TOTAL SCORE = $\dfrac{89 + 46 + 59}{296} = 66.0$

6.02 Digital Communication Systems—Spring 2009

Quiz 1 - 7:30-9:30pm (Two Hours)
Thursday, March 5, 2009

Check your section	Section	Time	Room	Rec. Instr.
☐	1	10-11	13-5101	Chris Terman
☐	2	11-12	13-5101	Chris Terman
☐	3	1-2	5-233	Mythili Vutukuru
☐	4	2-3	5-233	Mythili Vutukuru
☐	5	1-2	38-166	Vladimir Stojanovic
☐	6	2-3	38-166	Vladimir Stojanovic

Directions: The exam consists of 6 problems on 10 pages. Please make sure you have all the pages. **Enter all your work and your answers directly in the spaces provided on the printed pages of this exam. Please make sure your name is on all sheets. DO IT NOW!** All sketches must be adequately labeled. Unless indicated otherwise, **answers must be derived or explained in the space provided**, not just simply written down. This examination is closed book, but students may use one 8 1/2 × 11 sheet of paper for reference. Calculators may not be used.

The probability density function for a zero-mean unit standard deviation Normal(Gaussian) random variable is:

$$f_X(x) \equiv \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}}$$

The cumulative distribution function for a zero-mean unit standard deviation Normal(Gaussian) random variable is:

$$\Phi(x) \equiv \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{\frac{-\hat{x}^2}{2}} d\hat{x}$$

Please leave the rest of this page blank for use by the graders:

Problem	No. of points	Score	Grader
1	5	5	
2	10	10	
3	17	17	
4	24	22	
5	24	22	
6	20	13	
Total	100	89	

 NAME: Sol Key

Problem 1 (5 points)

In the following plot of a voltage waveform from a transmitter, the transmitter sends 0 Volts for a zero bit and 1.0 Volts for a one bit, and is sending bits with with a certain number of samples per bit.

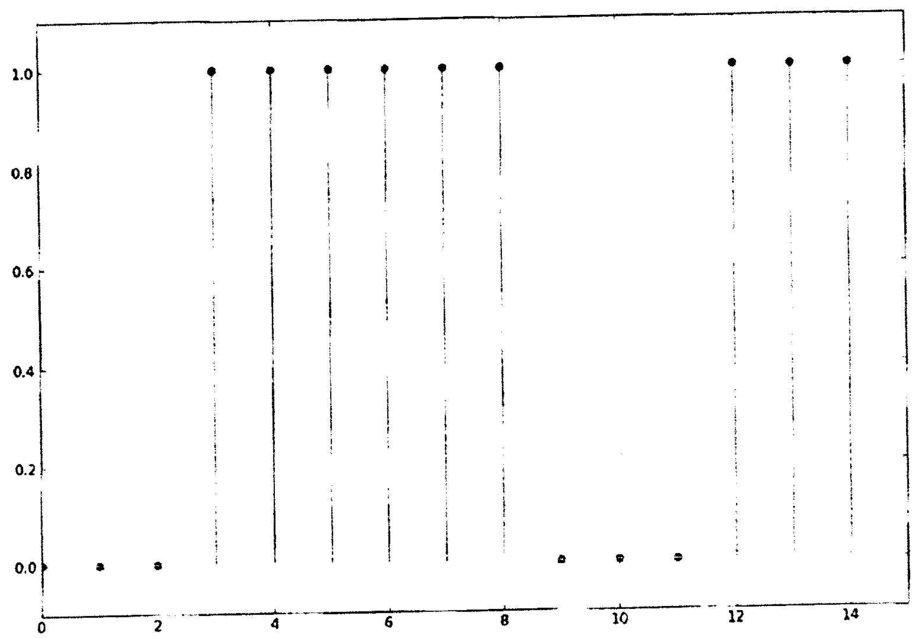

1A. (2 points) What is the largest number of samples per bit the transmitter could be using?

3 ✓

1B. (3 points) What is the sequence of bits being sent? ✓

01101

Problem 2 (10 points)

The input sequence to a linear time-invariant (LTI) system is given by

$x[0] = 0$,
$x[1] = 1$,
$x[2] = 1$ and
$x[n] = 0$ for all other values of n

and the output of the LTI system is given by

$y[0] = 1$,
$y[1] = 2$,
$y[2] = 1$ and
$y[n] = 0$ for all other values of n.

2A. (3 points) Is this system causal? Why or why not?

No, it isn't causal because the system response precedes the input in time ✓

2B. (7 points) What are the nonzero values of the output of this LTI system when the input is

$x[0] = 0$,
$x[1] = 1$,
$x[2] = 1$,
$x[3] = 1$,
$x[4] = 1$ and
$x[n] = 0$ for all other values of n?

$y = h * x$

$y[n] = \sum_i h[i]x[n-i]$

$y[n] = x[n+1] + x[n]$

$y[0] = 1$

$y[1] = 2$

$y[2] = 2$

$y[3] = 2$

$y[4] = 1$

$y[n] = 0$ for all other n ✓

Problem 3 (17 points)

Suppose the voltage sampled by a receiver is $1.0 + noise$ volts when the transmitter sends a one bit, and $0.0 + noise$ volts when the transmitter sends a zero bit, where $noise$ is a zero-mean Normal(Gaussian) random variable with standard deviation σ.

3A. (4 points) If the transmitter is equally likely to send zeros or ones, and 0.5 volts is used as the threshold for deciding the bit value, give an expression for the bit-error rate (BER) in terms of the zero-mean unit standard deviation Normal cumulative distribution function, Φ, and σ.

$$BER = \Phi\left(\frac{-0.5}{\sigma}\right)$$

the probability of an error given 1's & 0's are equally likely

3B. (7 points) If the transmitter is equally likely to send zeros or ones, and 0.5 volts is used as the threshold for deciding the bit value, for what value of σ is the probability of a bit error approximately equal to $\frac{1}{5}$? Note that $\Phi(0.85) \approx \frac{4}{5}$ (see cover page for definition of Φ).

$\Phi(0.85) = 4/5$ w/ $\sigma = 1$

$\Phi\left(\frac{0.5}{\sigma}\right) = 4/5$ $\quad 0.85 = \frac{0.5}{\sigma}$

$$\boxed{\sigma = 0.588}$$

3C. (3 points) Will your answer for σ in part 3B change if the threshold is shifted to 0.6 volts? Do not try to determine σ, but justify your answer.

Yes, p10 (probability of a 1 being read as a zero) will increase more than p01 decreases, so the system is more susceptible to noise errors.

3D. (3 points) Will your answer for σ in part 3B change if the transmitter is twice as likely to send ones as zeros, but the threshold for deciding bit value is still 0.5 volts? Do not try to determine σ, but justify your answer.

No. Because the threshold is in the middle, the relative proportions don't affect the BER. (Although lowering the threshold would decrease BER)

 NAME: Sol Key

Problem 4 (24 points)

Consider three linear time-invariant systems, denoted I, II, and III, each characterized by their unit-sample responses:

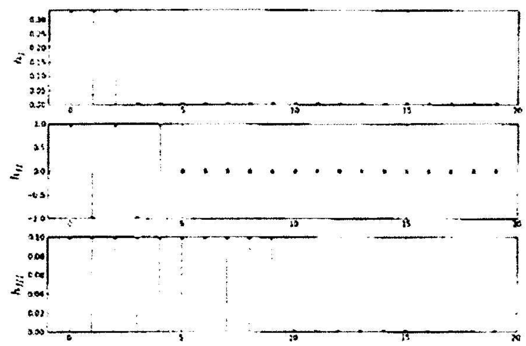

There are also three possible inputs to each of these linear time-invariant systems:

Sequence X_1:$x_1[n] = u[n]$
Sequence X_2:$x_2[n] = (-1)^n(u[n] - u[n-5])$ →
Sequence X_3:$x_3[n] = u[n] - u[n-5]$

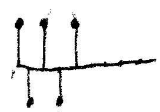

Note that $u[n]$ denotes the unit step. That is $u[n] = 1$ for $n \geq 0$ and is zero otherwise.

4A. (12 points) Which system (I, II or III) and which input (X_1, X_2 or X_3) produced the following output, **AND** what is the numerical value of D? In addition to justifying your answer, for possible partial credit, briefly explain why you ruled out particular combinations of the given LTI systems and responses.

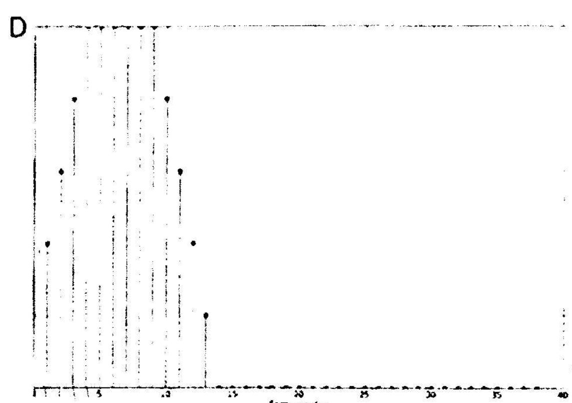

III & X_3

The system rises over five values, which is characterized by system III. Also, the system is finite and monotonic, only corresponding to input X_3

$D = y[4] = h[0]x[4] + h[1]x[3] + h[2]x_{\ldots} + h[3]x[3] + h[4]x[0]$

$y[4] = 0.35 \cdot 1 + 0.35 + 0.35 + 0.35 + 0$

$y[3] =$

$= 1.75$

−2 ~~Arr~~ ARITHMATIC ERROR

10

Problem 4 (continued)

Consider the same three linear time-invariant systems as in 4A, denoted I, II, and III, each characterized by their unit sample responses (repeated here for your convenience):

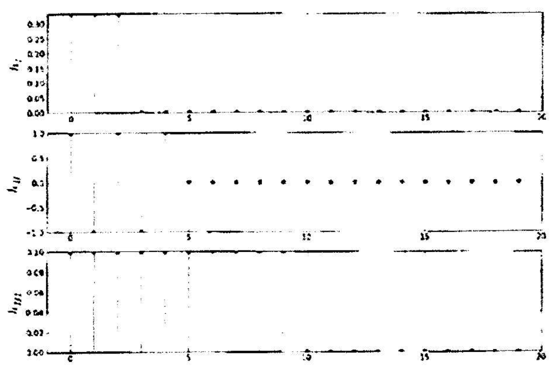

There are also the same three possible inputs to each of these linear time-invariant systems (again repeated here for your convenience):

Sequence X_1:$x_1[n] = u[n]$
Sequence X_2:$x_2[n] = (-1)^n(u[n] - u[n-5])$
Sequence X_3:$x_3[n] = u[n] - u[n-5]$

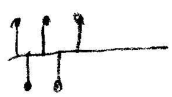

4B. (12 points) Which system (I, II or III) and which input (X_1, X_2 or X_3) produced the following output, **AND** what is the numerical value of D in the plot below? In addition to justifying your answer, for possible partial credit, briefly explain why you ruled out particular combinations of the given LTI systems and responses.

