

Midterm Exam

Directions: Write your name on the exam and on every page you submit. Write something for every question. Students who do not write something for everything lose out over students who write down wild guesses.

P1	40
P2	19.5
P3	20
P4	8
Total	87.5

1, Overview, 40 points: Broad overview, Multiple Choice. 4 points each:

- **1. Layering and Interfaces:** When the data link software in computer receives a packet P , it directly calls the transport layer (through an interface) to ask where P should be placed in memory. The software copies P to the specified buffer, and then hands it over to the routing layer for normal processing. a) This is not a layer violation because neither transport or data link is looking at each other's header b) This is a layer violation because calling between layers should only flow downwards c) This is a layer violation but can easily be fixed by the Data Link calling routing who in turn calls Transport. d) This is a layer violation and cannot be fixed because two non-adjacent layers should not pass information to each other

- **2. Fourier Analysis:** Consider a signal S whose Fourier analysis has been done for you such that $I = 2\sin(\pi t) + \sin(3\pi t)$ What is the value of I at time $1/2$? (In case you need it, recall that π corresponds to 180 degrees, and a sine wave reaches its maximum at 90 degrees, goes back to 0 at 180, then goes negative in symmetric fashion till it comes back to zero at 2π .) a) -1 b) 5 c) 3 d) 2

$$2\sin\left(\frac{\pi}{2}\right) + 3\sin\left(\frac{3\pi}{2}\right) = 2(1) + 3(-1)$$

$$2 - 3 = -1$$

- **3. Media:** Most laptops use 802.11 (WiFi) which uses radio transmission. What would go wrong (or right) if the 802.11 standard decided to use infrared instead of radio? a) The 802.11 bit rate would reduce because Infrared has lower frequency b) The 802.11 bit error rate because Infrared has more noise c) 802.11 would have less interference because neighboring Access Points would not interfere with each other just as 2 infrared keyboards do not interfere with each other in adjacent cubicles. d) 802.11 would not work if there was any obstacle between the sender and receiver

would increase

- **4. Bit Stuffing and Transitions:** Bit stuffing and physical layer coding to ensure transitions are both forms of coding. If you use bit stuffing at the Data Link can you ensure sufficient transitions at the physical layer? a) Yes because we ensure that for every 5 1's we add a 0 b) No, because the data could be all zeroes c) No, because the data could be 010101 . . . and this will confuse the receiver about clock boundaries. d) Yes because we can rely on the data to be sufficiently random
- **5. Clock Recovery:** Phased Locked loops have the following property: a) They quickly correct drift in the receiver clock b) They are fooled by noise spikes that appear to be transitions c) They do not need a preamble to set up synchronization when the frame first starts d) They are slow to adapt but resilient to noise


- **6. Framing:** Why is 00000000 a bad flag for framing regardless of the stuffing rule used?
 - a) Because the data could be a long sequence of 0's and we can't tell the data from the flag
 - b) Because the overhead would be 100 percent using the homework definition of overhead
 - b) Because the overhead would be 50 percent using the homework definition of overhead
 - b) Because if the data is two zeroes it could form a false flag with the first six bits of the false flag

- **7. Latency vs Throughput:** Many modems do time-consuming compression algorithms before they send data so that
 - a) The latency increases and the throughput increases
 - b) The latency decreases and the throughput increases
 - c) The latency increases and the throughput ~~increases~~
 - d) The latency decreases and the throughput ~~decreases~~

- **8. Multiplexing:** When a voice call is made, 64kbps of data bandwidth is reserved across any digital line in the path of the call. Why is strict multiplexing reasonable for a voice call instead of statistical multiplexing?
 - a) Voice calls are not bursty
 - b) Voice calls are bursty
 - c) Voice calls need predictable bandwidth guarantees compared to say email
 - d) Both a) and c)
 - e) Both b) and c)


- **9. Error recovery versus error correction:** On a satellite link, sometimes engineers use error correction instead of error recovery because:
 - a) Satellites have a high error rate and its always good to do error recovery when the error rate is high
 - b) It reduces the latency to recover from a small error
 - c) It costs too much bandwidth to retransmit messages when there is an error as in the example of CDs
 - a) It reduces the overhead for the checksum

- **10. Ethernet and the End-to-end argument:** Ethernets retransmit collided frames hop-by-hop, but let frames lost (due to bit errors) be retransmitted end-to-end. This can be explained as follows in terms of the end-to-end argument.
 - a) Collisions are frequent but bit errors are rare
 - b) Retransmission is an artifact of an older technology (like HDLC) but its not worth it because errors are much less frequent today
 - c) Retransmission of collided frames is needed for correctness.
 - d) Doing retransmission at the Data Link removes the need for a TCP end to end acknowledgement



000
001
010
011
100
101
110
111

INPUT FOR BIT = 1



OUTPUT FOR BIT = 1

2. Nyquist and Shannon Limits Reconsidered, 20 points: Suppose the sender sends bit 1 using 4 Volts for 1 usec, and sends bit 0 by using 0 Volts for 1 usec. Unfortunately, the channel is quite sluggish and the response to a 1 bit takes 3 usec to die out. As shown in the figure, the output of the channel (when a 1 is sent) stays at 4V for 1 usec, goes to 2 Volts for 1 usec, then stays at 1 Volt for 1 usec, before dying to 0 Volts after 3 usec. Sending 0 Volts will result in an output of 0 Volts.

