

CSC175 Data Communications & Networking

Spring 2026 Midterm Exam (Version B Backup)

Student Name: _____ ID: _____

| | |
|-----------------|--|
| Total Points | |
|-----------------|--|

Note: This exam paper should be kept confidential and any dissemination violates copyright.

| | | | | | | | |
|-----|-----|-----|-----|-----|-----|-----|-----|
| Q0 | Q1 | Q2 | Q3 | Q4 | Q5 | Q6 | Q7 |
| /20 | /10 | /10 | /10 | /15 | /15 | /10 | /10 |

Q0 (20 points) Multiple-choice questions: enter your answer keys here:

| | | | | | | | | | |
|----|----|----|----|----|----|----|----|----|----|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| B | B | A | C | B | B | C | B | C | B |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| B | B | A | B | B | B | B | A | C | B |

For the following multiple-choice questions, each question has exactly one correct answer key. If multiple choices are correct, choose the option “All of the above”. Fill in the answer keys in the table above. (Answer keys written in the question area will not be counted.)

1. What distinguishes a Host from a Switch or Router?
 - A) Hosts forward packets; Switches generate data
 - B) Hosts are end machines (like laptops); Switches forward packets toward destinations
 - C) Hosts operate at Layer 1; Switches operate at Layer 4
 - D) Hosts are only for receiving; Switches are only for sending

ANS: B

2. Which of the following is true about Layer 3 (The IP Layer)?
 - A) It strictly uses only wired links
 - B) It connects many local networks to form a global network
 - C) It handles the physical voltage on the wires
 - D) It guarantees the order of packet arrival

ANS: B

3. What term is used to describe a stream of packets exchanged between two endpoints at Layer 4?

- A) Flow (or Connection)
- B) Circuit
- C) Voltage
- D) Spectrum

ANS: A

4. Which of the following is a characteristic of Layer 3 (Best-Effort service)?

- A) Packets are guaranteed to arrive in order
- B) The network notifies you if a packet is lost
- C) Packets are limited in size and could get lost, reordered, or corrupted
- D) The network uses encryption to protect all packets

ANS: C

5. How does demultiplexing at Layer 3 work?

- A) By encrypting packets so only the intended recipient can read them
- B) By adding a header field that tells the receiving end what the next higher-layer protocol is
- C) By splitting large packets into smaller pieces
- D) By routing packets along multiple paths simultaneously

ANS: B

Demultiplexing uses header fields to indicate which Layer 4 protocol (e.g., TCP or UDP) should handle the packet. This allows the IP code to forward the packet to the correct higher-layer protocol handler.

6. What is a "logical port" in networking?

- A) The physical cable connector on a network interface card
- B) A number identifying an application, existing in software
- C) A dedicated circuit established between two hosts
- D) The maximum speed a link can transmit data

ANS: B

A logical port is a software-based number (0-65535) that uniquely identifies an application or connection on a computer. This is distinct from a physical port, which is the hardware connector where cables are plugged in.

7. Why do well-known ports (like port 80 for HTTP) have fixed numbers?

- A) To reduce the size of packets
- B) To increase network security
- C) So that clients can find the server without pre-arrangement
- D) To limit the number of simultaneous connections

ANS: C

8. What is the bandwidth-delay product?

- A) The total number of bits transmitted in one second
- B) The capacity of the link (how many bits fit in the link at a given instant)

- C) The maximum packet size that can be transmitted
- D) The ratio of transmission time to propagation time

ANS: B

The bandwidth-delay product equals bandwidth \times propagation delay and represents the capacity of the link—the total number of bits that can be "in flight" on the link at any given moment.

9. For an 800-bit packet on a link with bandwidth = 1 Mbps and propagation delay = 1 ms, what is the total packet delay?

- A) 0.0008 ms
- B) 0.8 ms
- C) 1.8 ms
- D) 1.0008 ms

ANS: C

Total packet delay = transmission delay + propagation delay. Transmission delay = 800 bits / 1,000,000 bps = 0.0008 seconds = 0.8 ms. Propagation delay = 1 ms. Total = 0.8 + 1.0 = 1.8 ms.

10. What is transmission delay?

- A) The time it takes for a bit to travel across the link
- B) The time it takes to put all bits of a packet into the link
- C) The time a packet waits in a router queue
- D) The total time for a packet to reach its destination

ANS: B

Transmission delay (also called serialization delay) is the time required to transmit an entire packet onto the link. It is calculated as: packet size / bandwidth.