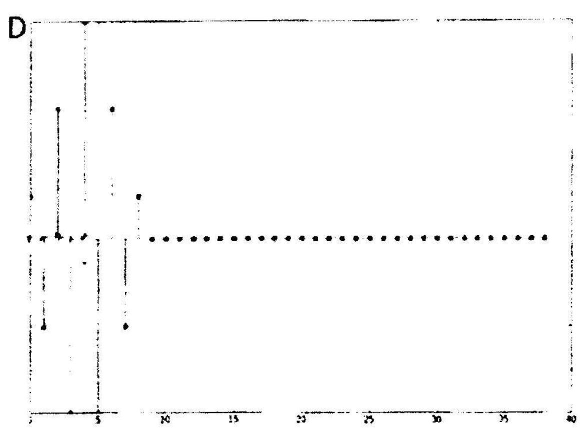

II & X_2

$y[4] = D = h[0]x[4] + h[1]x[3] + h[2]x[2] + h[3]x[1] + h[4]x[0]$

$= 1\cdot 1 + (-1)(-1) + (1)(1) + (-1)(-1) + 1\cdot 1$

$y[3] = 1\cdot(-1) + (-1)$

$\boxed{D = 5}$

It must be oscillating so it must contain either II or X_2. the signal has less than 10 values so it can't be III and the frequency amplifies suggesting X_2 & II

Problem 5 (24 points)

This question refers to the LTI systems, I, II and III, whose unit sample responses were given in questions 4 (again repeated here for your convenience):

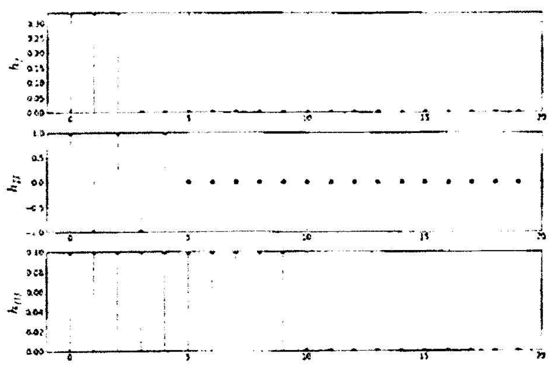

In this question, the input to these systems are bit streams with eight voltage samples per bit, with eight one-volt samples representing a one bit and eight zero-volt samples representing a zero bit.

5A. (12 points) Which system (I, II or III) generated the following eye diagram? To ensure at least partial credit for your answer, explain what led you to rule out the systems you did not select.

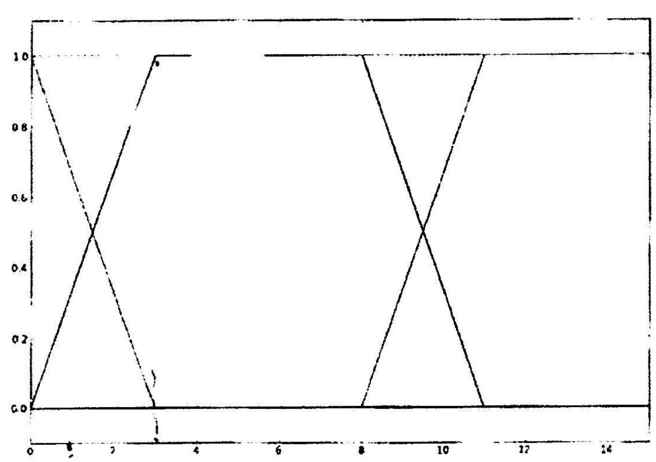

I, there ~~are~~ too few bits of 1s
∴ it can't be III,
there is no ringing
so it can't be II

Problem 5 (continued)

This question refers to a fourth LTI system whose unit sample response, $h_{IV}[n]$ is given below:

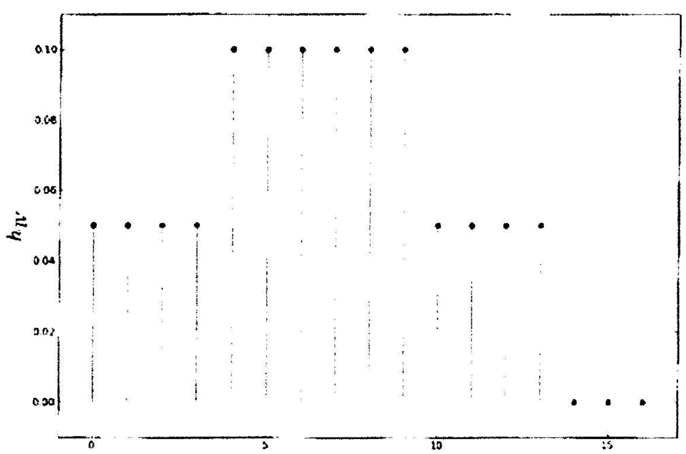

where, just like in problem 5A, the input to this system is a bit stream with eight voltage samples per bit, with eight one-volt samples representing a one bit and eight zero-volt samples representing a zero bit.

5B. (12 points) Determine the voltage level denoted by D in the eye diagram generated from the system with unit sample response $h_{IV}[n]$.

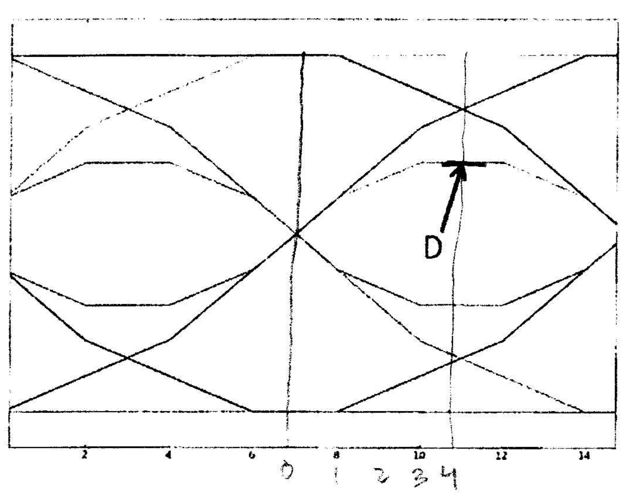

voltage level D =

is $y[4]$, occurs when preceding 2 bits were 0 and current bit is a 1

$$y[4] = h[0]x[4] + h[1]x[3] + h[2]x[2] + h[3]x[1] + h[4]x[0]$$

$$y[4] = 0.045\cdot 1 + 0.045\cdot 1 + 0.045\cdot \cancel{1} + 0.045\cdot 1 + 0.1\cdot 1$$

$$\boxed{= 0.28V}$$

$\boxed{-2}$ MISREAD GRAPH AS VALUES OF 0.045, INSTEAD

10

Problem 6 (20 points)

Consider a transmitter that encodes pairs of bits using four voltage values. Specifically:

00 is encoded as zero volts,
01 is encoded as ${}^{1}/_{3}V_{high}$ volts,
10 is encoded as ${}^{2}/_{3}V_{high}$ volts and
11 is encoded as V_{high} volts.

For this problem we will assume a wire that only adds noise. That is,

$$y[n] = x[n] + noise[n]$$

where $y[n]$ is the received sample, $x[n]$ the transmitted sample whose value is one of the above four voltages, and $noise[n]$ is a random variable.

Please assume all bit patterns are **equally likely** to be transmitted.

Suppose the probability density function for $noise[n]$ is a constant, K, from -0.05 volts to 0.05 volts and zero elsewhere.

6A. (3 points) What is the value of K?

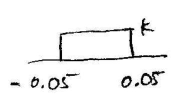

0.1 width * height = area
0.1 * h = 1
h = 10
$\boxed{K = 10}$

6B. (3 points) What is the smallest numerical value for V_{high} such that a threshold detector can determine the transmitted bit pair from $y[n]$ WITHOUT ERROR. Assume the threshold detector needs a forbidden zone of 0.1 volt.

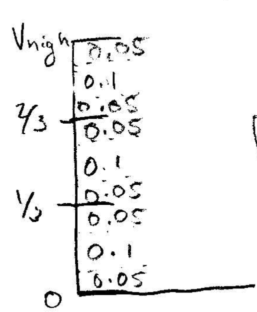

$\boxed{= 0.6\ V}$

Problem 6 (continued)

Suppose now $V_{high} = 1.0$ volts and the probability density function for $noise[n]$ is a zero-mean Normal with standard deviation σ.

6C. (7 points) If $\sigma = 0.001$, what is the **approximate** probability that $\frac{1}{3} < y[n] < \frac{2}{3}$? You should be able to give a numerical answer.

