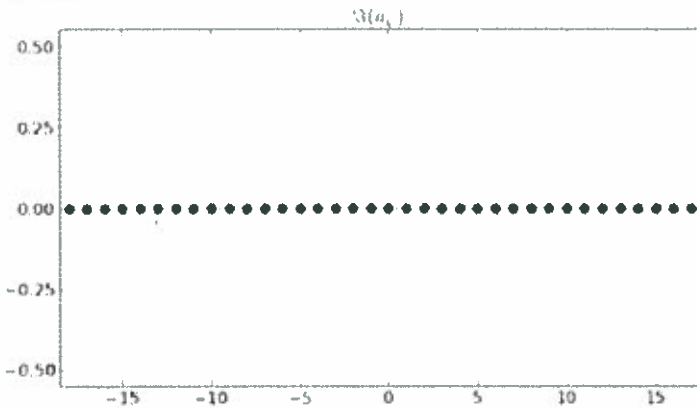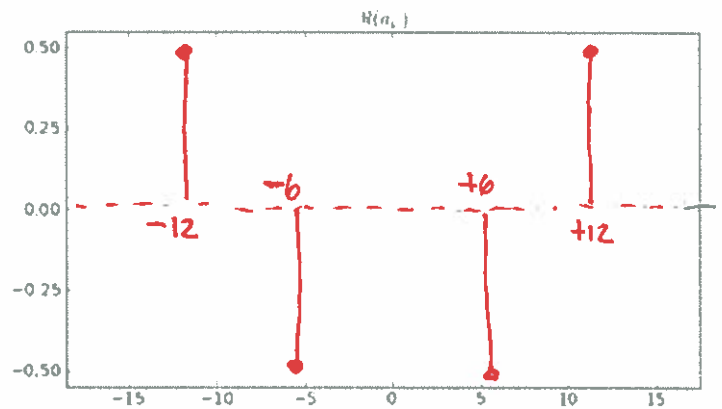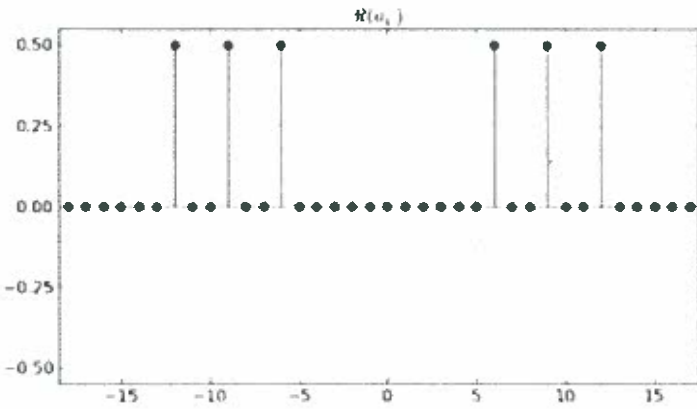
$\cos(6 \frac{2\pi}{36} n) \Rightarrow$ 6 cycles in 36 samples $\Rightarrow$ 1 cycle in 6 samples $\therefore$ 3 samples = ½ cycle
↳ cos → −cos

$\cos(9 \frac{2\pi}{36} n) \Rightarrow$ 9 " $\Rightarrow$ " 4 samples $\therefore$ 3 samples = ¾ cycle
↳ cos → −sin

$\cos(12 \frac{2\pi}{36} n) \Rightarrow$ 12 " $\Rightarrow$ " 3 samples $\therefore$ 3 samples = 1 cycle
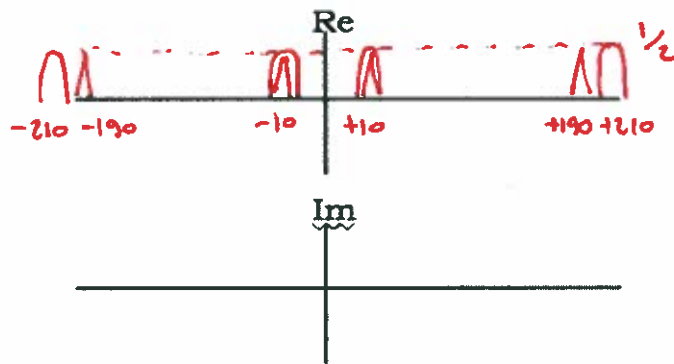↳ cos → cos

