CSC 458/2209 (Section L0101): 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. CSCMA/CD. In Ethernet using CSMA/CD, what does a sender do immediately after detecting a collision?

a) Waits for the receiver to resend the frame
b) Retransmits the frame immediately without delay
[✓ c)] Sends a jamming signal and then backs off for a random period
d) Aborts the transmission permanently

2. IP Prefixes. What is the subnet mask corresponding to a /24 IPv4 prefix?

[✓ a)] 255.255.255.0
b) 255.255.255.128
c) 255.255.255.192
d) 255.255.255.224

3. Nagel's Algorithm. What is the main purpose of Nagel's algorithm in TCP?

a) To reduce congestion by dropping packets early
b) To improve throughput by sending multiple segments in parallel
[✓ c)] To minimize the number of small packets by buffering and combining them
d) To ensure reliable delivery of packets through selective acknowledgments

4. IP Fragmentation. An IP datagram of size 2000 bytes (including a 20 bytes IP header) must be sent over a link with an MTU of 1000 bytes. How many fragments will be created, and what will be the payload (data) size of each fragment (excluding headers)? You can assume IP header size is 20 bytes and ignore all other layer headers.

a) 2 fragments; each with 980 bytes of payload
b) 2 fragments; both with 1000 bytes of payload
[✓ c)] 3 fragments; first two with 980 bytes of payload, last with 40 bytes of payload
d) 3 fragments; first two with 1000 bytes of payload, last with 20 bytes of payload

5. Routing. Which of the following problems is commonly associated with distance-vector routing protocols?

a) Link state flooding
b) Long path lengths
[✓ c)] Counting to infinity
d) Congestion in the network

Part II – Comparisons [8 points]

Compare the following pairs of terms/concepts very briefly (in at most 3-4 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. Distance Vector vs. Link State

Distance Vector routing (e.g., RIP) uses the Bellman–Ford algorithm, where routers exchange distance information only with neighbors to compute shortest paths. Link State routing (e.g., OSPF) uses Dijkstra's algorithm, where each router floods link information to all others and builds a full network map. Distance Vector is simpler but slower to converge and suffers from count-to-infinity, while Link State converges faster.

7. MAC address vs. IP Address

A MAC address is a permanent, physical identifier assigned to hardware operating at the Data Link Layer (Layer 2), used strictly for local delivery within a single network segment (LAN). In contrast, an IP address is a logical, often temporary identifier operating at the Network Layer (Layer 3), used to route packets across different networks (internetworking). While switches rely on MAC addresses to direct frames to specific devices, routers use IP addresses to determine the path to the destination network.

8. ARP vs. DNS

ARP (Address Resolution Protocol) is used within a local network (LAN) to map a known Network Layer address (IP) to a physical MAC address, operating at the boundary of the Data Link Layer (Layer 2). In contrast, DNS (Domain Name System) is an Application Layer (Layer 7) service used globally to translate human-readable domain names (like www.example.com) into IP addresses. While ARP is essential for the final hop of data delivery between hardware interfaces on the same link, DNS is primarily a directory service used to initiate connections across the internet.

9. Flow Control vs. Congestion Control

Flow control protects the receiver from buffer overflow by limiting the sender with the advertised receive window (rwnd). Congestion control protects the network from overload by limiting the sender with algorithms such as AIMD that determine the congestion window (cwnd). In TCP, the sender's allowed in-flight data is the minimum of the two.

Part III - Longer Questions

10. Packet Switching. Let us consider two end-host **A** and **B** connected through routers **R1** and **R2**. End-host **A** has two packets to transmit to **B**. Each packet (including all headers) is 1000 bytes. The transmission rates and latencies are as follows:

- **A** to **R1**: 1 MB/s (1 MB is 1,000,000 bytes) with 1ms (milliseconds) latency; and
- **R1** to **R2**: 500 KB/s (here we assume 1 KB is 1,000 bytes) with 1ms latency;
- **R2** to **B**: 500 KB/s and 2ms latency.