$y[n] = x[n] + noise[n]$

0

$p = 1/3$ X

b/c the probability of all bit sequences are equal and Gaussian distribution is symmetrical.

6D. (7 points) If $\sigma = 0.1$, is the probability that a transmitted 01 ($\frac{1}{3}$ volts) will be incorrectly received the same as the probability that a transmitted 11 (1.0 volts) will be incorrectly received? Explain your answer.

No, they are not the same because a bit error for 11 can only occur if the noise is $\leq -1/6$ whereas a bit error can occur for a 01 if either noise $\leq -1/6$ OR $\geq 1/6$, twice the error rate.

7

End of Quiz 1!

Problem 1. (15 points)

(A) (1 point) Two code words A and B have a Hamming distance of 2. If an even parity bit is appended to each code word, producing A' and B', what is the Hamming distance between A' and B'?

Hamming distance between A' and B': 3 ✗

(B) (3 points) Consider a channel encoding that simply replicates each bit 8 times, i.e., "0" is encoded as the code word "00000000" and "1" is encoded as the code word "11111111".

(n,k,d) designation for this code: (1, 8, 8) ✓

code rate for this code: 1/8 ✓

If one uses this code to detect and correct errors of up to C bits, what's the maximum possible value for C?

Maximum possible value for C: 2 ✓

(C) (3 points) Three parity bits (P_1, P_2, P_3) calculated as specified below are concatenated with five message bits $(D_1, \ldots, D_5)$ to form an 8-bit code word. $\oplus$ means XOR (addition modulo 2).

$$P_1 = D_1 \oplus D_2 \oplus \cancel{D_3}$$
$$P_2 = D_2 \oplus \cancel{D_3} \oplus \cancel{D_4}$$
$$P_3 = \cancel{D_3} \oplus \cancel{D_4} \oplus \cancel{D_5}$$

After a transmission involving at most a single bit error, checking the received parity bits indicates parity errors involving P_1 and P_2 but not P_3. What correction (if any) is indicated?

Correction (if any): flip D_2 ✓

Can this code be used to perform single-bit error correction assuming that the error can occur in any one of the eight code word bits? Briefly explain your reasoning.

Single-bit error correction possible (YES or NO): NO ✓

if an error only occurs in D_1, then the P_1 will be wrong so correction is ambiguous ([illegible] for D_5)

(5/7)

(D) (3 points) A $RS(n, k)$ Reed-Solomon code encodes a block of k message symbols into a block of n code word symbols; in this problem each symbol is an 8-bit byte. A $RS(n, k)$ code can detect and correct any combination of E symbol errors and S symbol erasures so long as $2E + S \leq n - k$.

Consider an encoding scheme where a 28-byte data block is encoded as a 32-byte message block using a $RS(32, 28)$ code. Then thirty-two 32-byte message blocks undergo a 32-way interleaving, transmitting the first byte of each of the thirty-two message blocks, then the second byte, and so on until all 1024 code word bytes have been sent.

If the transmission is corrupted by one error burst of B bytes, what is the longest burst that can be corrected by the encoding scheme? Briefly explain your reasoning.

2

max 64 errors correctable =

Maximum value of B: 128

$\boxed{-1}$

$2S = n - k$ (no erasures)

$S = 32 - 28 = 4$ errors

ALGEBRA ERROR

4 errors × 32 = 128

before one block contains 5 errors and is corrupted.

(E) (5 points) Congratulations! You've been hired by the Registrar to come up with a binary encoding for the gender field of the records database. They want an encoding that allows detection of errors of up to three bits in the gender field (no correction required; they'll use back-up tapes to restore fields where errors have been detected). Assume that three (!) genders (Male, Female, Other) need to be encoded. Please indicate the required Hamming distance between your code words and specify an appropriate code word for each gender. Your code doesn't have to be optimal, but it must support 3-bit error detection.

Required Hamming distance: 4 ✓

Binary code word for Male: 00000000 ✓

5

Binary code word for Female: 11111111 ✓

Binary code word for Other: 01010101 ✓

Problem 2. (18 points)

Consider a convolutional code with three generator polynomials

$$G_0 = 11 \Rightarrow p_0[n] = x[n] \oplus x[n-1] \quad 11$$
$$G_1 = 10 \Rightarrow p_1[n] = x[n] \quad 10$$
$$G_2 = 01 \Rightarrow p_2[n] = x[n-1] \quad 01$$

The figure below is a snapshot of the decoding trellis showing a particular state of a maximum likelihood decoder implemented using the Viterbi algorithm. The labels in the boxes show the path metrics computed for each state after receiving the incoming parity bits at time t. The labels on the arcs show the expected parity bits for each transition; the actual received bits at each time are shown above the trellis.

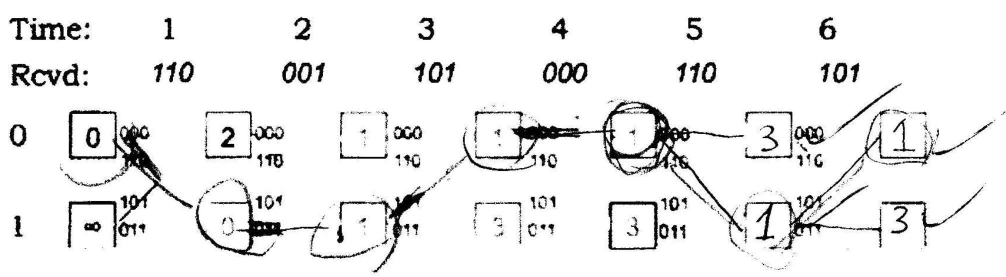

(A) (2 points) What is the code rate r and constraint length k for this code?

r: 1/3

k: 2

(2)

(B) (3 points) Draw the state transition diagram for a transmitter that uses this convolutional code. The states should be labeled with the binary string $x[n-1]...x[n-k+1]$ and the arcs labeled with $x[n]/p_0p_1p_2$ where $x[n]$ is the next message bit and p_0, p_1 and p_2 are the three parity bits computed from $x[n]...x[n-k+1]$ using G_0, G_1 and G_2.

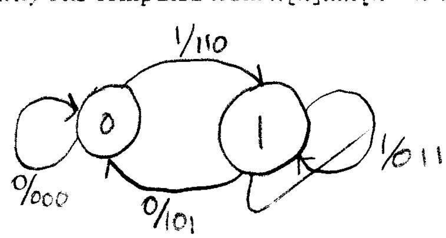

(Draw state transition diagram)

(3)

(C) (4 points) Fill in the path metrics in the empty boxes in the trellis diagram on the previous page (corresponding to the Viterbi calculations for times 5 and 6).

4

(Fill in path metrics)

(D) (2 points) Considering the original trellis (i.e., the parts that had already been filled in before you did part C), what is the most-likely final state through time 4? How many errors were detected along the most-likely path to that state?

Most-likely final state through time 4: 0 ✓

Number of errors detected: 1 ✓

2

(E) (4 points) Again, considering only through time 4, what's the most-likely decoded message? You may find it helpful to mark the most-likely path through the trellis up until time 4.

Most-likely decoded message through time 4: 01100

(OMIT B/C GRAPH COULDN'T BE CLEANED UP)

(F) (3 points) Based on your choice of the most-likely path through the trellis up until time 4, at what time(s) did the error(s) occur?

Time(s) error(s) occurred: $\boxed{2}$ ✓

14/14

Problem 3 (33 points)

In answering the four parts of this question, consider three linear time-invariant systems, denoted I, II, and III, each characterized by the magnitude of their frequency responses, $|H_I(e^{j\Omega})|$, $|H_{II}(e^{j\Omega})|$, and $|H_{III}(e^{j\Omega})|$, given in the plot below. It may be helpful to recall that $\cos\frac{\pi}{4} = \frac{\sqrt{2}}{2}$ and $\cos\frac{5\pi}{6} = -\frac{\sqrt{3}}{2}$.