Just to go against tradition, Henry Nyquist (the grandson of Harry Nyquist) decides to send bits once every usec. Thus Henry is deliberately incurring intersymbol interference! However, he feels that a controlled amount of Intersymbol Interference can be detected at the receiver.

Assume the sender is sending a bit every 1 usec. When the voltage of the current bit's output interferes with the voltages of the previous bits, assume that the resulting output is just the sum of the voltages.

- a) By writing down all possible combinations of the current bit and the previous bits that can interfere, show the resulting voltage for each such combination. (8 points)
- b) How can the receiver use the results of a) to correctly detect the current bit despite intersymbol interference. (2 points)
- c) What is the effective bit rate of the sender. Does this violate the Nyquist limit? (5 points)
- d) Cyril Shannon, the grandson of Claude Shannon, hears about Harry's new scheme and is not very impressed. He claims that his grandfather could have done the same thing in a much simpler way. Explain. (5 points)

+8

BIT PATTERN	VOLTAGE
100	1V
010	2V
110	3V
101	5V
011	6V
111	7V

+1.5
b) 0 if voltage < 4V
1 if voltage ≥ 4V -0.5


c) The bit rate of the sender is 1 bit/usec. ✓
The Nyquist limit says that the max signal rate is $2/T$, where T is the period of one bit. In this case, the max rate = $2/1 = 2$ bit/usec. So, no, this does not violate the Nyquist limit. →


Yes, this violates the Nyquist limit because there is intersymbol interference.

d) Instead of sending a V_{ref} every μsec , we could eliminate intersymbol interference altogether while still transmitting 1 bit/ μsec by using levels. For example, we could use the following voltage coding: for some voltage X :

$0 < X \leq 0.5 : 000$
 $0.5 < X \leq 1 : 001$
 $1 < X \leq 1.5 : 010$
 $1.5 < X \leq 2 : 011$
 $2 < X \leq 2.5 : 100$
 $2.5 < X \leq 3 : 101$
 $3 < X \leq 3.5 : 110$
 $3.5 < X \leq 4 : 111$

In this way, we only have to sample once every 3 μsec , but we are coding 3 bits every 3 μsec . Effectively, this is the same as 1 bit / 1 μsec .






3. CRCs, 20 points A sender sends a frame to a receiver over 3 physical lines using the following round-robin strategy. Bit 0 is sent on line 0, bit 1 on line 1, bit 2 on line 2; then bit 3 is sent again on line 0, bit 4 on line 1 etc. Line 0 is flaky and can have burst errors. However, the frame burst error differs from the burst error that a single line sees.


- Suppose Line 0 has a burst error of size 3 that corrupts up to 3 consecutive bits that Line 0 receives. What size frame error burst does the line error burst correspond to? (4 points)
- Suppose Line 0 has a burst error of size 3 and the frame is protected by CRC-8. Argue why the burst error will be detected by the receiver. (6 points)
- Suppose Line 0 has a burst error when 4 consecutive bits are *all* corrupted. Write down the error polynomial assuming the first corrupted bit starts at position 0 in the frame. Will the receiver catch such an error using CRC-8? Assume CRC-8 is $x^8 + x^6 + x^4 + x^2 + 1$. Show your working and the 8-bit remainder you get. Major points are for the process and not for the answer (2 points for error polynomial, 2 points for answer, and 6 points for correct working of the division. If you prefer, please convert the polynomials to bit strings and do the division using bit strings.)

(a) A burst error of size 3 in line 0 corresponds to a burst error of size 7⁺⁴ in the frame.



(b) CRC-8 = $x^8 + x^6 + x^4 + x^2 + 1$ (101010101)⁺⁶
 A burst error of length 7 adds $x^i(x^6 + x^5 + \dots + 1)$ to the message. CRC-8 can catch a burst error of this size because $7 < 8$. This is because any multiple of the CRC-8 generator will have a term of x^7 or higher, so the corrupted result after any 7 burst error will not be divisible by CRC-8.

(c) Error polynomial: $x^9 + x^6 + x^3 + 1$ (1001001001)⁺²



Remainder: 111100011⁺⁶

~~Yes~~ Yes, CRC-8 will catch this error.⁺²

Yes, CRC-8 will catch this error.

4. **Error Recovery, 20 points:** Peter Protocol has been consulting with an Internet Service Provider (ISP) and finds that they use a standard go-back N error recovery protocol shown in the Figure below. As usual, when D(1) is lost, packets D(2) and D(3) are rejected and all three packets must be discarded. The ISP feels that single packet losses are very common and wishes to avoid retransmitting the entire window in such a case.