10a) [2 points] The figure below shows the first packet sent from Host **A** to **R1**. Complete the figure to show all other packets going from **A** to destination **B** in the figure. Please make sure the start and end of transmissions accurately represented (based on the time-axis on top of the figure). Here, we assume the **R1** and **R2** use store-and-forward switching.

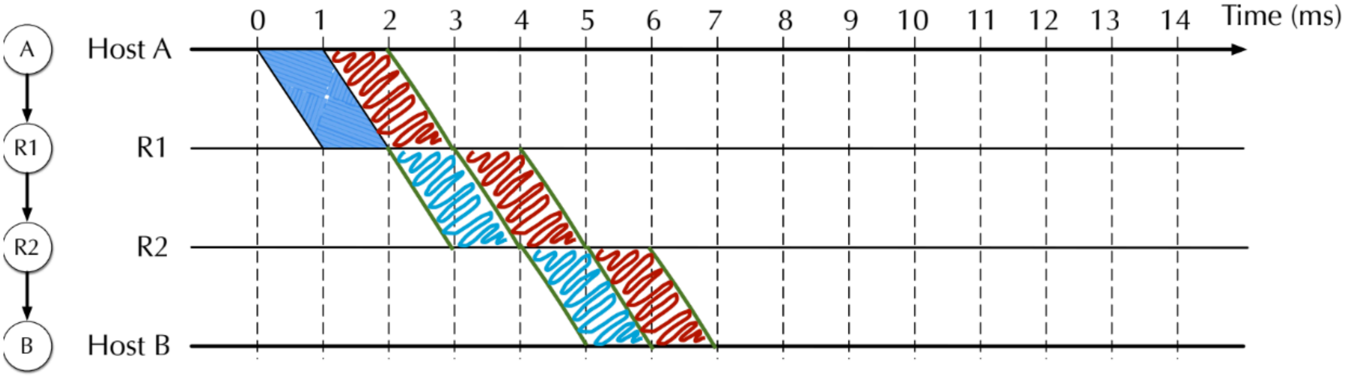

Use the space below to do any calculations you need.

10b) [2 points] Now, let us assume we change the two links (**R1**, **R2**) and (**R2**, **B**) as follows:

- **A** to **R1**: remains the same as before, i.e., 1 MB/s, and 1ms latency.
- **R1** to **R2**: 1 MB/s and 1ms latency.
- **R2** to **B**: 1 MB/s and 1ms latency.

In the figure below show all packets going from **A** to **B**. Again, please make sure the times of departure and arrival of each packet to/from each host and router is accurate. Here, we assume the **R1** and **R2** use store-and-forward switching.

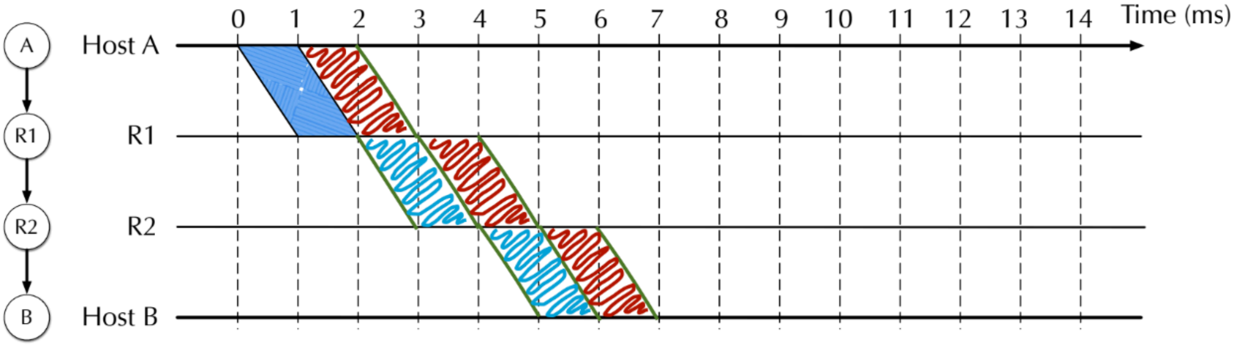

Use the space below to do any calculations you may need.