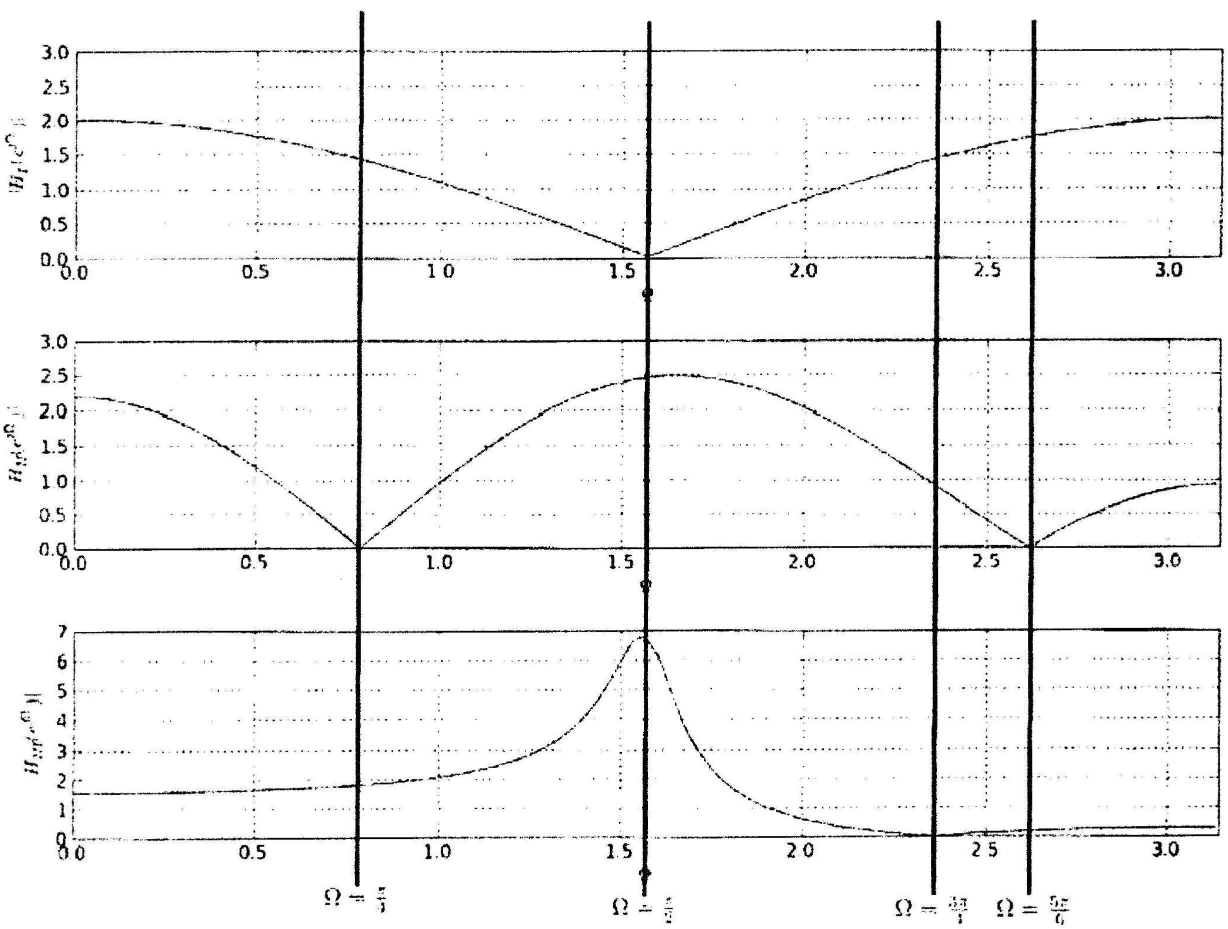

(A) (7 points) Which frequency response (I, II or III) corresponds to the system

$$(1.1)^2 y[n] + \alpha y[n-1] + y[n-2] = x[n] + \sqrt{2}x[n-1] + x[n-2]$$

and what is the numerical value of α?

III

$-2\cos(\Omega)$

$\alpha = 0$ $\quad -2\cos(\pi/2) = 0$

7

(B) (9 points) Which frequency response (I, II or III) corresponds to the system

$$y[n] = x[n] + (\sqrt{3.0} - \sqrt{2.0})x[n-1] + \beta x[n-2] + (\sqrt{3.0} - \sqrt{2.0})x[n-3] + x[n-4]$$

and what is the numerical value of β?

II $\quad \beta =$

$h_1 = \{1, -2\cos(\pi/4), 1\}$

$h_2 = \{1, -2\cos(5\pi/6), 1\}$

9

$h = h_1 * h_2$

$h[0] = h_1[0]h_2[0] = 1$

$h[1] = h_1[0]h_2[1] + h_1[1]h_2[0] = -2\cos\pi/4 + (-2\cos\frac{5\pi}{6}$

$h[2] = h_1[0]h_2[2] + h_1[1]h_2[1] + h_1[2]h_2[0] = 1 + 1 + (-2)\left(\frac{\sqrt{2}}{2}\right)\left(\frac{\sqrt{3}}{2}\right)$

$= 2 + \sqrt{6} = \boxed{2 - \sqrt{6}}$

(C) (7 points) Suppose the input to each of the above three LTI systems is zero for $n < 0$ and

$$x[n] = \sin(\frac{\pi}{2.0}n) + 1.0 = \sin(\frac{\pi}{2.0}n) + \cos(0 \cdot n)$$

for $n \geq 0$. Which system, (I, II or III), produced the following plot of the output $y[n]$, and what is the value of C in the plot?

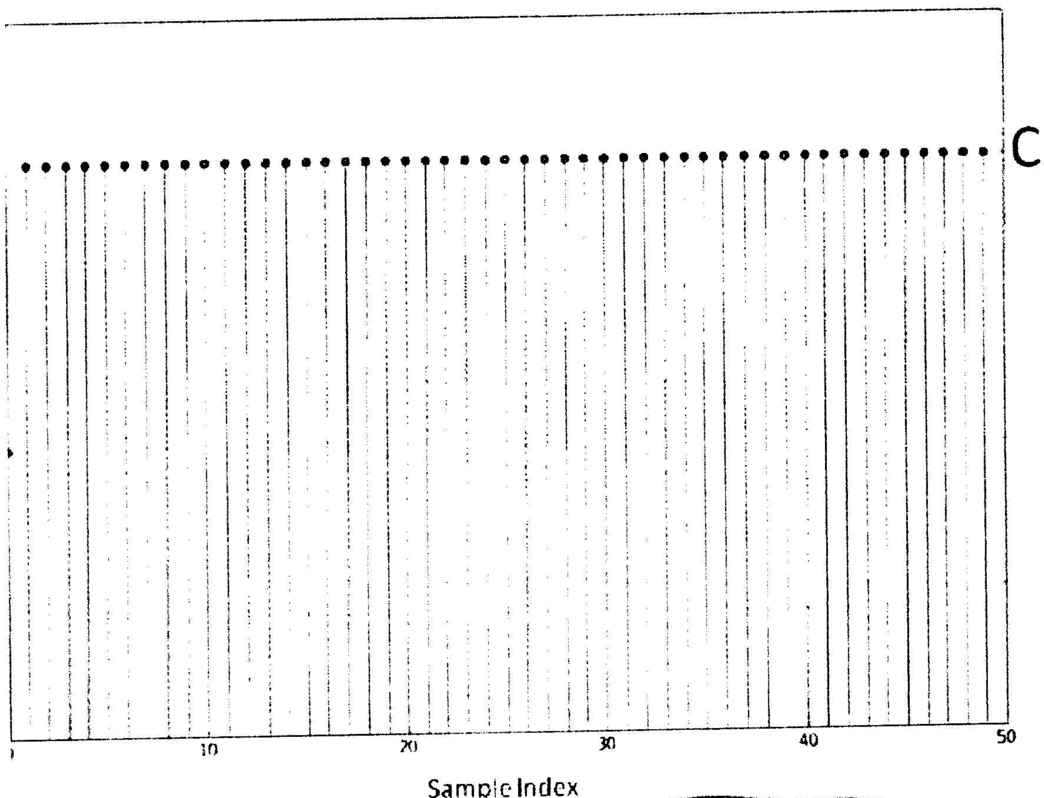

$\boxed{\text{System I}}$ ✓

3/7

C = 1 ×

(D) (10 points) Suppose the input to each of the above three LTI systems is zero for $n < 0$ and

$$x[n] = \sin(\frac{\pi}{2.0}n) + 1.0 = \sin(\frac{\pi}{2.0}n) + \cos(0 \cdot n)$$

for $n \geq 0$. Which system, (I, II or III), produced the following plot of the output $y[n]$, and what, approximately (within 10 percent) is the value of D in the plot?

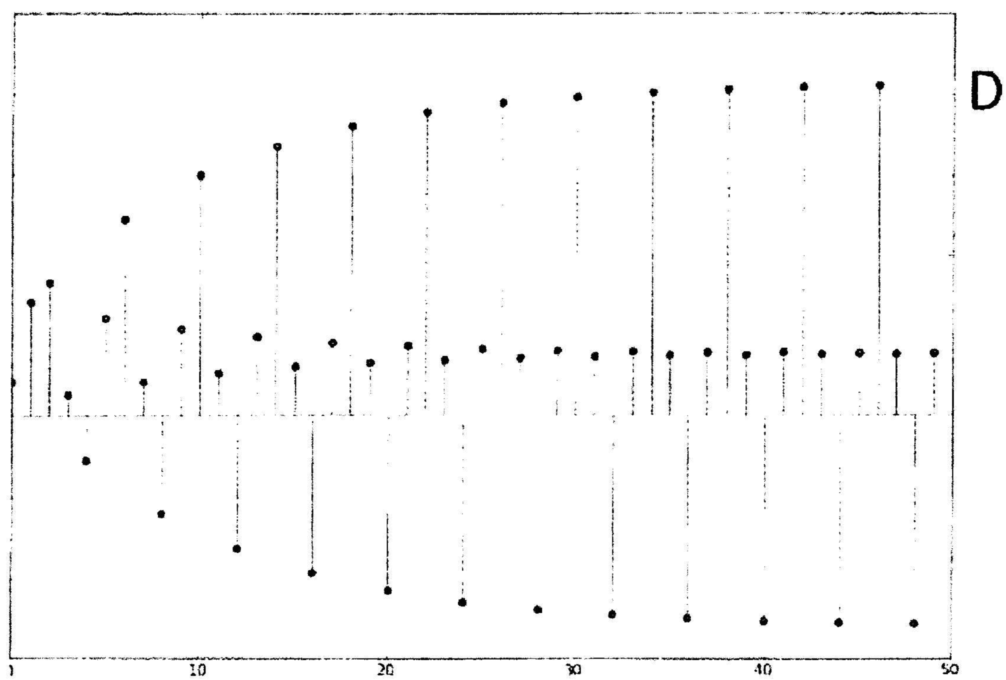

$\boxed{\text{II}}$ X

$\boxed{D = 2.5}$

$\frac{19}{33}$

Problem 4 (20 points)

Consider the simple modulation-demodulation system below, where $\Omega_a = \frac{1000 * 2 * \pi}{10001}$.