11. When choosing between two links, which factor dominates for a small packet?

- A) Bandwidth
- B) Propagation delay
- C) Packet size
- D) Buffer capacity

ANS: B

For small packets, the transmission delay is negligible compared to propagation delay. Therefore, propagation delay dominates, making a link with lower propagation delay better for small packets, even if it has lower bandwidth.

12. When choosing between two links, which factor dominates for a large packet?

- A) Propagation delay
- B) Bandwidth (which determines transmission delay)
- C) Physical distance
- D) Router processing time

ANS: B

13. What is the 'Count-to-Infinity' problem?

- A) A loop where routing updates bounce back and forth, incrementing costs indefinitely
- B) When a router runs out of memory
- C) When a packet takes too many hops to reach the destination
- D) When the network bandwidth is infinite

ANS: A

Count-to-Infinity happens when bad news (like a link failure) propagates slowly. Routers believe they can reach the destination through each other, creating a loop where the cost keeps increasing.

14. Which technique prevents a routing loop between two immediate neighbors?

- A) Poisoned Reverse
- B) Infinite Bandwidth
- C) Manual Configuration
- D) Using IPv6

ANS: A

Poisoned Reverse prevents 2-node loops. If Z routes through Y to reach X, Z tells Y that its distance to X is infinity, so Y won't try to route back through Z to get to X.

15. Does Poisoned Reverse solve all counting-to-infinity loops?

- A) Yes, it solves all loops permanently
- B) No, it only works for loops involving 2 nodes
- C) No, it only works for loops involving 3 or more nodes
- D) It creates more loops

ANS: B

Poisoned Reverse works for loops between two immediate neighbors. It cannot detect or prevent larger loops involving three or more nodes (e.g., A -> B -> C -> A).

16. What defines the 'state' of a node in Distance-Vector routing?

- A) Its physical location
- B) Its distance vector (costs to all destinations) and its forwarding table
- C) The amount of RAM it has
- D) The number of cables plugged in

ANS: B

The node's state consists of its own distance vector estimates (D_x), the vectors it received from neighbors (D_v), and its local link costs ($c(x,v)$).

17. How does route aggregation (supernetting) help reduce routing table size?

- A) By deleting routes to unpopular websites
- B) By combining multiple adjacent small prefixes into one larger prefix
- C) By converting all IPv4 addresses to IPv6 automatically
- D) By storing tables in the cloud instead of on routers

ANS: B

Aggregation allows a router to advertise a single summary prefix (e.g., 4.0.0.0/8) that covers many smaller, specific networks (like 4.12.0.0/16 and 4.29.0.0/16) owned by the same ISP.

18. What is 'Multi-homing' and how does it affect aggregation?

- A) Connecting to multiple ISPs for redundancy; it breaks perfect aggregation
- B) Using multiple computers at home; it improves aggregation
- C) Moving a laptop between work and home; it has no effect
- D) Hosting multiple websites on one server; it reduces table size

ANS: A

Multi-homing means a network connects to two different providers (e.g., Stanford connecting to AT&T and Verizon). This forces both providers to advertise the specific route to Stanford, preventing them from just summarizing it into their general blocks.

19. If a packet matches two entries in a forwarding table (e.g., 4.0.0.0/8 and 4.29.0.0/16), which one is used?

- A) The first one in the list
- B) The one with the shortest prefix length
- C) The one with the Longest Prefix Match (most specific)
- D) The router randomly picks one to load balance

ANS: C

Routers always use Longest Prefix Match. The /16 entry is more specific (longer match) than the /8 entry, so it is preferred.

20. Approximately how many addresses does IPv4 provide?

- A) 1 million
- B) 4 billion (2^{32})
- C) 340 undecillion (2^{128})
- D) 16 million

ANS: B

IPv4 uses 32-bit addresses, providing 2^{32} unique combinations, which is roughly 4.3 billion addresses. This is not enough for the modern world.

Q1. (10 pts) Links and End-to-End Delay

A packet of size **1,000 bits** is sent from Host A to Host D through three links (store-and-forward). Each link has different transmission rates and propagation speeds.

| Link | Distance | Propagation Speed | Transmission Rate |
|-------|----------|---------------------|-------------------|
| A → B | 1,000 km | 1×10^8 m/s | 1 Mbps |
| B → C | 2,000 km | 2×10^8 m/s | 1 Mbps |
| C → D | 3,000 km | 1×10^8 m/s | 2 Mbps |

Compute the end-to-end delay (transmission + propagation), assuming store-and-forward switching and ignoring processing/queuing delays.

