

CSE 123: Computer Networks - Spring 2022 - Midterm Exam

Name: _____

Student ID: _____

You can use a single 8.5x11” “cheat sheet” (double sided).
The sheet must be a standard thickness with no staples, tape, etc.

Section	Points
Multiple Choice	40
Short Answer	30
Spanning Tree	20
Feedback	10
Total	100

Multiple Choice (40 points)

Please fill in one circle for the correct answer for each question. Some question have multiple parts, make sure you answer all of the parts.

1. (3 points) The Internet's capability of carrying data for many different applications across many different link types is primarily due to which of the following design decisions made by the Internet's pioneers?
 - ☐ End-to-end design: functionality is implemented in end hosts as much as possible rather than in the lower layers (i.e., in the network).
 - ☐ The IP protocol is the sole protocol in the network layer providing a "narrow waist".
 - ☐ Both of these features provide this capability.
 - ☐ Neither of these features provide this capability.
2. (3 points) Which transport protocol is the most efficient (requires the least amount of header bytes sent) for applications that need to perform short data transfers that fit into a single packet, and also need reliability (to be able to handle a packet loss). An example of this type of application is the Domain Name System (DNS).
 - ☐ Transport Control Protocol (TCP)
 - ☐ Universal Datagram Protocol (UDP)
 - ☐ Both of these protocols are equally efficient for this type of application.
 - ☐ Neither of these protocols can provide this application, this is why DNS is its own protocol.
3. (3 points) Which of the following is an example of a correct use of timeouts in networking protocols?
 - ☐ To tell a bridge using the Spanning Tree Protocol when a new root needs to be elected.
 - ☐ To decide when a transport protocol that uses Go-Back-N should initiate a retransmission.
 - ☐ To tell a learning bridge when to delete a mapping from a host's link-layer address to physical port.
 - ☐ To decide when host A using the Address Resolution Protocol can consider an IP address that was used by host B available for A to use.
 - ☐ All of the above are correct uses of timeouts
4. (3 points) Cumulative acknowledgements provide which of the following information to senders:
 - ☐ When a loss occurs, identifying what sequence number is missing so that packet can be retransmitted.
 - ☐ Informing the sender when they can forget about packet transmissions that have completed (i.e., moving the sender window forward).
 - ☐ After a loss is retransmitted, Informing the sender what packets do not need to be retransmitted (i.e., selective retransmission).
 - ☐ All of the above

Extra space to do work (answers in this area will not be graded)

5. (10 points) Host A has established a TCP connection to Host B. The Maximum Transmission Unit at the link layer results in at most 527 bytes in each TCP segment (i.e., TCP packet). At a single time instant during this connection, the state of the sliding window on the sender is as follows (hint. TCP is a sliding window-based transport protocol where the sequence number is indexed per-byte):

A sending to B: A's LBA is 527 and LBS is 4,998. B's LBR is 527 and LAB is 10,000.

B sending to A: B's LBA is 43 and LBS is 43. A's LBR is 43 and LAB is 120.

The notation for the sender's window state is:

LBA: Last Byte Acknowledged by the receiver

LBS: Last Byte Sent by the sender (note: A re-sent frame does not affect LBS value).

The notation for the receiver's window state is:

LBR: Last in-order Byte received

LAB: Largest Acceptable Byte at the receiver

Answer the following True/False questions about what the sliding window state tells us about the state of the data transfer between A and B.

- (a) (2 points) A is transferring a file that is greater than or equal to 4,913 bytes long to B.
- ☐ True
- ☐ False
- (b) (2 points) A and B are very likely to be randomizing their initial sequence numbers (hint. TCP uses 32-bit sequence numbers).
- ☐ True
- ☐ False
- (c) (2 points) B intends on sending 43 bytes to A, but either has not sent them yet, or does not know if A has received them yet.
- ☐ True
- ☐ False
- (d) (2 points) It is possible that B has received (but not yet acknowledged) 9 TCP segments from A.
- ☐ True
- ☐ False
- (e) (2 points) TCP uses cumulative acknowledgements and fast retransmission. It is possible that A has received three duplicate acknowledgements from B and therefore performed fast retransmission of the TCP segment starting at sequence number 528.
- ☐ True
- ☐ False

Extra space to do work (answers in this area will not be graded)