(D) (16 Points) The following figure shows an image rejection mixer. The frequency responses of the two filter components (the 90° phase shift and the low-pass filter) are as shown.
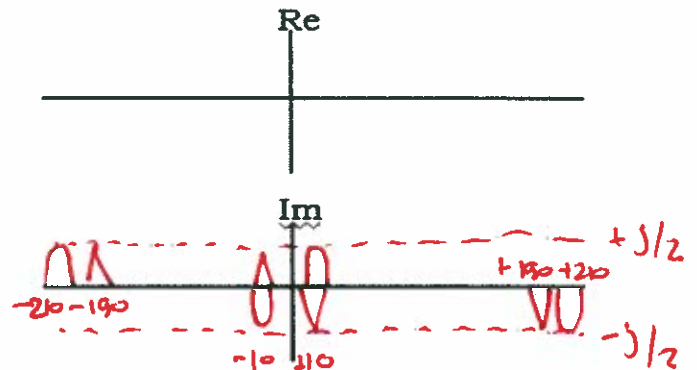
Spectrum for x[n]

Re 1

-110 -90   90 110

Im

p[n]

×

$\cos\left(100\dfrac{2\pi}{N}n\right)$

x[n] →

×   q[n]

$\sin\left(100\dfrac{2\pi}{N}n\right)$

90° phase shift   r[n]

$H(e^{jk2\pi/N})$

+j

0

−j

Compute p[n]−r[n]

+
+   → LPF → s[n]
−

$H(e^{jk2\pi/N})$
.... 1
-100 -90   90 100

The spectral plot to the left in figure above shows the spectrum of the input signal, x[n]. Using the same icon representation of a spectrum, draw the spectrum for signals p[n], q[n], r[n] and s[n] below, taking care to label the center frequency and magnitude of each spectral component. If two different icons overlap, simply draw them on top of one another. If identical icons overlap, perform the indicated addition/subtraction, showing the net result with a bold line.

Draw labeled spectrum for signals p[n], q[n], r[n] and s[n]

Spectrum for p[n]
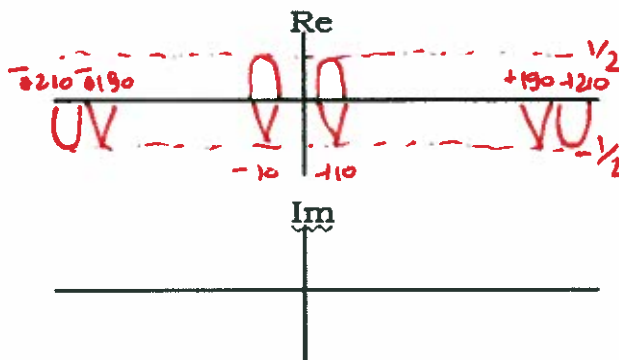
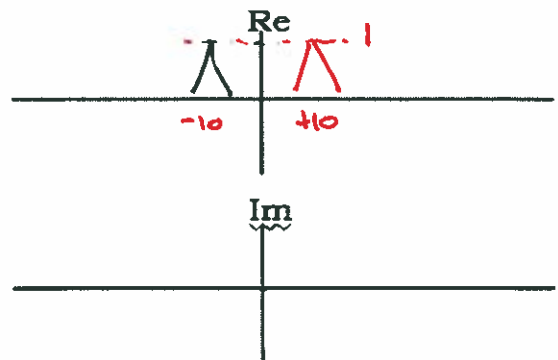Re   ½

−210 −190   −10 +10   +190 +210

Im

Spectrum for q[n]

Re

Im

−210 −190   −10 +10   +190 +210   +j/2

−j/2

Spectrum for r[n]

Re   ½

−210 −190   +190 +210

−10 +10   −½

Im

Spectrum for s[n]

Re   1

−10   +10

Im

Problem 2. Potpourri (9 points)

(A) (4 Points) On the average, 220 students become EECS majors each year. 40% of the students leave with their SB after 3 years. The other 60% continue for an additional 1.5 years to complete their MEng. What is the average number of EECS students (undergraduate majors plus MEng students)?

