

Welcome to the midterm examination!

Please read everything on this page before you begin.

Every page has two sides – blank pages are explicitly marked as such.

As you should already know, you may not consult any materials during the exam other than the exam itself: no textbooks, no crib sheets, no calculators, etc.

The midterm contains **TEN** problems, including **SIX** short-answer problems and **FOUR** parametrized problems. The value of each problem is indicated below and at the start of each question.

Show all work and reasoning, writing both work and solution legibly, and box all answers. If the course staff cannot read a solution, no credit will be given. All short-answer questions should be stated in **THIRTY-FIVE** words or less; additional words will be ignored when grading. Please write legibly. If we cannot read your answer, then it is wrong!

The **exam will end promptly at 9:00 p.m.** by the clock in the auditorium, or, in the absence or failure of such a clock, by the proctor's watch. Please cease writing immediately when time is called.

Print your name and ID number neatly in the space provided below and print your name and ID in the space provided on every page.

PROBLEM	POINTS	SCORE
1	12	
2	12	
3	6	
4	9	
5	9	
6	13	
7	13	
8	13	
9	13	
Total	100	

Name and Net ID: James Bond_007

This page intentionally left blank.

1. Direct link networks (both sides of sheet) (12 points)

This question has **FOUR** parts. Answer each of the three short answer parts using **TWENTY-FIVE** words or less.

- a) (**3 points**) Consider a point-to-point link between two hosts. Suppose one host sends a signal composing a sum of 2 Hz, 4 Hz, and 7 Hz signals. How fast should the receiving host sample the signal, to correctly capture the signal?

By Nyquist's sampling theorem, we must sample at twice the maximum sending rate, so need to sample at 14 Hz.

- b) (**3 points**) Name the OSI layer or layers in which **the following functions** are addressed in the Internet architecture and state whether they are typically handled in hardware, in software, or in both.

i) **congestion control**

Congestion control is addressed in the OSI Transport layer and handled in software.

Transport Layer (1 point)

Software (1 point)

ii) **CSMA/CD**

CS, MA, and CD are all addressed in the OSI Data Link Layer and handled in hardware.

Data Link Layer (1 point)

Hardware (1 point)

- c) (3 points) In class, we discussed sentinel-based, length-based, and clock-based framing. First, describe (clearly!) what each of these schemes is. Then, give one benefit of sentinel-based over length-based framing.

Sentinel-based: a well-known byte/bit pattern is used to signal end/start of frame

Length-based: the length is given in the packet header

Clock-based: endpoints are synchronized, new fixed-length packet every fixed number of cycles.

Benefit of sentinel-based over length-based: if length field is corrupted, hard to recover. (less packet overhead isn't true, bc sentinel has to include sentinel value)

- d) (3 points) What is CSMA/CA and CSMA/CD? Give a separate definition for each of "CS", "MA", "CA" and "CD". Why do we need CD if we already do CS?

CS: carrier sense (check to see if another host is transmitting)

MA: multiple access (multiple hosts may transmit on the wire)

CA: collision avoidance (eg transmit a signal in advance to make sure no one else starts sending, then send)

CD: collision detection (detect if a collision already occurred)

You need CD even if you already do CS, since there's delay on the line - a collision might not be detected until after you already start transmitting. Note however that CA obviates the need for CD.

2. MAC and Internet architecture (both sides of sheet) (12 points)

- a) **(3 points)** Define “hub”, “switch” and “router” – what is the difference between these devices? Under what circumstances would you prefer a hub over a switch?

hub: used as a repeater to extend networks

switch: can create point to point connections, better for bandwidth savings

router: this is a network-level device, forwards at IP layer (while switch/hub are MAC layer)

- b) **(3 points)** Give 3 benefits of circuit switching over packet switching, and 3 benefits of packet switching over circuit switching.

circuits: *Guaranteed bandwidth, simple Abstraction, Simple forwarding, Low per-packet overhead*

packets: *Statistical multiplexing, Offers “okay” service to everyone, No set-up delay, Routers only store aggregated routes*

- c) **(3 points)** Describe how bridges/switches learn. In particular, describe specifically (a) what data they store, i.e., what fields/information from packets they store (b) the algorithmic process they use to discover/learn that data

when a switch observes a frame from a particular source address, it knows the frame lies out that particular port. so, it stores (source address, port) mappings.

- d) **(3 points)** Why do Ethernet networks use a “spanning tree”? What purpose does it serve?

3. Direct Link Challenges (both sides of sheet) (12 points)

This question has **FOUR** parts. Answer each of the four short answer parts using **TWENTY-FIVE** words or less.

- a) (**3 points**) Name a benefit of “error detection” over “error correction”. Under what circumstances would you design your protocol to perform error detection instead of error correction?