6. (6 points) The DNS name a Netflix server in Los Angeles serves movies to customers in San Diego is `ipv4_1.lagg0.c003.lax004.ix.nflxvideo.net`. Answer the following questions about what a DNS resolver will need to do to resolve the IPv4 address for this name.
- (a) (2 points) What is the *maximum* number of *Authoritative DNS servers* (DNS servers that are in charge of each part of the name) that will need to be contacted to resolve this name. Hint: root DNS servers are also authoritative DNS servers.
- ☐ 6
 - ☐ 7
 - ☐ 8
 - ☐ 1
 - ☐ None of the answers are correct
- (b) (2 points) Since Netflix is a very popular service at UCSD (e.g., many students enjoy watching movies as a way to relax after doing CSE123 projects), how does that affect *how many* Authoritative DNS servers that *typically* need to be contacted to resolve this name.
- ☐ It does not affect it, DNS always requires contacting all of the authoritative servers.
 - ☐ UCSD's DNS resolvers will need to reach out to every authoritative server, every time this name is resolved because Netflix is such a critical service for UCSD students.
 - ☐ Often no authoritative DNS servers need to be contacted, because UCSD's DNS resolvers have cached the first resolution from the authoritative servers, significantly decreasing the load on them from UCSD students.
 - ☐ None of these answers are correct.
- (c) (2 points) Assume that the `lax004` part of the DNS name indicates the name is resolving a server in Netflix's fourth datacenter in Los Angeles (LAX is the airport code of Los Angeles Airport). How is including the datacenter name in part of the DNS name a good example of the benefits of the hierarchical nature of the DNS system?
- ☐ This hierarchy allows Netflix to reuse the same public IP address prefixes for servers in different datacenters. This saves Netflix money because they only have to purchase one IP space.
 - ☐ They are not using the hierarchical property of DNS resolution to their advantage by including the name of the datacenter as a subdomain.
 - ☐ This hierarchy allows Netflix to run independent authoritative DNS resolvers in each of datacenters that are only in charge of resolving the IP addresses of servers in each datacenter, making them easier to manage and improving failure redundancy.
 - ☐ None of these answers are correct.
7. (4 points) A network engineer has modified a network protocol that only used checksums for error detection to use **both** a 32-bit checksum and a 32-bit Cyclic Redundancy Check (CRC) in each frame. Answer the following questions about what the benefits and costs are of this dual approach:
- (a) (2 points) CRC is calculated with complex polynomial arithmetic and checksum is calculated with a simple 1's complement sum. Therefore, the engineer must add much more hardware (i.e., logic gates and memory) to perform the CRC in addition to the simple checksum.
- ☐ True
 - ☐ False
- (b) (2 points) Including both the CRC and checksum of the same payload will significantly improve the error detection capability of the protocol. Checksums can *only detect individual bit errors*, and CRC can *only detect burst errors*? (i.e., errors of multiple bits in a row).
- ☐ True
 - ☐ False

8. (8 points) The Ethernet protocol uses framing that is much simpler than the bit-stuffing HDLC protocol that you studied in-depth in your homework. Ethernet just uses a unique 8-byte start-of-frame delimiter consisting of two parts: 7-bytes of a repeating 0/1 byte pattern followed by a single byte start-of-frame delimiter (10101011). This 64-bit delimiter is 1010101010101010101010101010101010101010101010101010101010101011. Ethernet knows when the end of a frame is by relying on a minimum empty time gap between frames (minimum length string of zeros). In other words, unlike HDLC, Ethernet *does not* use a single start-of-frame and end-of-frame delimiter byte and it *does not* use bit-stuffing.

Answer the following questions about how Ethernet's framing compares to HDLC. Note: Efficiency in this context refers to the *# of bytes in the original data / # of bytes sent after framing*.

- (a) (2 points) What will HDLC's efficiency be for a payload that is 16-bytes long and all 16 bytes are a pattern that matches either HDLC's start-of-frame or end-of-frame delimiter byte?
- ☐ 3/5 (60%)
 - ☐ 1/3 (33%)
 - ☐ 4/5 (80%)
 - ☐ None of these answers are correct
- (b) (2 points) What will Ethernet's efficiency be for the same payload?
- ☐ 2/3 (66%)
 - ☐ 1/4 (25%)
 - ☐ 1/3 (33%)
 - ☐ None of these answers are correct
- (c) (2 points) HDLC's framing always has superior efficiency compared to Ethernet.
- ☐ True
 - ☐ False
- (d) (2 points) Ethernet is significantly more prone to misidentifying the beginning of frames than HDLC because it relies on detecting a unique 64-bit pattern after a minimum length time gap, rather than looking for a single start-of-frame byte *that is guaranteed to never exist in the payload*.
- ☐ True
 - ☐ False

Extra space to do work (answers in this area will not be graded)

Short Answer (30 points)

Please write clearly and ensure that your answer fits in the box provided.

