CSC 458/2209 (Section L5101): Computer Networks, Fall 2025

Department of Computer Science, University of Toronto

Midterm Exam – 100 Minutes **Date:** Tuesday, October 21st, 2025

(i) This test has 14 questions (some with multiple parts). Make sure to skim through all the questions before starting. This will help you pace yourself. This exam has 50 points in total, and you have 100 minutes (*i.e.,* 2 minutes per point).

(ii) This exam is closed book and closed notes. You can use a non-programmable calculator.

(iii) Write your answers on this questions paper. Make sure to put your name on this page.

(iv) Show your reasoning clearly. If your reasoning is correct, but your final answer is wrong, you will receive some credit. If you just show the answer without reasoning, and your answer is wrong, you may receive no points.

Part I - Multiple Choice Questions [10 points]

For each of the following questions, only one assertion is correct. Selecting the correct assertion earns you 2 points, while selecting an incorrect assertion deducts 1 point from your total score for this section. If you are unsure of the answer and do not select any assertion, you will neither gain nor lose points. Your total score for this section cannot go below 0.

1. Hamming Distance. In the context of error correction, consider the set of codewords {0000, 0011, 1100, 1111}. What is the Hamming distance of this set?

a) 1
b) 2
c) 3
d) 4

2. Shortest Path Routing. A network administrator increases the OSPF cost of a particular link. What happens to routes using this link?

a) All OSPF routers will ignore the link permanently.
b) Routers will prefer alternate paths with a lower total cost, if available.
c) The link will only be used for forwarding multicast packets.
d) The routing protocol will encounter the "counting-to-infinity" problem.

3. Subnet Size. What is the maximum number of IP addresses that can be assigned in a /28 subnet (i.e., when the network ID has 28 bits)?

a) 8
b) 16
c) 32
d) 128

4. BGP. Why does BGP rely on policies instead of using shortest-path metrics like OSPF?

a) Because shortest-path routing is computationally expensive for large networks.
b) Because BGP operates on an inter-domain level, where routing policies are more important than path length.
c) Because BGP can only store a limited number of routes per prefix.
d) Because BGP routers do not exchange reachability information.

5. Fast Retransmission. Which of the following statements is true about the fast retransmission mechanism in TCP?

a) Fast retransmission is triggered when the sender receives two duplicate ACKs for the same packet.
b) Unlike timeout-based retransmission, fast retransmission helps reduce latency by quickly detecting lost packets.
c) Fast retransmission is part of transport layer's flow control mechanism the prevents the sender from overwhelming the receiver.
d) Fast retransmission does not help much when flows are long.

Part II – Comparisons

Compare the following pairs of terms/concepts very briefly (in at most 2-3 sentences). For each pair, explain the key differences – the context they are defined at, protocol(s) they are related to, when/where they are used, etc.

6. [2 points] DHCP vs. DNS

DNS resolves human-readable domain names into IP addresses, allowing hosts to locate services on the Internet. DHCP dynamically assigns IP addresses and configuration parameters (gateway, subnet mask, DNS server) to hosts when they join a network. DNS is a lookup service, while DHCP is a configuration/assignment protocol.

7. [2 points] Internal Border Gateway Protocol (iBGP) vs. External Border Gateway Protocol (eBGP)

iBGP (Internal BGP) is used to route data within a single Autonomous System (AS) and relies on a Full Mesh or Route Reflectors to prevent loops (Split Horizon), as it does not modify the AS-Path. In contrast, eBGP (External BGP) is used to exchange routes between different Autonomous Systems and prevents loops by rejecting updates that already contain its own AS number in the AS-Path. While eBGP peers are typically directly connected (TTL=1) and change the "next-hop" attribute to themselves, iBGP peers are often multiple hops away (TTL=255) and preserve the original next-hop.

8. [2 points] Point of Presence (POP) vs. Autonomous System (AS)

A PoP (Point of Presence) is a physical location where an ISP or network operator houses equipment—routers, switches, servers—to provide local connectivity and access to its network. An Autonomous System (AS) is a logical administrative domain: a collection of IP prefixes and routers operated under a single organization and using a unified routing policy.