Benefits of Detection

Pro: Overhead only on messages with errors

Con: Cost in bandwidth and latency for retransmissions

Benefits of Correction

Pro: Quick recovery

Con: Overhead on all messages

So, detection is better if low overhead is important, or if there's high cost in bandwidth/latency for retransmissions.

- b) (3 points) In a sliding window protocol, on the sender's side, what is the relationship between the Last Frame Sent (LFS), the Last Acknowledgement Received (LAR), and the Send Window Size?

$LFS - LAR \leq SWS$

(note it's *less than* or equal - but just take off 1 point if they missed the less than)

- c) (3 points) Under what circumstances would you choose the receiver's window size (RWS) to be smaller than the sending window size (SWS)?

Selecting RWS Depends on buffer capacity of receiver. so if your devices have small buffers, you'd prefer to have a small RWS.

d) **(3 points)** In a sliding window protocol, the receiver sends back an ACK for frame 15. What does this mean if the following types of acknowledgements are used?

i) **cumulative acknowledgements?**

All frames up to and including frame 15 have been received correctly

ii) **selective acknowledgement**

Only frame 15 has been received correctly

4. Encoding and Framing Challenges (both sides of sheet) (6 points)

This question has **TWO** parts. Answer each of the three short answer parts using TWENTY-FIVE words or less.

- a) **(3 points)** Why is byte stuffing necessary with some sentinel-based framing schemes? Is a byte of all ones a good choice for a sentinel? Justify your answer.

The sentinel may occur in the data.

one is very common in data.

Would have to stuff a lot (*excessive overhead*).

- b) **(3 points)** Consider a “voting” transmission scheme, where when the sender wants to send a bit X, it sends it four times: XXXX.

- i) What is the hamming distance of this code?

The hamming distance for a code is the **minimum** distance over all pairs of **valid** codewords. In this case: 4 (0000 and 1111 differ by 4)

- ii) How many bit errors can this scheme detect? How many can it correct?

It can correct up to 1 bit error, can detect up to 3 bit errors.

iii) Suppose we want to modify this scheme so that it can **correct** k bit errors. How can we change it? The only change you may make to the scheme is to change the number of bits it sends (i.e., it must still use voting).

If it transmits $2k+1$ bits, it can correct k bit errors.

iv) Suppose we want to modify this scheme so that it can **detect** k bit errors. How can we change it? The only change you may make to the scheme is to change the number of bits it sends (i.e., it must still use voting).

If it transmits $k+1$ bits, it can detect k bit errors.

5. Medium Access Control Challenges (both sides of sheet) (9 points)

This question has **THREE** parts. Answer each of the three short answer parts using **TWENTY-FIVE** words or less.

- a) (**3 points**) Ethernet frames must be at least 64-bytes long. Why did the designers of Ethernet require there to be a maximum frame size?

To ensure that the transmitter is still sending in the event of a collision at the far end of the cable.

- b) (**3 points**) Describe the binary exponential backoff algorithm. Then, list the possible transmission slots a host may use under this algorithm after (i) one, (ii) two, and (iii) seven collisions.

After one collision: slots 0 and 1

After two collisions: slots 0,1,2,3

(in general, 0 through 2^k-1 where k is num of collisions)

After seven collisions, 0,1,..., 2^7-1 .

c) **(3 points)** How does a sender deal with collisions in each of the following schemes:

i) **Collision Avoidance (e.g., CSMA/CA)**

Transceiver sends avoidance signal in advance.

ii) **Collision Detection (e.g., CSMA/CD)**

Transceiver detects collisions after they occur

iii) How does the transceiver detect collisions?

Transceiver sees sum of voltages, looks for voltages impossible for only outgoing signal

iv) **Does Ethernet use acknowledgements at the MAC layer? Why/why not?**

No - absence of collision implies successful transmission.

6. Packet Switching Issues (9 points)

This question has **THREE** parts. Answer each of the three short answer parts using **TWENTY-FIVE** words or less.

- a) **(3 points)** What is head-of-line blocking? How can we change our switch design to mitigate or avoid it from happening?

Packet at head of input queue stalled and blocks packets later in queue that could otherwise cross switch (to distinct output).

Can use virtual queues.

- b) **(3 points)** Explain **two** benefits of forwarding packets with **source routing**.

- (i) hosts can control paths through the network
- (ii) highly flexible, can use any possible path through the network

- c) (**3 points**) Ethernet's spanning tree algorithm repeatedly sends messages over time, even after the tree is computed. Why is this done?

As a heartbeat to verify the tree and make sure none of the links or bridges on the tree have failed.

This page intentionally left blank.

7. Channel Rates and Shared Media (both sides of sheet) (13 points)

You are entrusted with the design of a network to interconnect a set of geographically distributed hosts within your corporation. After some research, you narrow the options to two choices, a fiber-based token ring or a copper-based switched network. The pertinent statistics appear in the table below.