Average number of EECS students: 858

$$N = \lambda \cdot D$$
$$\lambda = 220 \text{ students/year}$$
$$D_{avg} = 3 + (.6)1.5 = 3.9 \text{ years}$$
$$\Rightarrow N = (220)(3.9) = 858$$

(B) (5 Points) We discussed how a reliable transport protocol sender computes an average round-trip time (RTT) for the connection using an exponential weighted moving average (EWMA) estimator:

$$s(n) = \alpha \cdot r(n) + (1-\alpha) \cdot s(n-1),$$

where $r(n)$ is the n^{th} RTT sample and $s(n)$ the n^{th} smoothed RTT estimate updated after $r(n)$ arrives.

Suppose that at time 0, the initial estimate, $s(0)$, is equal to the true value, $r(0)$. Suppose that immediately after this time, the RTT for the connection increases to a value R much larger than $r(0)$ and remains at that value for the remainder of the connection.

If the retransmission timeout value at step n, TIMEOUT(n), is set to $2 \cdot s(n)$, and $\alpha = 0.2$, calculate the number of RTT samples before we can be sure that there will be no spurious retransmissions. Note that your answer should be an integer.

Number of RTT samples: 4

want TIMEOUT(n) ≥ R to avoid spurious retransmissions

brute force approach:

n	r(n)	s(n)	TIMEOUT(n)
0	r(0)	r(0)	$2 \cdot r(0)$
1	R	.2R	.4R
2	R	.2R + (.8)(.2R) = .36R	.72R
3	R	.2R + (.8)(.36R) = .488R	.976R
4	R	.2R + (.8)(.488) = .5904R	1.18R ✓

Problem 3. Routing (27 points)

Consider the 5-node network shown below. The number shown near the center of each link is the link's cost for transmissions in either direction. The labels at the end of each link are the local name of the link used by the nearby node.
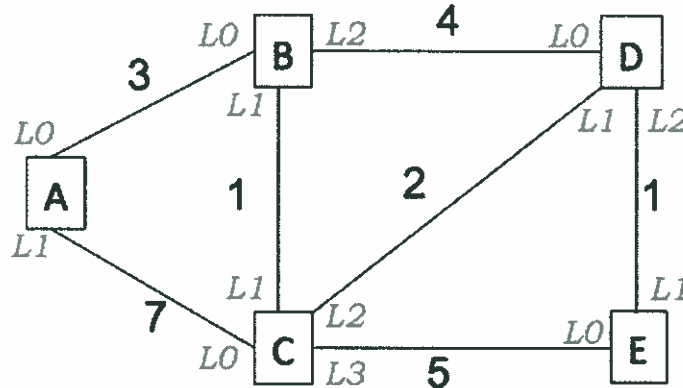
Routing advertisements are originated by each node at intervals of 100 time units: the first advertisement is originated by all nodes at time t=100, the second advertisement at t=200, and so on. Assume that it takes <u>one time step to deliver all queued packets</u> across a link and that the appropriate processing of all incoming packets then happens instantaneously. Thus, <u>the first advertisement arrives at the neighbors of the originating node at t=101 and is processed at that time</u>. If running a link-state protocol, the forwarded LSA will arrive at the neighbor's neighbors at t=102, i.e. flooding happens at one time step per link. Assume that LSA sequence numbers begin at 0.

The questions below ask about the behavior of two different routing protocols operating on the network above: distance-vector (DV) and link-state (LS). Assume that at each time step nodes process all the announcements arriving at that time and then compute a new or updated routing table.

(A) (8 Points) For each routing protocol give the contents of the <u>second</u> routing advertisement generated by node B, i.e., the advertisement originated at time t=200. Include <u>all</u> the information contained in the advertisement.

    includes itself!

second DV advertisement from node B: $[(A,3), (B,0), (C,1), (D,4)]$

    ↳ *routing table only has info about immediate neighbors*

second LS advertisement from node B: { *seqnum=1*   ← *count began at 0.*

    ↳ *all adverts contain list of (neighbor, cost)* $[(A,3), (C,1), (D,4)]$

(B) (8 Points) Assume that the links are stable and that no packets are lost. Then after some number of time steps each node will produce a routing table that will remain unchanged thereafter. For each protocol, what is the earliest time at which all nodes will have produced their final, unchanging routing table?