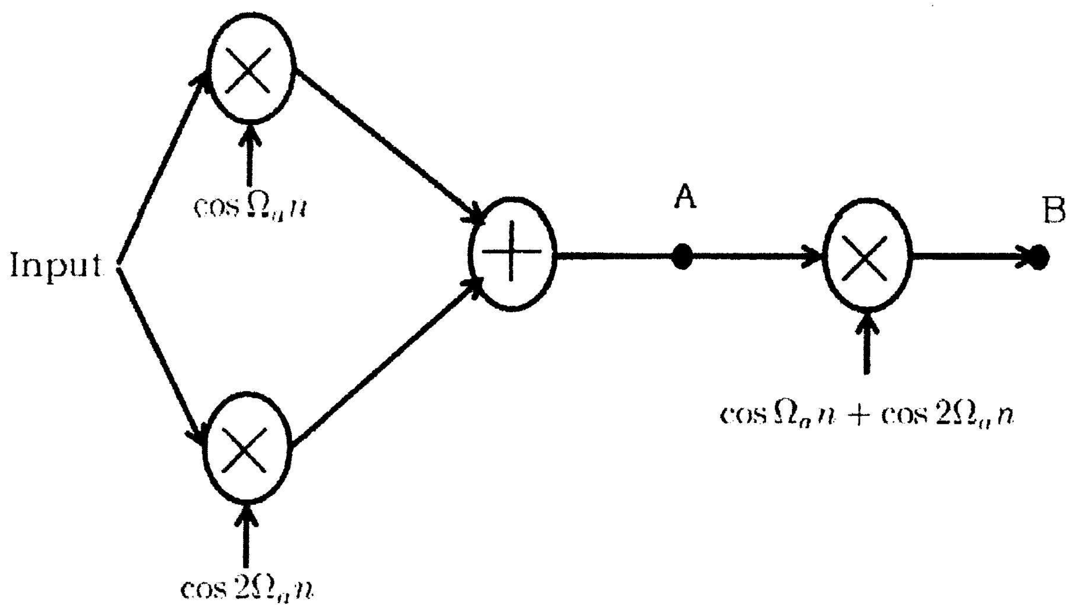

The Fourier Series coefficients for the input to the modulation-demodulation system is plotted below for the case $N = 10001$. Note that the Fourier coefficients are nonzero only for $-400 \leq k \leq 400$.

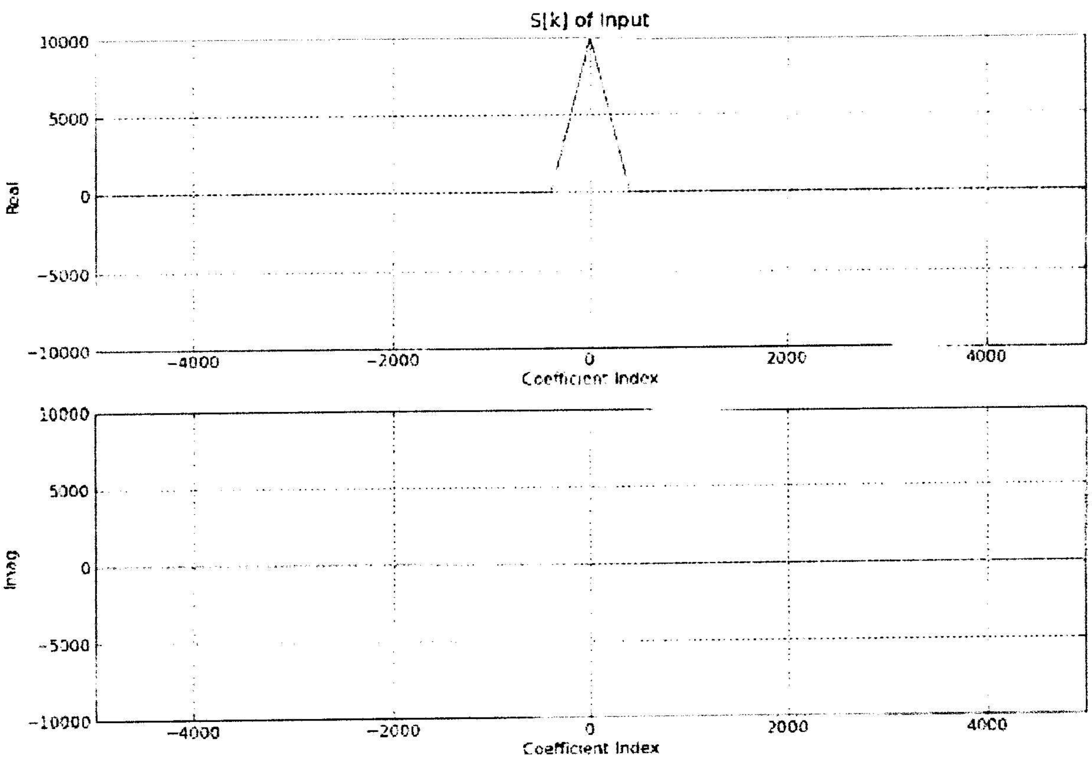

On the two sets of axes below, please plot the Fourier series coefficients for the signal at location **A** and **B** in the above diagram. Be sure to label key features such as values and coefficient indices for peaks.