10c) [2 points] In the scenario described in part **(9b)** how would packet transmissions change if we use cut-through switching at both **R1** and **R2**? Show the packet transmissions in the figure below:

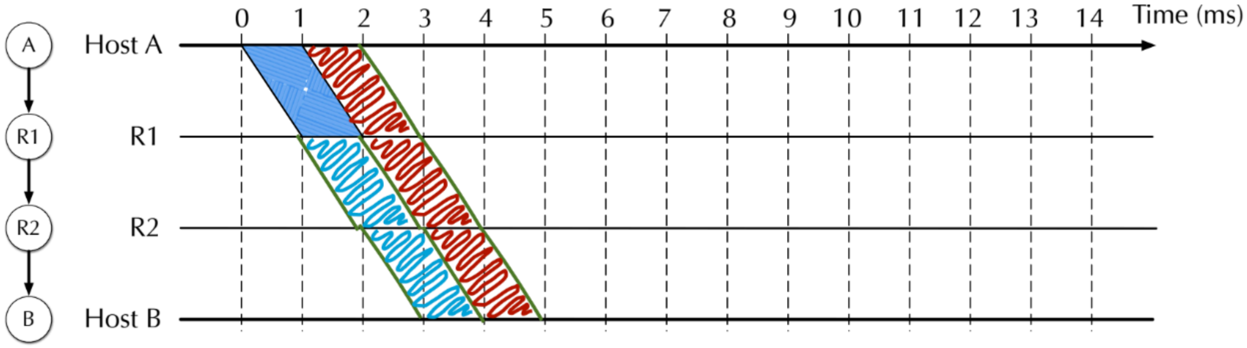

Use the space below to do any calculations you may need.

11) Misbehaving Hosts. Many Internet protocols implicitly assume that end hosts follow protocol rules in good faith. However, when hosts intentionally or unintentionally misbehave, the performance or security of other users can degrade significantly.

11a) [2 points] A malicious host decides to send IP packets with a spoofed source IP address, i.e., the sender uses someone else's IP address. In 2-3 sentences describe two potential negative impacts this can have on the legitimate owner of the spoofed IP address (Be brief, specific and relate to networking behavior.)

The legitimate owner may suffer a Denial of Service (DoS) because the return traffic (backscatter) from the targeted servers will be routed to them, consuming their bandwidth and processing resources. Additionally, the owner may face service blocking or blacklisting, as security systems at the destination will flag the legitimate IP as the source of the attack and block it from accessing services.

11b) [2 points] Consider a host that runs a non-compliant TCP stack, which ignores congestion control and always transmits at maximum rate regardless of network conditions. Why is it challenging for today's routers to deal with such a misbehaving sender?

It is challenging because core Internet routers are designed to be stateless and optimized for high-speed forwarding, meaning they do not track the history or rate of individual flows. To catch a misbehaving sender, a router would need to maintain per-flow state (tracking which packets belong to which connection and calculating their rates), which is computationally expensive and does not scale to the millions of flows a core router handles. Furthermore, the primary tool routers have to signal congestion is dropping packets, which a non-compliant sender will simply ignore (unlike standard TCP which backs off), rendering the punishment ineffective without complex active queue management (AQM).

12. 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.

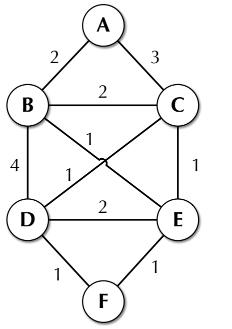

12a) [4 points] Calculate the shortest path and its cost from node **A** to node **F** using Dijkstra's algorithm. Clearly show each step of the algorithm, including the evolution of the shortest-path set, **S**. Write your answer in the table below. Each entry in the second column should show a triple (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	(B, B, 2), S= {A,B}
3	(C, C, 3), S= {A,B,C}
4	(E, B, 3), S= {A,B,C,E}
5	(D, C, 4), S= {A,B,C,E, D}
6	(F, B, 4), S= {A,B,C,E, D, F}

12b) [2 points] Suppose link **B-E** fails. What is the new shortest path from **A** to **F** and its cost?