ANS:

Packet size = 1000 bits, so transmission delays:

- 1 Mbps \rightarrow 1000 bits / 10^6 bits/s = 0.001 s = 1 ms
- 2 Mbps \rightarrow 1000 bits / (2×10^6) bits/s = 0.0005 s = 0.5 ms

Propagation delays:

- 1×10^6 m / (1×10^8 m/s) = 0.010 s = 10 ms
- 2×10^6 m / (2×10^8 m/s) = 0.010 s = 10 ms
- 3×10^6 m / (1×10^8 m/s) = 0.030 s = 30 ms


Transmission 1 + 1 + 0.5 = 2.5 ms

Propagation: 10 + 10 + 30 = 50 ms

Total delay = 2.5 + 50 ms = 52.5 ms

Q2. (10 pts) Dijkstra Shortest Paths

Consider an Autonomous System with 5 routers (A, B, C, D, E) connected by bidirectional (undirected) point-to-point links. The link costs are shown in the figure. Use Dijkstra's Algorithm to compute the shortest path tree from the **Source Router A**. Break ties in alphabetical order.



a) List the Visit Order of the routers.

b) Fill out the Routing Table, keeping track of the Shortest Distance (SD) and Previous Node (PN). When a shorter path is found to a router, cross out the old value and write the new one.

| Node | SD | PN |
|------|----|----|
| A | 0 | |
| B | | |
| C | | |
| D | | |
| E | | |


ANS:

a) Visit order: A, D, E, C, B

| Node | SD | PN |
|------|-----|----------------|
| A | 0 | |
| B | 7 4 | A C |
| C | 4 3 | D E |
| D | 1 | A |
| E | 2 | D |

Q3. (10 pts) Distance-Vector Steady State and Failure

Alice (A), Bob (B), Connie (C), and Diego (D) are connected to the local network, which runs the distance-vector algorithm.




a) Fill in R1's table at steady state. If a host is directly connected, the next hop is "Direct".

| To | Next-hop Router | Cost |
|----|-----------------|------|
| | | |


| | | |
|---|--------|---|
| A | Direct | 1 |
| B | | |
| C | | |
| D | | |

b) After the network converges, the R1 - R2 link is broken. Fill in R1's table at the new steady state.



| To | Next-hop Router | Cost |
|----|-----------------|------|
| A | Direct | 1 |
| B | | |
| C | | |
| D | | |

c) After the network converges, the R1 - R4 link is broken. Fill in R1's table at the new steady state.



| To | Next-hop Router | Cost |
|----|-----------------|------|
| A | Direct | 1 |
| B | | |
| C | | |

| | | |
|---|--|--|
| D | | |
|---|--|--|

ANS:

a)

| To | Next-hop Router | Cost |
|----|-----------------|------|
| A | Direct | 1 |
| B | R3 | 3 |
| C | R2 | 2 |
| D | R4 | 3 |

b)

| To | Next-hop Router | Cost |
|----|-----------------|------|
| A | Direct | 1 |
| B | R3 | 3 |
| C | R4 | 7 |
| D | R4 | 3 |


c)

| To | Next-hop Router | Cost |
|----|-----------------|------|
| A | Direct | 1 |
| B | R3 | 3 |
| C | R3 | 8 |
| D | R3 | 4 |

Q4. (15 pts) Link-State

Consider the following network graph with three hosts (A, B, C) and six routers (R1 - R6). For the following questions, assume that the routers run a link-state routing protocol and the routing state has converged. Every link is up unless otherwise noted. When picking between equal-cost paths, the routers pick the route through the neighbor with the lower ID number. For each answer, please provide a concise explanation. Note that

all subparts are independent questions (changes made in one subpart do not affect the subsequent ones).



a) (2 pts) Suppose that the link between **R2 and R4** goes down. R2 and R4 have recomputed their routes, but have not yet sent updates. What route will a packet from A to C take?

b) (3 pts) Suppose that the link between **R4 and R5** goes down. R4 and R5 have recomputed their routes, but have not yet sent updates. What route will a packet from A to C take?

c) (10 pts) Assume that A sends a packet to C, and at time $t = 0$, it arrives at R1. At $t = 1.5$ seconds, the link between R4 and R5 goes down, and R4 and R5 instantaneously recognize and recompute their routes. Assume that link-state advertisements are processed and propagated instantaneously. A link's propagation delay is equal to the link costs in the diagram (in seconds). You can ignore all processing and queuing delays. Does the packet reach its destination? If so, write down the route the packet from A to C takes, and fill in the table for the timing of each event, including both packet arrivals and LSA arrivals at each router or host.

| t (s) | Event |
|-------|----------------------|
| 0 | Packet arrives at R1 |
| | |
| | |

| | |
|--|--|
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |

ANS:

Original path: A → R1 → R2 → R4 → R5 → C

a) A → R1 → R2 → R3 → R4 → R5 → C

b) A → R1 → R2 → R4 → R2 (routing loop between R2 and R4)

c) A → R1 → R2 → R4 → R2 → R6 → R5 → C

| t (s) | Event |
|-------|---|
| 0 | Packet arrives at R1 |
| 1 | Packet arrives at R2 |
| 1.5 | R4-R5 link breaks. R4 reconfigures its route to C as R4-R2-R6-R5-C |
| 3 | Packet arrives at R4 |
| 3.5 | LSA from R4 arrives at R2. R2 reconfigures its route to C as R2-R6-R5-C |
| 5 | Packet arrives at R2 |
| 7 | Packet arrives at R6 |
| 15 | Packet arrives at R5 |

| | |
|----|---------------------|
| 20 | Packet arrives at C |
|----|---------------------|

Q5. (15 pts) Longest Prefix Match

A router has 4 ports, with routing table shown below. Use these rules in order: longest prefix match, then lowest cost, then smallest port number. If no prefix matches, use the **default route on port 4**. Given the routing table below, for each destination IP address, determine **the output port** to forward it, and **the routing table entry** used to make that routing decision. Give a brief explanation for each item.

| Port | Destination | Cost |
|------|-----------------|------|
| 1 | 20.0.0.0/8 | 8 |
| 1 | 180.16.0.0/12 | 11 |
| 1 | 200.10.128.0/17 | 10 |
| 2 | 20.5.0.0/16 | 6 |
| 2 | 200.10.0.0/16 | 9 |
| 3 | 180.16.32.0/19 | 5 |
| 3 | 200.10.192.0/18 | 7 |
| 4 | 20.5.7.0/24 | 6 |
| 4 | 200.10.128.0/17 | 8 |

Some useful binary conversions in the table below.

| Decimal | Binary |
|---------|----------|
| 192 | 11000000 |
| 128 | 10000000 |
| 96 | 01100000 |
| 208 | 11010000 |
| 50 | 00110010 |
| 32 | 00100000 |

| Destination IP | Output port | Matched routing table entry and brief explanation |
|----------------|-------------|---|
| (a) 20.5.7.42 | | |
| (b) 20.6.1.1 | | |

| | | |
|------------------|--|--|
| (c) 180.16.50.7 | | |
| (d) 180.31.2.4 | | |
| (e) 200.10.220.5 | | |

ANS: (Note: although the router's global rule includes lowest cost and smallest port number, those tie-breakers are not needed for any of the five listed destinations.)

| IP | Output port | Routing entry used and brief explanation |
|------------------|-------------|--|
| (a) 20.5.7.42 | Port 4 | 20.5.7.0/24. This address matches 20.0.0.0/8, 20.5.0.0/16, and 20.5.7.0/24, but /24 is the longest and therefore the most specific match. |
| (b) 20.6.1.1 | Port 1 | 20.0.0.0/8. This address starts with 20, so it matches 20.0.0.0/8, but it does not match 20.5.0.0/16 or 20.5.7.0/24 because the second octet is 6, not 5. |
| (c) 180.16.50.7 | Port 3 | 180.16.32.0/19. This address matches 180.16.0.0/12 and also 180.16.32.0/19; the /19 is longer, and its third-octet range is 32–63, which includes 50. |
| (d) 180.31.2.4 | Port 1 | 180.16.0.0/12. A /12 fixes the first 8 bits of 180 and the first 4 bits of the second octet, so 180.16.0.0/12 covers second-octet values 16 through 31; therefore 180.31.2.4 is still inside that prefix. |
| (e) 200.10.220.5 | Port 3 | 200.10.192.0/18. This address matches 200.10.0.0/16, both copies of 200.10.128.0/17, and 200.10.192.0/18, but /18 is the longest match, and it covers third-octet values 192–255, which includes 220. |

More brief explanation style:

(a) 20.5.7.42 → Port 4, using 20.5.7.0/24. It matches /8, /16, and /24, and the /24 entry is the longest prefix.

- (b) 20.6.1.1 → Port 1, using 20.0.0.0/8. Only the /8 entry matches.
- (c) 180.16.50.7 → Port 3, using 180.16.32.0/19. It matches both /12 and /19, and the /19 entry is more specific.
- (d) 180.31.2.4 → Port 1, using 180.16.0.0/12. The /12 block covers second-octet values 16 – 31, so this address is included.
- (e) 200.10.220.5 → Port 3, using 200.10.192.0/18. It matches /16, /17, and /18, and the /18 entry is the longest prefix.