Time when all nodes have their final routing table using DV: __401__

Time when all nodes have their final routing table using LS: __102__

DV: # of advert cycles required = # hops in ~~longest~~ longest shortest path
  = 4 in this topology (A ↔ E)

LS: correct table constructed when flooding of
  LSA is complete
  = 2 timesteps in this topology

(C) (4 Points) Please fill in the contents of the final, unchanging routing table for node C.

Fill in table

Node	Link	Cost
A	L1	4
B	L1	1
C	---	0
D	L2	2
E	L2	3

(D) (7 Points) Now consider a 5-node network with the same link topology but possibly different link costs. This network uses the LS routing protocol where the shortest paths are determined using Dijkstra's algorithm. The final routing table for <u>node A</u> is shown on the left. Please annotate each link in the diagram on the right with the most specific deduction you can make about the link's cost: give the exact cost if that can be determined, or indicate "$\geq k$" if you determine that k is a lower bound on the possible link cost. The costs obey two constraints: (1) *all costs are integers ≥ 1* and (2) *each shortest path is unique*, i.e., there aren't two possible shortest paths between any two nodes. Note that the link costs apply to transmissions in either direction.

Annotate links with cost deduction

Node	Link	Cost
A	---	0
① B	L0	5
② C	L1	6
③ D	L0	7
④ E	L1	9
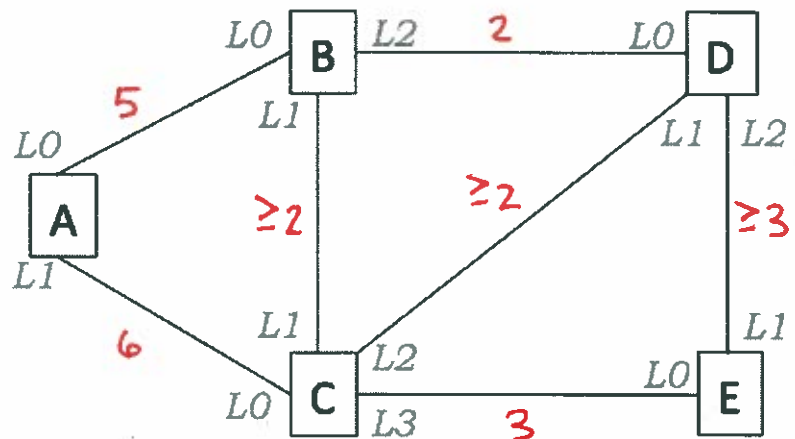
① direct link to B is shortest, so A-B has cost 5

② direct link to C is shortest, so A-C has cost 6

∴ B-C has cost ≥ 2 (if it was 1, the shortest path from A-C wouldn't be unique)

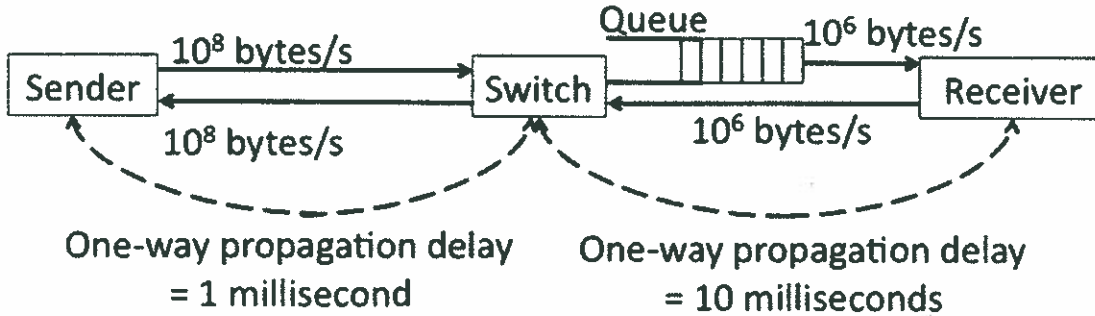③ possible paths A→D:
- A-B-D ⇒ B-D link has cost 2 ✓
- A-B-C-D
- A-B-C-E-D ⇒ A-B-C has cost ≥7 so these alternatives aren't possible