A,C,D,F - 5

A,C,E,F - 5

12c) [2 points] Now suppose two links fail: **B-E** and **D–E**. What is the new shortest path from **A** to **F**? If multiple shortest paths exist, specify at least one and its cost.

same as b

13. Autonomous System Relationships. Consider eight autonomous systems **AS1**, **AS2**, …, **AS8** shown in the figure. Here, an arrow indicates a customer-provider relationship (i.e. X→Y shows **X** is a provider for **Y**), and a dashed line represents a peer-peer relationship between two autonomous systems (e.g. **AS1** is a provider for **AS2** in our example, and **AS2** and **AS3** are peers).

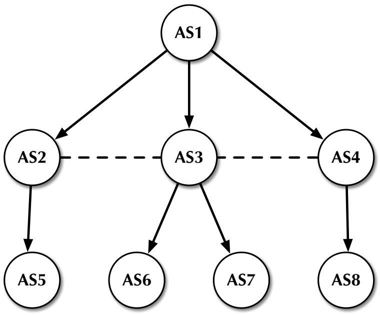

13a) [2 points] Can we have a flow (e.g. a video stream) going through **AS5** to **AS2** to **AS3** and finally to **AS1?** Explain why in 1-2 sentences.

No. The path goes from a peer edge (across) to a provider edge (up), which violates valley-free routing. More precisely, nobody pays transit provider AS3.

13b) [2 points] Can we have a flow going from **AS5** to **AS2** to **AS3** to **AS4** and finally to **AS8**? Explain why in 1-2 sentences.

No, it goes through two peer edges and nobody pays transit provider AS3 as well.

13c) [2 points] If we have a flow that starts at **AS7** and is destined to **AS8**, which one of the two paths will it take:

- **AS7→AS3→AS1→AS4→AS8**; or
- **AS7→AS3→AS4→AS8**

Explain why in 1-2 sentences.

AS7→AS3→AS4→AS8 is chosen because it prioritizes no-cost peering over paid transit, adhering to the fundamental principles of BGP economic routing and valley-free routing.

14. Small Flows. Consider a TCP connection between two applications running on two end-hosts **A** and **B**. Let us assume **A** wants to send a single byte to **B**, and **B** sends only 1 byte of data to **A** in return.

14a) [2 points] In the diagram below show the sequence of ALL packets exchanged between **A** and **B** (including those required to initiate and tear down the connection).

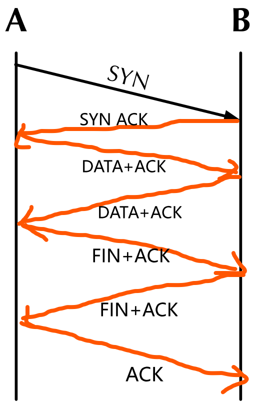

14b) [2 points] Assuming the initial sequence number (ISN) is 0 for both directions, show the sequence number for each packet, and each acknowledgement.

SYN -> SN 0, A None

SYN ACK -> SN 0, A 1

DATA+ACK -> SN 1,A 1

DATA+ACK -> SN 1,A 2

FIN+ACK -> SN 2, A 2

FIN+ACK -> SN 2, A 3

ACK-> SN 3, A 3

14c) [2 points] Assuming the header size is 40 bytes (20 bytes for IP and 20 for TCP), how many bytes in total are exchanged in this process?

7*40 + 2 = 282 B

14d) [2 points] If we used UDP instead, the header size would be 28 bytes (20 bytes for IP and 8 bytes for UDP). How many bytes would be exchanged for A and B in this case?

29*2 = 58B

Extra Page