Type	Fiber-based token ring	Copper-based switched network
Signal bandwidth	6 GHz	50 MHz
Signal-to-noise ratio at transmitter	26 dB	64 dB
Attenuation rate	2 dB/km	5 dB/km

The longest link in the network in either case is 15 km.

The following may be of use to you:

$$\log_2(0.4) = -1.32, \log_2(3.5) = 1.8, \log_2(1.4) = 0.48, \log_2(2.5) = 1.32$$

$$\log_{10}(0.4) = -0.4, \log_{10}(2.5) = 0.4$$

(a) What link bandwidth is possible according to Shannon's Law

i) (4 points) For the fiber network?

$$S/N = 26\text{dB} - (15\text{km} \cdot 2\text{dB/km}) = -4\text{dB}$$

$$-4\text{ dB} = 10\log_{10}(S/N)$$

$$S/N = \log_{10}^{-1}(-0.4) = .4$$

So, S/N at receiver is .4.

$$\begin{aligned} C &= B \cdot \log_2(1+S/N) \\ &= 6\text{GHz} \cdot \log_2(1+.4) \\ &= 6\text{GHz} \cdot 0.48 \\ &= 2.42\text{ Gbps (in fiber)} \end{aligned}$$

ii) (4 points) For the copper network?

$$S/N = 64\text{dB} - (15\text{km} \cdot 4\text{dB/km}) = 4\text{dB}$$

$$4\text{ dB} = 10\log_{10}(S/N)$$

$$S/N = \log_{10}^{-1}(0.4) = 2.5$$

So, S/N at receiver is 2.5.

$$\begin{aligned} C &= B \cdot \log_2(1+S/N) \\ &= 50\text{MHz} \cdot \log_2(1+2.5) \\ &= 50\text{MHz} \cdot 1.8 \\ &= 90\text{Mbps (in copper)} \end{aligned}$$

- (b) **(2 points)** Assuming that hosts in the copper network can all transmit at their link rate (the values found in part (a)) simultaneously, roughly how many hosts are necessary for the networks to provide equal aggregate bandwidth (the sum of bandwidth for all hosts)?

The formula is $\lceil (\text{fiber_data_rate}/\text{copper_data_rate}) \rceil$, which is $\lceil (2420/90) \rceil = 27$.


Therefore, we would need 27 hosts on the copper network to equal the aggregate bandwidth of the fiber network. (note you must round up!)

- (c) **(3 points)** Using the copper-based network with a 256-point QAM encoding, what modulation rate (baud) is necessary to obtain the bandwidth found in part (a)?

$$\begin{aligned} 90\text{Mbps} / (8 \text{ bits/symbol}) &= 11.25 \text{ M symbols/sec} \\ &= 11.25 \text{ Mbaud} \end{aligned}$$

8. Medium Access Control (both sides of page) (13 points)

Consider the problem of medium access control using CSMA/CD (the algorithm used with Ethernet) for the network shown below. There is exactly **10km** between neighboring stations. Assume that signals travel directly from any sender to all receivers, propagating at the speed of light in a fiber: 2×10^8 m/sec.



- (a) (4 points) If a transmitter sends at **100 Mbps**, how long must packets be to guarantee collision detection by the transmitter?

Longest link = $10 \times 7 = 70$ km, so maximum one-way delay is:

$$\begin{aligned} (70 \times 10^3 \text{ m}) / (2 \times 10^8 \text{ m/s}) &= 35 \times 10^{-5} \text{ sec} \\ &= 350 \text{ } \mu\text{sec} \end{aligned}$$

Therefore, maximum RTT is $350 \times 2 = 700$ μsec . So at a rate of 100 Mbps, we get:

$$\begin{aligned} (100 \times 10^6 \text{ bps}) * 700 \text{ } \mu\text{sec} &= 80000 \text{ bits} \\ &= 10000 \text{ bytes} \end{aligned}$$


- (b) Divide time into slots the length of the maximum round-trip propagation delay in the network. One packet may be transmitted each time slot. Assume that each of the hosts attempts to **transmit with probability p** in each time slot. What is the probability of a successful transmission in any given slot if

- i) (2 points) $p = 1/2$?

$$8 \times (1/2) \times (1 - 1/2)^7 = 1 / 32$$

- ii) (2 points) $p = 2/3$?

$$8 \times (2/3) \times (1 - 2/3)^7 = 8 \times 2 / 6561 = 16 / 6561$$



For the last part of the problem, assume that each packet requires **1024 bytes** of header/trailer and **15000 bytes** of data. Also assume the probability of successful transmission calculated for part (b)(i) (using $p=1/2$).