9. (10 points) A Google server (i.e., a host) will often have millions of established TCP connections to serve webpages to hosts on networks outside of Google's LAN. All of these connections will require substantial resources in the network as well as on the hosts (e.g., clients and servers). List all of the resources that these connections will consume, what protocol features cause them to be consumed, and where they are consumed (e.g., where in the network or on which host). You should consider all layers of the network protocol stack that you have learned about thus far. Only a paragraph or two is needed.

10. (6 points) Suppose your home Internet Service Provider's (ISP) IPv6 network addresses have a prefix of 55:2300/24. They use this prefix to address millions of homes across the country. All of these homes use dozens of the ISP's routers as gateways. They assign your home in California an IPv6 network prefix of 55:2302:1500:0000:0000:0000:0000/48. They assign your neighbor's home a prefix of 55:2302:1700:0000:0000:0000:0000/48.

- (a) (3 points) What is the longest prefix (i.e., most specific prefix) that the ISP's router for these two homes can have in its routing table to route to both of these home networks? Please write the prefix in IPv6 Classless Interdomain Routing (CIDR) notation (XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX/Y): where X's are the 16-bit parts of the address prefix (in hexadecimal format) and the Y is the prefix length in bits. (hint: $16 \times 3 = 48$)

- (b) (3 points) Since IPv6 has a much larger address space compared to IPv4, explain why it might be possible to switch back to Class-based Addressing rather than Classesless Interdomain Routing (CIDR)? What downsides would class-based routing have, if any. For instance, would it still be possible to have a specific route in the ISPs network for the two homes described in the question?

11. (10 points) Answer the following questions based on the following proposed 4-bit hamming codebook.

$\{000000 \rightarrow 00, 010101 \rightarrow 01, 101010 \rightarrow 10, 111111 \rightarrow 11\}$

- (a) (5 points) Can this codebook can detect and correct bit errors? If so, how many bit errors can this codebook handle? Show an example of how it can detect the largest number of bit errors, and an example of how it can correct the largest number of bit errors.

- (b) (2 points) What approach is used to build this codebook? For example, some codebooks are created by adding parity bits to the data bits.

- (c) (3 points) Explain why this codebook does not have optimal efficiency? In other words, why it uses more bits to detect and correct errors than absolutely necessary. *Extra Credit (+5pts): What is a more efficient codebook with the same error detect and correction capability?*

12. (4 points) Suppose a router receives an IP packet from a network with an MTU of 2,000 bytes. The packet has the following fields set in its header:

Length = 200	ID = 441	MF = 1	DF = 0	Offset = 187	TTL = 2
--------------	----------	--------	--------	--------------	---------

- (a) (2 points) What does this packet tell the router about the smallest MTU of the network path that the packet has passed through up to the point where it reached this router? Note that the “offset” indicates position in the current fragment (in bytes/8) and $187 \times 8 = 1496$. The IP header always consumes 20 bytes of the link-layer payload, and the “length” field represents (header + payload).


- (b) (2 points) Now this router has to send this packet to a network with a smaller 100 byte MTU. In this case, the MTU refers to the largest link-layer payload.

For the first fragment, list the fields in their IP headers: Length, ID, MF, DF, offset, and TTL.

Extra space to do work (answers in this area will not be graded)

Spanning Tree Protocol (20 points)

13. (20 points) Consider the LAN topology shown below, where B1-B9 are bridges with their corresponding links numbered.



- (a) (10 points) After the spanning tree algorithm converges, what links will be deactivated? Show your work. Note: ties are broken by using the lower numbered bridge (ties are there more than one path from a LAN to the root bridge that are the same length).

- (b) (5 points) After the spanning tree algorithm converges, what are the configuration messages sent by bridge B2 and B9? Use the following notation for spanning tree messages from bridge BX which claims to be distance d from root bridge BY : “ (Y, d, X) ”. For example, the first message sent by bridge B8 is (B8, 0, B8).

- (c) (5 points) Why is the Spanning Tree Protocol needed in this network? In other words, why does this network have links that do not end up being used?

Extra space to do work (answers in this area will not be graded)

Course feedback (10 points)

Please tear off this page of the exam. Answer now, or bring to class on Wednesday.

- What is one thing about the way networks work that you enjoyed learning about?
- What is one thing that is not working well for you in this class? What can we do to help?
- Is there something you would like the staff to do differently?
- Please estimate the time you spent on the homework and project.
- Please estimate the time you spend reading and attending lecture + listening to the podcasts.