Plot of Fourier Coefficients of signal at Point A

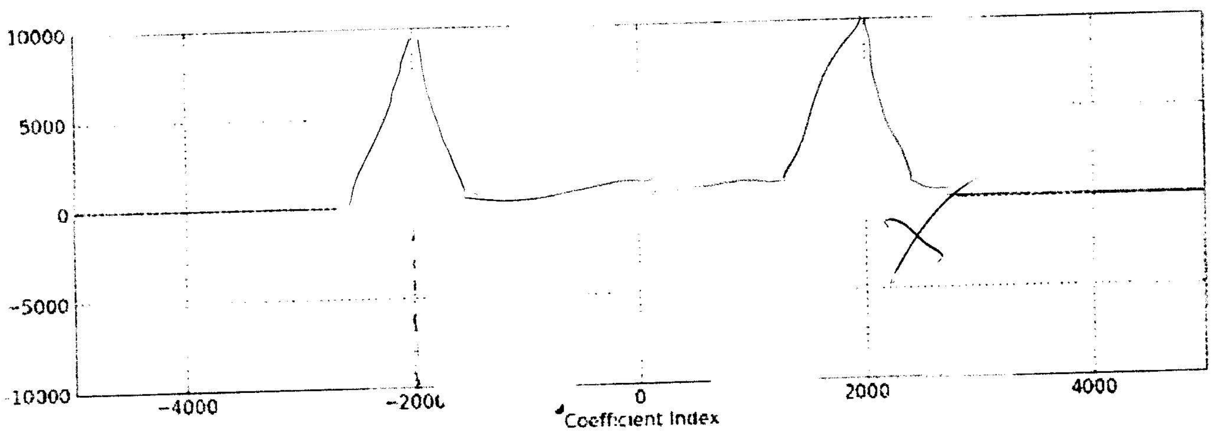

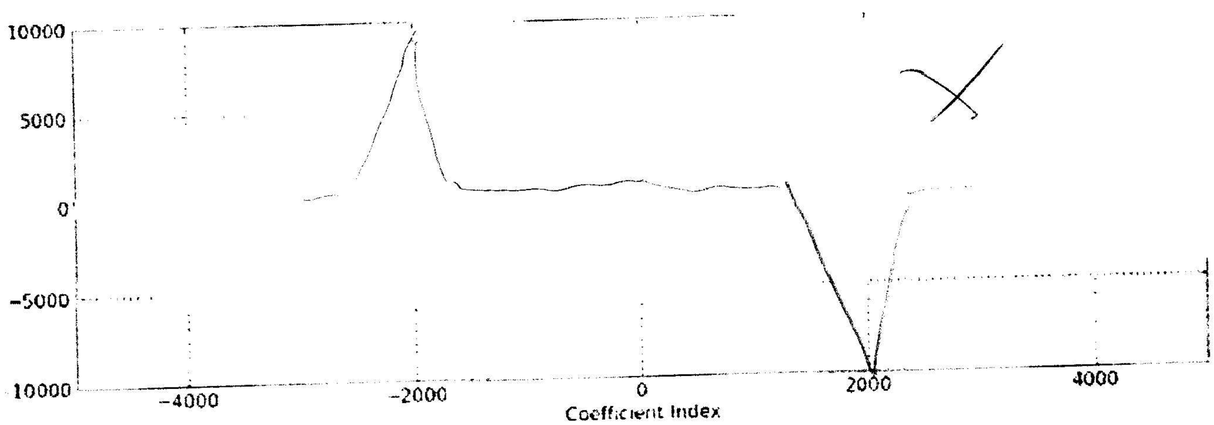

Plot of Fourier Coefficients of signal at Point B

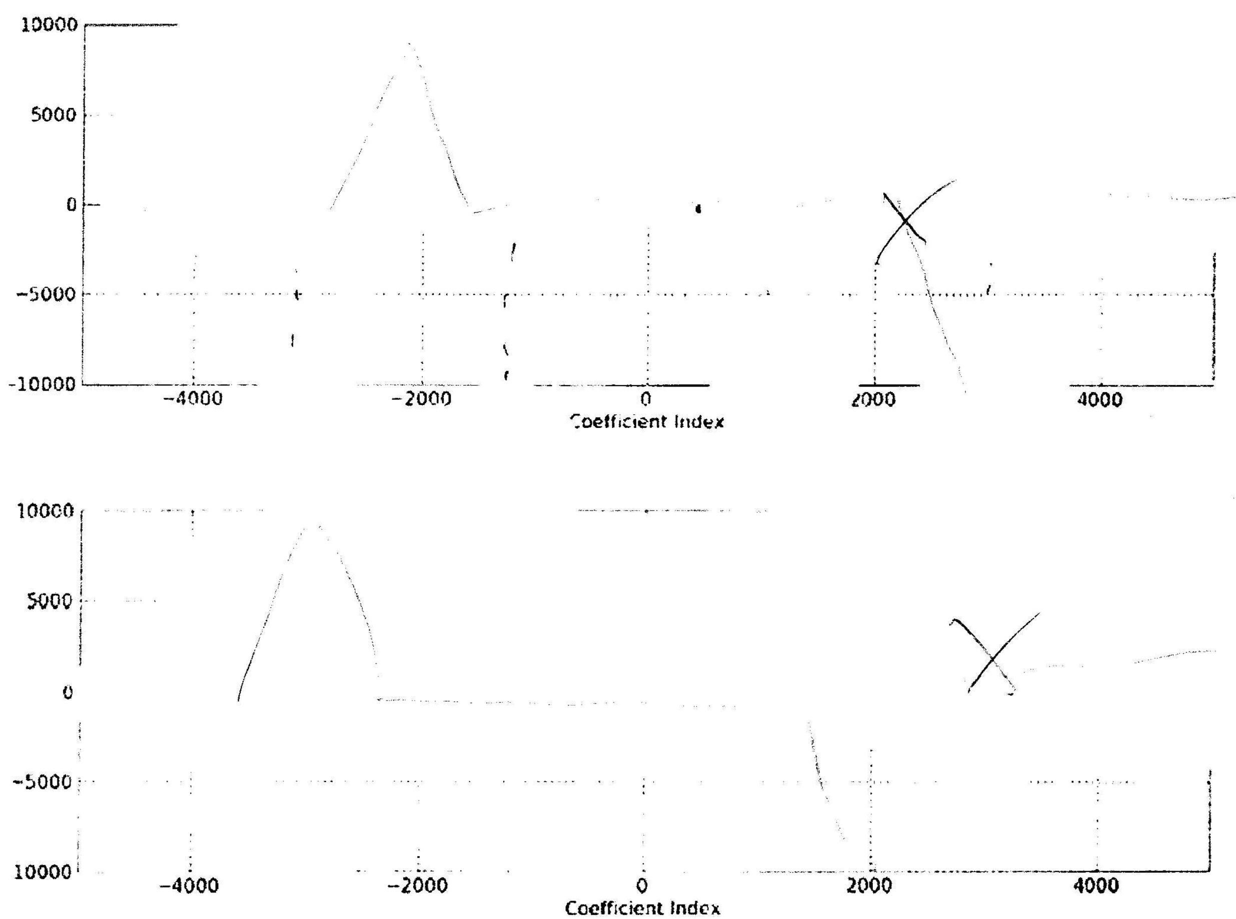

Problem 5 (14 points)

Consider the multiple delay system diagrammed below.

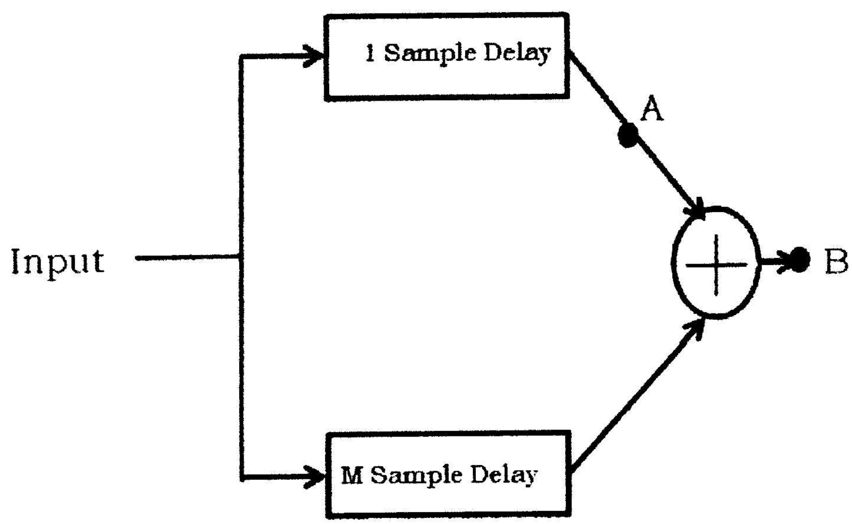

The input to the multiple delay system is a modulated signal that is periodic with period $N = 8001$. The Fourier Series coefficients for this modulated signal are plotted below.

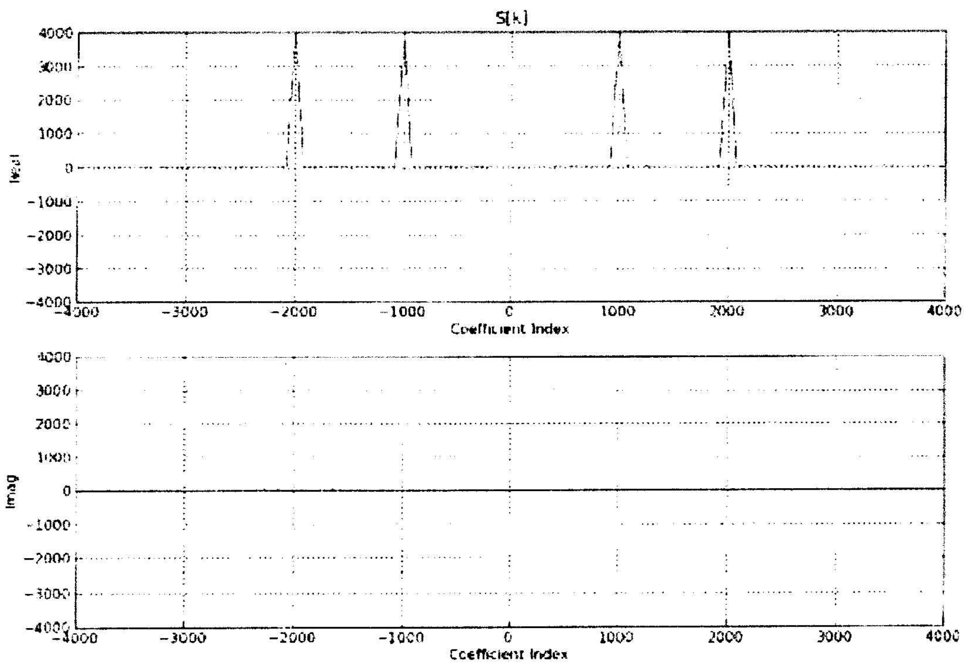

(A) (10 points) Below are plots of the real and imaginary parts of the Fourier coefficients for point A in the multiple delay system. Determine the numerical values for the six peaks in the plots.

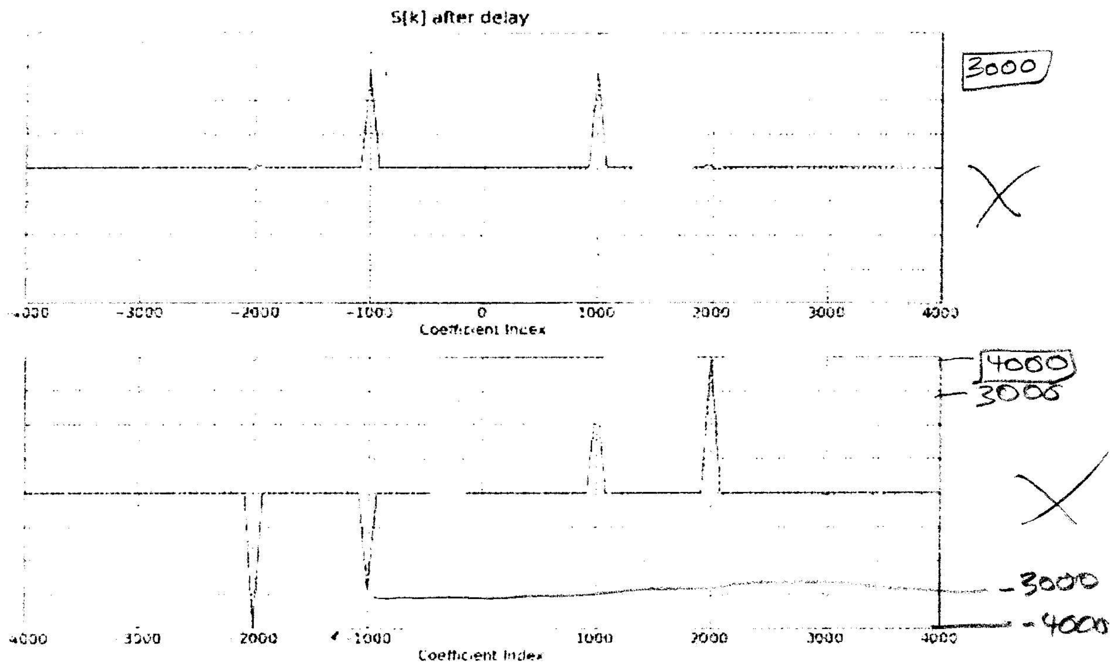

(B) (4 points) Use the following plot of the Fourier series coefficients for the sum of the delayed signals (point **B** in the multiple delay diagram), to determine the smallest integer value for M, the number of samples in the second delay.