(c) (**5 points**) A packet transmission may require more than one slot to complete. Assume that other hosts sense the continuation of the transmission (carrier sense) and simply discard any packets that arrive during those slots. Calculate the average throughput for the network.

$$\lceil 25240 \text{ B} / 10000 \text{ B} \rceil = 3 \text{ slots to transmit}$$

C: competition, T: continue transmission, F: finish transmit

Success = CTF with no collisions

	probability	event	explanation	len	transmitted	prob.
C→C	31/32	C	failure	1	0	31/32
T→F	1	CTF	success	3	1	1/32*1*1
F→C	1					

average length of cycle = $1 * 31/32 + 3 * 1/32$ slots = $34/32 = 1.0625$ slots

average messages sent per cycle = $0 * 31/32 + 1 * 1/32$

average messages sent per slot = $(1/32) / (34/32) = 1/34$

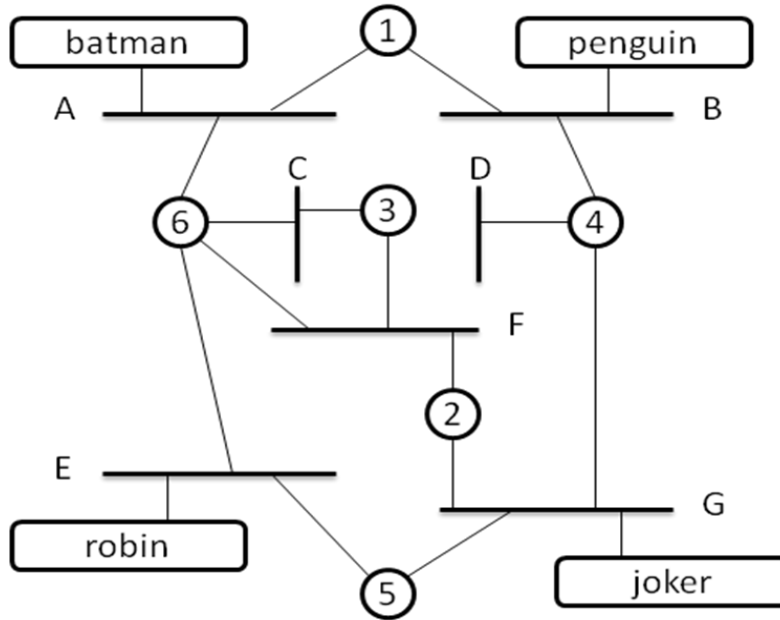
average throughput per slot is $1000 \text{ B} \times 1/34 = 29 \frac{14}{34} \text{ B}$

average throughput = $1000/34 \text{ B/slot time} = 340/6 \text{ B}/(560 \mu\text{sec}) = 52.5 \text{ KBps}$

average throughput = 420 Kbps

9. Spanning Tree Algorithm for Intelligent Bridges (both sides of sheet) (13 points)

The Perlmans spanning tree algorithm and the bridge learning algorithm for forwarding are used for the network shown below.



a) (6 points) Fill in the tables below indicating the root port for each bridge (as a LAN) and the designated bridge for each LAN.

Bridge #	Root Port LAN
1	NA
2	G
3	C
4	B
5	G
6	A

LAN	Designated Bridge
A	1
B	1
C	6
D	4
E	6
F	6
G	4

b) After the spanning tree algorithm has settled:

i) (1 point) Which bridge transmits configuration messages directly onto LAN G?

4

ii) (1 point) Which bridge(s) listen for configuration messages from LAN G?

2, 4, 5

iii) (1 point) Which bridge(s) forward configuration messages heard from LAN G?

4

10.Bit-Stuffing and Byte-Stuffing (both sides of sheet) (13 points)

Model a particular data stream as a stream of bytes, with each byte taking on any value from **48 to 127** with probability **1/80**, independently of the other bytes.

- a) **(4 points)** Use the following variation of the bit-stuffing protocol discussed in class: Stuff (and also reset the count of consecutive 1's) **after four consecutive 1's** instead of after five. This prevents four consecutive 1's from occurring. What is the average number of **bits** that must be stuffed (inserted) **per bit** in the stream?

count them...

range	binary	binary stuffed	# of stuffed bits
48-63	0011 xxxx	0011 11xx	4
64-79	0100 xxxx	0100 1111	1
80-95	0101 xxxx	0101 111x	2
96-111	0110 xxxx	0110 1111	1
112-127	0111 xxxx	0111 1xxx	8

total # of stuffed bits = 16

average stuffed bits per byte in stream = $16/80$

average stuffed bits per bit in stream = $(16/80) \times (1/8)$

- b) **(2 points)** Using the byte-stuffing protocol in which the DLE character (value 48) must be escaped by stuffing a second DLE byte, what is the average number of **bytes** that must be stuffed **per byte** in the stream?

chance of 48 = $1/80$ -> stuff $1/80$ byte per byte in original stream