▦ ∴ C-D has cost ≥2 (if it was 1, the shortest path wouldn't be unique)

④ possible paths A→E:
- A-C-D-E ⇒ D-E cost is 1, but then A-B-D-E cost is 8, a contradiction ⊗
- A-C-E ↘ C-E cost is 3

Problem 4. Reliable Data Transport (33 points)

Consider the network topology shown below. Assume that the processing delay at all the nodes is negligible. Data packets are 1000 bytes long, acknowledgements are 10 bytes long (short enough that you can treat the transmission time of an ACK as 0). The queue has a capacity of 30 packets. The receiver will generate an acknowledgement only _after_ the entire incoming packet has arrived.

Sender — 10^8 bytes/s → **Switch** — Queue — 10^6 bytes/s → **Receiver**

10^8 bytes/s ← 10^6 bytes/s

One-way propagation delay = 1 millisecond

One-way propagation delay = 10 milliseconds

(A) (9 Points) Assume no packets are lost. What is the round-trip time (RTT) for communications between the Sender and Receiver? If the sender uses a stop-and-wait transport protocol, what is the maximum throughput it can achieve? Using this protocol what is the utilization of the bottleneck link?

Round-trip time (sec): **23**

Throughput (packets/sec): $\frac{1}{RTT} = \frac{1}{.023} = 43.5$

Utilization of bottleneck link (%): $\frac{43.5}{100} \cdot 100 = 4.35\%$

$t_{prop} : 1 + 10 + 10 + 1 = 22ms$

$t_{xmit} = fast + 1 + fast + fast \approx 1 \ ms$ $\Big\}$ $RTT \cong 23 \ ms$ (actually a little larger)

(B) (4 Points) Now suppose the per-link loss rate is $p = .02$ for each of the four hops in the round-trip path between sender and receiver (no packets are lost between the switch and the queue). If the sender's timeout is set to 50 milliseconds, what is the maximum throughput the sender can expect to achieve using a stop-and-wait protocol?

Throughput (packets/sec): $\frac{1}{T} = 36.75$

T = time to successfully transmit a packet

$= RTT + \frac{L}{1-L} \cdot TIMEOUT = .023 + \frac{.078}{.922}(.050) = .0272$

$\hookrightarrow L = (1 - (1-.02)^4) = .078$

For the following questions assume that the sender is using a sliding window protocol with window size W. Assume no packet losses occur for any reason.

(C) (4 Points) If W is set to 5, what is the throughput the sender can expect to achieve?

Throughput (packets/sec): _217_

$$\text{throughput} = \min\left(B, \frac{W}{RTT}\right)$$

$$= \min\left(1000, \frac{5}{.023}\right) \cong 217$$

(D) (8 Points) What is the smallest value of W that gives the maximum possible throughput? What is the maximum possible throughput?

Value of W for maximum throughput: _23 or 24_

Maximum possible throughput (packets/sec): _1000_

↳ is equal to B.

$$W_{\text{for max throughput}} = \lceil B \cdot RTT_{min} \rceil$$

$$= \lceil 1000(.023) \rceil = 23 \text{ packets or } 24 \text{ packets}$$

↑ slightly larger   ↑ using exact RTTmin

(E) (8 Points) If W is set to 30, what is the round-trip time experienced by a packet? What is throughput the sender can expect to achieve? Hint: some of the 30 outstanding packets will be in-flight on the links (how many?) and the rest will be sitting in the queue. The number of packets in the queue will determine the queuing delay.

Round-trip time: _30 ms_

Throughput (packets/sec): _1000_

↑ limited by B

$$RTT = \underbrace{(t_{propagation} + t_{xmit})}_{23\,ms} + t_{queuing}$$

of the 30 outstanding packets, 23 will be in-flight and the other 7 will be in the queue:

$$= 23 + 7 = 30\,ms$$

$$D = \frac{N}{x} = \frac{7\,\text{pkts}}{1000\,\text{pkts/sec}} = 7\,ms$$

END OF QUIZ 3! Have a good summer ☺