Q6. (10 pts) Route Aggregation

Consider a router running longest prefix matching to forward packets. Given the current routing table, use route aggregation to build a new table with the minimum number of entries, such that both tables produce the same forwarding decisions. Write one IP prefix per box. You may not fill all rows, or you may need to add additional rows, in the merged routing table.

a) (3) Current routing table:

| Destination | Port |
|--------------|------|
| 10.0.16.0/24 | 2 |
| 10.0.17.0/24 | 2 |
| 10.0.18.0/24 | 2 |
| 10.0.19.0/24 | 2 |

Merged routing table:

| Destination | Port |
|-------------|------|
| | |
| | |
| | |

Explanation:

b) (3) Current routing table:

| Destination | Port |
|--------------|------|
| 10.0.16.0/24 | 2 |

| | |
|--------------|---|
| 10.0.17.0/24 | 2 |
| 10.0.18.0/24 | 3 |
| 10.0.19.0/24 | 2 |

Merged routing table:

| Destination | Port |
|-------------|------|
| | |
| | |
| | |

Explanation:

c) (4) Current routing table:

| Destination | Port |
|--------------|------|
| 10.0.16.0/24 | 1 |
| 10.0.17.0/24 | 1 |
| 10.0.18.0/24 | 2 |
| 10.0.19.0/24 | 2 |
| 10.0.20.0/24 | 1 |
| 10.0.21.0/24 | 1 |
| 10.0.22.0/24 | 1 |
| 10.0.23.0/24 | 1 |

Merged routing table

| Destination | Port |
|-------------|------|
| | |
| | |
| | |

Explanation:

ANS:

a) Current routing table:

| Destination | Port |
|--------------|------|
| 10.0.16.0/24 | 2 |
| 10.0.17.0/24 | 2 |
| 10.0.18.0/24 | 2 |
| 10.0.19.0/24 | 2 |

Merged routing table:

| Destination | Port |
|--------------|------|
| 10.0.16.0/22 | 2 |

Explanation: All four /24s share the first 22 bits of the third octet (10.0.16.0 through 10.0.19.255). Aggregate: 10.0.16.0/22 → port 2 (single entry).

b) Current routing table:

| Destination | Port |
|--------------|------|
| 10.0.16.0/24 | 2 |
| 10.0.17.0/24 | 2 |
| 10.0.18.0/24 | 3 |
| 10.0.19.0/24 | 2 |

Merged routing table:

| Destination | Port |
|--------------|------|
| 10.0.16.0/22 | 2 |
| 10.0.18.0/24 | 3 |

Explanation: /22 covers 16–19. Three of them go to port 2; the exception 10.0.18.0/24 → port 3.

c) Current routing table:

| Destination | Port |
|--------------|------|
| 10.0.16.0/24 | 1 |
| 10.0.17.0/24 | 1 |

| | |
|--------------|---|
| 10.0.18.0/24 | 2 |
| 10.0.19.0/24 | 2 |
| 10.0.20.0/24 | 1 |
| 10.0.21.0/24 | 1 |
| 10.0.22.0/24 | 1 |
| 10.0.23.0/24 | 1 |


Merged routing table

| Destination | Port |
|--------------|------|
| 10.0.16.0/21 | 1 |
| 10.0.18.0/23 | 2 |

Explanation: Subnets 16,17,20,21,22,23 → port 1; 18,19 → port 2. Aggregate the 8 /24s 16–23 with /21 = 10.0.16.0/21 → port 1, then override the two going to port 2 with 10.0.18.0/23 → port 2.

Q7. (10 pts) Inter-Domain Routing under Gao-Rexford

Consider the AS graph below, where each AS follows the Gao-Rexford import and export policies. Provider to Customer relationship is denoted by arrows; peer to peer relationship is denoted by horizontal lines with dots. For each source/destination pair, select whether it is possible for packets to be sent from the source AS to the destination AS. Does the AS graph have full reachability from every AS to every other AS? If not, fill out the following table of reachability. (Hint: **Every intermediate AS on a legal path must have at least one customer neighbor along that path.**)



| Source AS | Reachable Destinations |
|-----------|------------------------|
| AS1 | |
| AS2 | |
| AS3 | |

| | |
|-----|--|
| AS4 | |
| AS5 | |

ANS:

| Source AS | Reachable Destinations |
|-----------|------------------------|
| AS1 | AS2 |
| AS2 | AS1, AS3, AS4, AS5 |
| AS3 | AS2, AS4, AS5 |
| AS4 | AS2, AS3, AS5 |
| AS5 | AS2, AS3, AS4 |