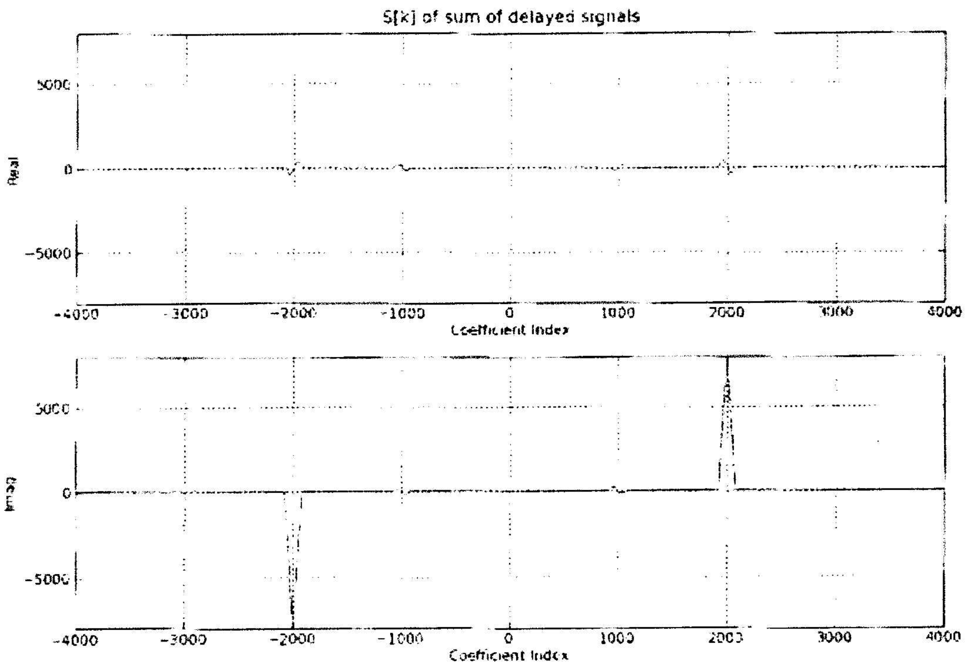

10

X

End of Quiz 2!

GRADING

1) 13 / 15
2) 14 / 14
3) 19 / 33
4) 0 / 20
5) 0 / 14

$\boxed{46/96}$ F

Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.02 Spring 2009

Quiz III

There are 24 questions and 9 pages in this quiz booklet. Answer each question according to the instructions given. You have **120 minutes** to answer the questions.

If you find a question ambiguous, be sure to write down any assumptions you make. **Please be neat and legible.** If we can't understand your answer, we can't give you credit!

Use the empty sides of this booklet if you need scratch space. You may also use them for answers, although you shouldn't need to. *If you do use the blank sides for answers, make sure to clearly say so!*

Before you start, please write your name CLEARLY in the space below.

One two-sided "crib sheet" and calculator allowed. No other notes, books, computers, cell phones, PDAs, information appliances, carrier pigeons carrying answer slips, etc.!

Do not write in the boxes below

1-8 (x/31)	9-14 (x/22)	15-17 (x/14)	18-20(x/19)	21-24 (x/14)	Total (x/100)
24	5	11	9	10	59

Name: ______________________________

I Warmup

1. [2 points]: To stabilize Aloha, a node should ___lower___ its packet transmission probability on a collision and ___increase___ it on a successful transmission.

2

2. [8 points]: A switch uses time division multiplexing (rather than statistical multiplexing) to share a link between four concurrent connections (A, B, C, and D) whose packets arrive in bursts. The link's data rate is 1 packet per time slot. Assume that the switch runs for a very long time.

A. The average packet arrival rates of the four connections (A through D), in packets per time slot, are 0.2, 0.2, 0.1, and 0.1 respectively. The average delays observed at the switch (in time slots) are 10, 10, 5, and 5. What are the average queue lengths of the four queues (A through D) at the switch?

(Answer legibly in the space below.)

Little's Law

$N = \lambda D$

4

A = 2
B = 2
C = 0.5
D = 0.5

B. Connection A's packet arrival rate now changes to 0.4 packets per time slot. All the other connections have the same arrival rates and the switch runs unchanged. What are the average queue lengths of the four queues (A through D) now?

(Answer legibly in the space below.)

~~A = 4~~
B = 2
C = 0.5
D = 0.5

2

3. [8 points]: Under some conditions, a distance vector protocol finding minimum cost paths suffers from the "count-to-infinity" problem.

(Circle True or False for each choice.)

8

A. True / False The count-to-infinity problem may arise in a distance vector protocol when the network gets disconnected.

B. True / False The count-to-infinity problem may arise in a distance vector protocol even when the network never gets disconnected.

C. True / False The "split horizon" technique *always* enables a distance vector protocol to converge without counting to infinity.

D. True / False The "path vector" enhancement to a distance vector protocol *always* enables the protocol to converge without counting to infinity.

4. [3 points]: Which of these statements is true of layering in networks, as discussed in 6.02?
(Circle True or False for each choice.)

1/3

A. True / False The transport layer runs only at the communicating end points and not in the switches on the path between the end points.

B. True / False The same transport layer protocol can run unchanged over a network path with a variety of different link technologies.

C. True / False The lower layers perform error detection on behalf of the higher layers, eliminating the need for higher layers to perform this function.

5. [4 points]: The *exponential weighted moving average* in a reliable transport protocol is:
(Circle True or False for each choice.)

3/4

A. True / False A single-pole low-pass filter to estimate the smoothed round trip time (RTT).

B. True / False A low-pass filter with a pole and a zero to estimate the smoothed RTT.

C. True / False Not necessary if the RTT is constant.

D. True / False Not necessary if the RTT samples are from an unknown Gaussian distribution.

6. [2 points]: In lecture we learned that the JPEG encoding for images was a "lossy" encoding. Which step of the JPEG encoding process loses information?
(Answer legibly in the space below.)

0 ?

7. [2 points]: The human genome consists of approximately 10^9 codons where each codon can be thought of as one of 21 "symbols" coding for one of the twenty amino acids or serving as a "stop" symbol. Give an expression for the number of bits of information in the genome assuming that each codon occurs independently at random with equal probability.
(Answer legibly in the space below.)

$\lg(21)\cdot 10^9 = 4.39\cdot 10^9$ bits ✓

8. [2 points]: You are trying to determine the registration number on Alice's Belize license plate. License plates in Belize have four characters, each either a digit or an upper-case letter, and are selected at random. Alice tells you that her license plate contains only digits. How much information has Alice given you about her license plate? You can give your answer in the form an expression.
(Answer legibly in the space below.)

info $= \lg(M/B)$

$\lg\left(\frac{36^4}{10^4}\right) = 7.4$ bits ✓

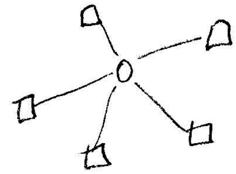

II Aloha

Ben Bitdiddle sets up a shared medium wireless network with one access point and N client nodes. *Unless mentioned otherwise,* assume that the N client nodes are backlogged and that the access point has no packets to send. Each of the N clients wants to send its packets to the access point. The network uses slotted Aloha with each packet fitting exactly in one slot. Recall that each backlogged node in Aloha sends a packet with some probability p.

9. [4 points]: Suppose each node uses the same, fixed value of p. If two or more client nodes transmit in the same slot, a collision occurs at the access point and both packets are lost. The graph below shows the utilization of this protocol as a function of p with N backlogged nodes. Fill in the three blanks shown in the graph below.

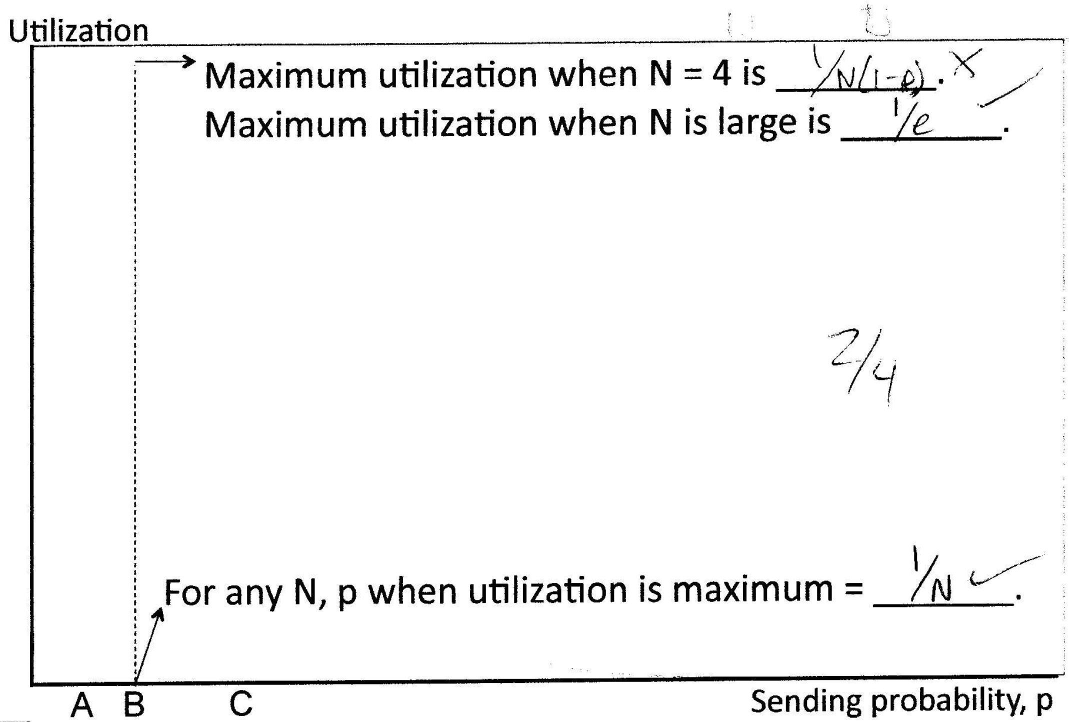

2/4

10. [2 points]: The graph above shows three values of p: A, B, and C. Rank them in *decreasing* order of collision probability.

1. A (lowest) ✓
2. B
3. C (highest)

(-1)

Now assume that the access point is also backlogged and each of its packets is destined for some client. As before, any two or more nodes (including the access point) sending in the same slot causes a collision. Ben sets the transmission probability, p, of each client node to $1/N$ and sets the transmission probability of the access point to a value p_a.

11. [5 points]: What is the utilization of the network in terms of N and p_a? (Note that the client node transmission probability $p = 1/N$.)

(Answer legibly in the space below.)

$1/N \cdot p_a$ X

12. [3 points]: Suppose N is large. What value of p_a ensures that the aggregate throughput of packets received successfully by the N clients is the same as the throughput of the packets received successfully by the access point? Explain your answer.