9. [2 points] Maximum Segment Size (MSS) vs. Maximum Transmission Unit

MTU: IP, TCP/IP Headers
MSS: TCP, no TCP/IP Headers

Part III - Longer Questions

10. Fragmentation Overhead. Hosts **A** and **B** are connected through an intermediate router **R**. The maximum transmission unit (MTU) for the link from **A** to **R** is 1500 bytes. We assume IP header is 20 bytes. We ignore the overhead of link and transport layer headers here.

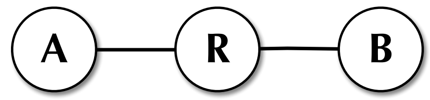

10a) [2 points] If **A** has 1000 bytes of data to send to **B**, how many bytes would **B** receive if the MTU of the link **R** to **B** is 1500 bytes?

1020 : 1000 byes of data + 20 byes header (1 fragment)

10b) [2 points] If **A** has 1000 bytes of data to send to **B**, how many bytes would **B** receive in total if the MTU of the link **R** to B is **500** bytes?

1060 : 1000 byes of data (480 + 480 + 20) + 3 * 20 byes header (3 fragment)

10c) [2 points] If **A** has 1000 bytes of data to send to **B**, how many bytes would **B** receive in total if the MTU of the link **R** to **B** is 100 bytes?

1260 : 1000 byes of data (80 * 12 + 40 * 1) + 13 * 20 byes header (13 fragment)

11. Shortest Path Routing. In the topology shown in the figure, the links are bidirectional (work in both directions) and the number next to each link shows the cost.

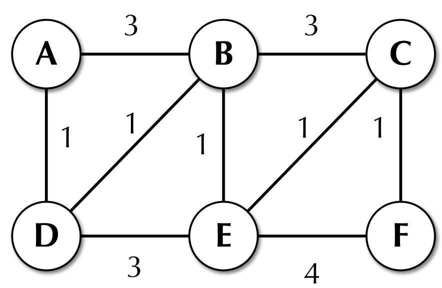

11a) [4 points] Calculate the shortest path and its cost from node **A** to node **F** using Dijkstra's algorithm. Clearly show each step of the algorithm, including the evolution of the shortest-path set, **S**. Write your answer in the table below. Each entry in the second column should show a triple (new router in the shortest path set, *next-hop from **A** to reach the new router*, cost to reach the router) as well as the set **S** at the end of that step. (4 points)

Step	**New entry in shortest path set, S** **(New Router, Next-hop, Cost), S**
1	(A, A, 0), S = {A}
2	(D,D,1), S= {A,D}
3	(B, D,2). S= {A,B,D}
4	(E,D,3), S= {A,B,D,E}
5	(C,D,4), S= {A, D, B, E, C}
6	(F,D,5), S = {A,B,C,D,E,F}

11b) [2 points] In the topology shown above, suppose the link between nodes **E** and **B** fails. Would that impact the shortest path from **A** to **F**? If so, what would be the new shortest path? You do not need to show the steps here, just the final path would suffice. (2 points)

Yes, the failure of the link between E and B would impact the shortest path because the original optimal path utilized that specific link.

New Shortest Path: A -> D -> B -> C -> F

(Note: A -> D -> E -> C -> F is also a valid shortest path with the same cost.)

New Cost: 6

12. TCP Sliding Window Optimization for High-Speed Networks. A TCP connection operates over a 10 GB/s (here, 1 GB is assumed to be 1,000,000,000 bytes) link with an RTT of 100 ms. The MSS is 1500 bytes.

12a) [2 points] Compute the bandwidth-delay product.

1GB

12b) [2 points] Let us assume we have an application which generates an extremely long stream of bytes (i.e., we will not run out of data in the middle of our transmission). The sender uses a sliding window to limit the rate of packet injections to the network. Assuming the size of the sliding window is 1 packet, what is the rate of traffic sent by the sender (in bytes per second)?

15KB/s

12c) [2 points] In class, we saw that a simple stop-and-wait strategy can lead to under-utilization of the network. What is the minimum sliding window size that would lead to a full link utilization? In other words, what should the sliding window size be (in packets) so that the rate of traffic injected by the sender to the network is at least 10 GB/s?