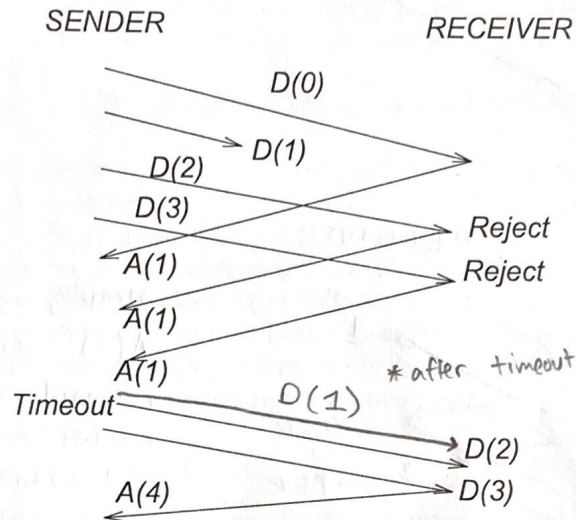
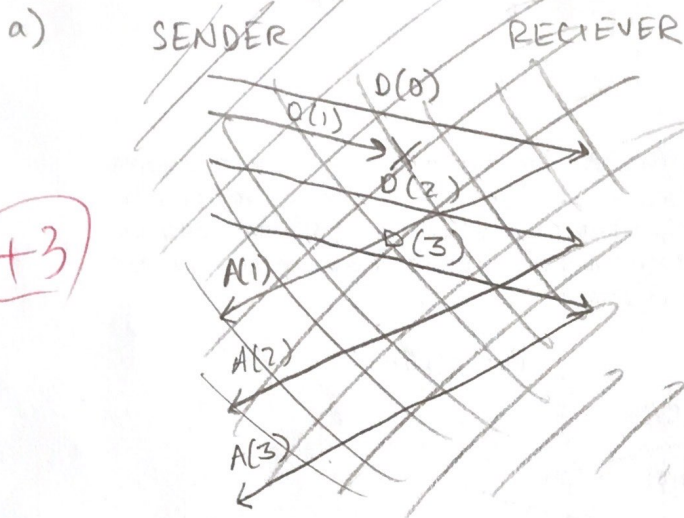
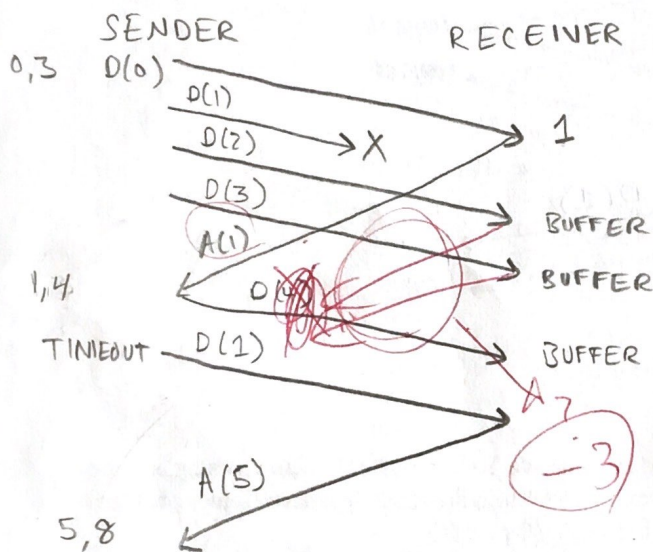


Figure 1:

- a) Peter Protocol first suggests going to a Selective Reject system. Show, using a time-space figure, how the same loss scenario would be handled using a Selective Reject protocol. What change in packet format is needed for this (6 points).
- b) The ISP is unwilling to change its packet formats and asks Peter for another solution. So Peter suggests the following solution: first buffer out-of-order packets; second, when a timeout occurs, only retransmit the oldest unacked sequence number. Will this solution avoid the problem in the above scenario? However, what is a disadvantage of this solution? (7 points).
- c) After some more measurements, the ISP determines that while single packet losses occur commonly, errors that lose a group of packets are also common because of congestion losses. To handle this case, Peter suggests the following alternate protocol. When three or more duplicate acks arrive, the sender retransmits the oldest unacked packet. For example, in the figure, the sender would retransmit D(1), after getting three A(1) packets. On a timeout, retransmission of the entire window is done as usual. What is an advantage of this new scheme over the scheme in b)? Describe your answer in terms of latency and throughput advantages. (7 points).



+3



Simply passing back $A(x)$ from the receiver is not enough. The receiver needs to send back a list/bitmap of packet numbers that it already has, so that the sender knows which remaining packets to send.

b) Yes, this solution will solve the problem in the scenario. However, the disadvantage is that the sender now requires acks for every packet, and the sender needs to have persistent memory of ~~older~~ older packets. which increases latency.

c) An advantage is that the sender no longer needs to keep track of ~~older~~ older packets. This method no longer requires acks for every packet. This method will have lower latency because of this with the same amount of throughput, compared to option (b).