(Answer legibly in the space below.)

$1/e$ X

From here on, only the client nodes are backlogged—the access point has no packets to send. Each client node sends with probability p (don't assume it is $1/N$).

Ben Bitdiddle comes up with a cool improvement to the receiver at the access point. If exactly one node transmits, then the receiver works as usual and is able to correctly decode the packet. If exactly two nodes transmit, he uses a method to *cancel the interference* caused by each packet on the other, and is (quite remarkably) able to decode both packets correctly.

13. [4 points]: What is the probability, P_2, of *exactly* two of the N nodes transmitting in a slot? Note that we want the probability of *any two* nodes sending in a given slot.

(Answer legibly in the space below.)

$$P_2 = \binom{N}{2} = \frac{1}{2!(N-2)!}$$

✓ 2

14. [4 points]: What is the utilization of slotted Aloha with Ben's receiver modification? Write your answer in terms of N, p, and P_2, where P_2 is defined in the problem above.

(Answer legibly in the space below.)

USE ANSWER FROM 12 as x

$x + P_2$

III Loss-minimizing routing

Ben Bitdiddle has set up a multi-hop wireless network in which he would like to find paths with high probability of packet delivery between any two nodes. His network runs a distance vector protocol similar to what you developed in Lab 9. In Ben's distance vector (BDV) protocol, each node maintains a *metric* to every destination that it knows about in the network. The metric is the node's estimate of the packet success probability along the path between the node and the destination. The packet success probability along a link or path is defined as 1 minus the packet loss probability along the corresponding link or path.

Each node uses the periodic HELLO messages sent by each of its neighbors to estimate the packet loss probability of the link from each neighbor. You may assume that the link loss probabilities are symmetric; i.e., the loss probability of the link from node A to node B is the same as from B to A. Each link L maintains its loss probability in the variable L.`lossprob` and $0 < L$.`lossprob` < 1.

15. [8 points]: The key pieces of the Python code for each node's `integrate()` function in BDV is given below. It has **three missing blanks**. Please fill them in so that the protocol will eventually converge without routing loops to the correct metric at each node. The variables are the same as in Lab 9: `self.routes` is the dictionary of routing entries (mapping destinations to links), `self.getlink(fromnode)` returns the link connecting the node `self` to the node `fromnode`, and the `integrate` procedure runs whenever the node receives an advertisement (`adv`) from node `fromnode`. As in Lab 9, `adv` is a list of (destination, metric) tuples. In the code below, `self.metric` is a dictionary storing the node's current estimate of the routing metric (i.e., the packet success probability) for each known destination.

```
# Process an advertisement from a neighboring node in BDV
def integrate(self, fromnode, adv):
  L = self.getlink(fromnode)
  for (dest, metric) in adv:
    my_metric = metric - L.lossprob  X                  # fill in the blank
    if (dest not in self.routes
       or self.metric[dest] < my_metric                 # fill in the blank
       or self.routes[dest] == L                      ): # fill in the blank
        self.routes[dest] = L
        self.metric[dest] = my_metric

   # rest of integrate() not shown
```

5/8

Ben wants to try out a link-state protocol now. During the flooding step, each node sends out a link-state advertisement comprising its address, an incrementing sequence number, and a list of tuples of the form `(neighbor, lossprob)`, where the `lossprob` is the estimated loss probability to the `neighbor`.

16. [2 points]: Why does the link-state advertisement include a sequence number?

(Answer legibly in the space below.)

Because the LSA is forwarded on each receipt; so without a seqnum old ads might overwrite newer ones and an old ad would never stop echoing through the network.

Ben would like to reuse, without modification, his implementation of Dijkstra's shortest paths algorithm from Lab 9, which takes a map in which the links have non-negative costs and produces a path that minimizes the sum of the costs of the links on the path to each destination.

17. [4 points]: Ben has to transform the `lossprob` information from the LSA to produce link costs so that he can use his Dijkstra implementation without any changes. Which of these transformations will accomplish this goal?

(Circle the BEST answer)

this will sum what should be products.

~~**A.** Use `lossprob` as the link cost.~~

this is negative

B. ~~Use $\frac{1}{\log(1-\text{lossprob})}$ as the link~~ cost.

C. Use $\log \frac{1}{1-\text{lossprob}}$ as the link cost.

✓

~~**D.** Use $\log(1-\text{lossprob})$ as the link cost.~~

$\log(1-\text{lossprob})$

this has cost higher for greater success

IV Reliable data delivery

18. [4 points]: Consider a best-effort network with packet losses and variable delays. Here, Louis Reasoner suggests that the receiver does not need to send the sequence number in the ACK in a correctly implemented stop-and-wait protocol, where the sender sends packet $k + 1$ *only after* the ACK for packet k is received. Explain whether he is correct or not.

(Answer legibly in the space below.)

Correct, because we are waiting for each packet acknowledgement before sending another, there can only ever be one ~~outstanding~~ packet and the receiver will never resend an ACK for a successful transmit

19. [8 points]: The 802.11 (WiFi) link-layer uses a stop-and-wait protocol to improve link reliability. The protocol works as follows:

- **A.** The sender transmits packet $k + 1$ to the receiver as soon as it receives an acknowledgment (ACK) for the packet k. Neither the packet nor the ACK incur any queueing delay.
- **B.** Right after the receiver gets the entire packet, it computes a checksum (CRC). The processing time to compute the CRC is T_p and you may assume that it does not depend on the packet size.
- **C.** If the CRC is correct, the receiver sends a link-layer ACK to the sender. The ACK has negligible size and reaches the sender instantaneously.

The sender and receiver are near each other, so you can ignore the propagation delay. The bit rate is $R = 54$ Megabits/s, the smallest packet size is 540 bits, and the largest packet size is 5,400 bits.

What is the maximum processing time T_p that ensures that the protocol will achieve a throughput of at least 50% of the bit rate of the link in the absence of packet and ACK losses, for any packet size?

(Answer legibly in the space below.)

540 bits =
10 million packets / second

2

5 000 000 T_p

$t_p = 2 \cdot 10^{-7}$ seconds

assuming all packets are 540 bits

20. [7 points]: Consider a sliding window protocol between a sender and a receiver. The receiver should deliver packets reliably and in order to its application.

The sender correctly maintains the following state variables:

`unacked_pkts` – the buffer of unacknowledged packets
`first_unacked` – the lowest unacked sequence number (undefined if all packets have been acked)
`last_unacked` – the highest unacked sequence number (undefined if all packets have been acked)
`last_sent` – the highest sequence number sent so far (whether acknowledged or not)

If the receiver gets a packet that is strictly larger than the next one in sequence, it adds the packet to a buffer if not already present. We want to ensure that the size of this buffer of packets awaiting delivery *never exceeds* a value $W \geq 0$. Write down the check(s) that the sender should perform before sending a new packet in terms of the variables mentioned above that ensure the desired property.

(Answer legibly in the space below.)

(last_sent − first_unacked) <= W

7

V Source coding

21. [6 points]: Consider a Huffman decoding tree for messages made up by randomly choosing one of 10 symbols (each symbol occurs with a non-zero probability). Let ℓ be the length of the encoding for the symbol least likely to occur.

A. What is the minimum possible value for ℓ? Give one example sequence of probabilities for the symbols that produces this minimum.

(Answer legibly in the space below.)

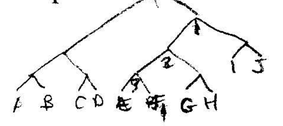

3

$p = \{0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1\}$

$\ell = 4$

B. What is the maximum possible value for ℓ? Give one example sequence of probabilities for the symbols that produces this maximum.

(Answer legibly in the space below.)

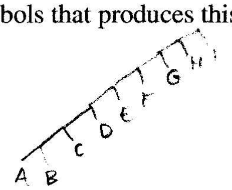

2/3

$\ell = 8$

$p = \{0.5, 0.25, 0.125, \frac{1}{16}, \frac{1}{32}, \frac{1}{64}, \frac{1}{128}, \frac{1}{256}, \frac{1}{512}, \frac{1}{1024}\}$

Alyssa P. Hacker's experiment consists of flipping a biased coin 100 times and encoding the resulting H/T sequence into a message. Each coin flip is independent of the others. The coin has a probability of landing heads, $P(H)$ of 0.75. Alyssa has chosen a Huffman code for encoding pairs of results (HH, HT, TH, TT), so the message is constructed by concatenating the codes for 50 pairs.

22. [4 points]: Give an encoding for each of the four possible pairs that would have resulted from running Huffman's algorithm to produce the decoding tree. Draw the tree in the space below.

Code for HH: 1

Code for HT: 01

Code for TH: 001

Code for TT: 000

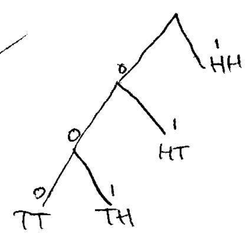

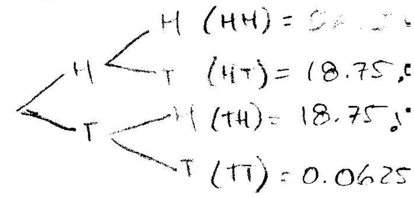

23. [2 points]: Using your code for the previous question, what is the average length of an encoded message obtained from 100 coin flips?

(Answer legibly in the space below.)

56.25% * 1
+ 18.75% * (2 + 3)
+ 6.25% * 3 $= 1.502 * 100 = \boxed{150.2 \text{ bits}}$

1/2 (forgot pairs of flips)

24. [2 points]: Give an expression for the minimum number of bits required to encode the results for a single flip of this biased coin whose $P(H) = 0.75$.

(Answer legibly in the space below.)

$\boxed{2 \text{ bits}}$

X