1GB/1.5KB = 666,667

12d) [2 points] Let us assume the sliding window size if 4 packets. For simplicity, we assume the sequence number of packets are 0, 1, 2, 3, … (and not the sequence number of the first byte of the packet as we talked about in class). If the sender has just received an ACK for sequence number 2 (meaning received acknowledges it has received packet number 2), what is the range of sequence numbers the sender can send next? If the sender receives ACKs for 3 and 4, show the new state of the sender's window (which packets can be sent)?

3,4,5,6

5,6,7,8

13. The Mystery of the Missing ARP Requests. Alice, a network engineer at CSC458-Security, is troubleshooting slow communication between two computers, Host **A** (10.0.0.5) and Host **B** (10.0.0.8), on the same subnet (10.0.0.0/24).

She suspects that excessive ARP requests might be flooding the network. To verify this, she captures network traffic and notices the following:

- Each ARP request and reply packet is 28 bytes in size.
- Host A sends an ARP request every 5 seconds to find Host B's MAC address.
- Host B, in turn, does the same for Host A.
- There are 50 computers on the subnet, and all of them communicate with each other in a similar way.

13a) [2 points] Total ARP Traffic per Hour. Calculate the total number of ARP request-reply exchanges happening in the entire network per hour. Also, determine the total ARP traffic in megabytes (MB) generated per hour.

There is a total of 2,450 directed pairs; each pair generates an ARP request every 5 seconds, which gives 720 exchanges per hour per pair, thus 1,764,000 exchanges per hour in total. If we only consider the traffic generated by the host, then an exchange consists of 1 request and 1 reply, which leads to a total of 56 bytes. In total, we will have approximately 98.78MB/h. If we also consider the traffic being broadcasted, then a request will be broadcasted and generate a total of 50 packets * 28 bytes = 1,400 bytes, with a unicast reply from B to switch to A, total of 56 bytes. This gives around 2,568.38 MB (approx. 2.57 GB). Both answers are acceptable depending on the interpretation.

13b) [2 points] Impact on Bandwidth. If the network operates at 100 MB (megabytes) per second, what percentage of bandwidth is consumed by ARP traffic?

Convert the unit and take the ratio, approximately 0.027% if only consider the host generated traffic. 0.71% considering the total network footprint.

13c) [2 points] Optimizing ARP Traffic. Alice decides to increase the ARP cache timeout to 10 minutes instead of 5 seconds. How does this affect the total ARP traffic per hour?

Reduced by 120x from the timeout ratio, relevant calculations are also awarded.

14. Robust Spanning Tree. Consider a network of 7 switches, **S1** to **S7**, as shown in the figure below. These switches run the spanning tree algorithm to eliminate any loops in this network. The root of this tree is the switch with the lowest ID, which is **S1** in this figure.

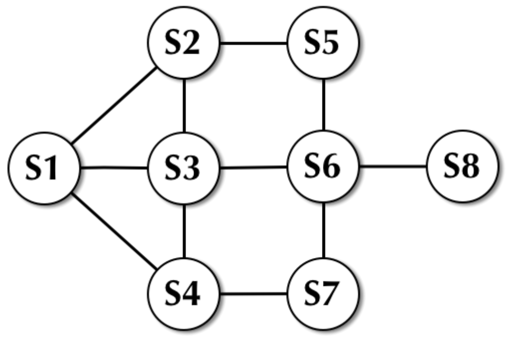

14a) [2 points] Which links are removed as a result of running this algorithm?

S2-S3
S3-S4
S5-S6
S6-S7

14b) [2 points] Let us assume the root of the tree, *i.e.* **S1**, crashes (i.e. it is completely removed from the network). How do other switches realize that the root is dead? How is a new round of spanning tree algorithm initiated?

S1 updates timer expires. Then everybody announces themselves as head, eventually S2 becomes new head

14c) [2 points] Now, let us assume **S1** is running properly, but **S3** fails instead. Here, some switches that are not directly connected to **S1** (for example **S8**) are affected by **S3**'s failure, but since they are not directly connected to **S3**, they have no direct way of realizing **S3**'s failure. How do these switches realize they need to use alternative paths towards the root? In other words, what causes these switches to start a new round of the spanning tree algorithm in this case?

They stop receiving HELLO packets from S3, timer expires. They select new shortest path from other neighbours' HELLOs.

Extra Page 1

Extra Page 2