CS 168 Summer 2026 — Introduction to the Internet Midterm

Solutions last updated: July 22, 2026

PRINT Your Name: _____

PRINT Your Student ID: _____

You have 110 minutes. There are 7 questions of varying credit. (100 points total)

Question:	1	2	3	4	5	6	7	Total
Points:	19	14	14	11	12	15	15	100

For questions with **circular bubbles**, you may select only one choice.

- Ⓐ Unselected option (Completely unfilled)
- Ⓑ Don't do this (it will be graded as incorrect)
- ● Only one selected option (completely filled)

For questions with **square checkboxes**, you may select one or more choices.

- ■ You can select
- ■ multiple squares
- ☑ Don't do this (it will be graded as incorrect)

Anything you write outside the answer boxes or you ~~cross out~~ will not be graded. If you write multiple answers, your answer is ambiguous, or the bubble/checkbox is not entirely filled in, we will grade the worst interpretation.

Honor Code: Read the honor code below and sign your name.

> I understand that I may not collaborate with anyone else on this exam, or cheat in any way. I am aware of the Berkeley Campus Code of Student Conduct and acknowledge that academic misconduct will be reported to the Center for Student Conduct and may further result in, at minimum, negative points on the exam.

SIGN your name: _____



This content is protected and may not be shared, uploaded, or distributed.

Clarifications

For Q4.1 and Q4.2: The topology has converged before the link goes down.

This content is protected and may not be shared, uploaded, or distributed.

Q1 *Potpourri* **(19 points)**

Q1.1 (2 points) In Project 1: Traceroute, why did we set an unreachable port as the destination?

Ⓐ To send packets directly to each router along the path.

Ⓑ To parse through duplicate packets and trigger a `recv`.

Ⓒ To trigger an ICMP Port Unreachable message from the final host.

Ⓓ To make UDP provide reliable delivery for traceroute probes.

Solution: We send the packet to an unreachable port to receive a definite message that we've finished traceroute.

Q1.2 (1 point) In Project 1: Traceroute, when receiving a packet with a non-empty Options field, we ignore the entire packet.

Ⓐ True Ⓑ False

Solution: If we receive a packet with non-empty Options, we ignore the options and process the packet as normal.

Q1.3 (2 points) In Project 2: Routing, why do we use the `history` data structure? Select all that apply.

[A] Implementing poison reverse.

[B] Tracking which neighbors have received which advertisements.

[C] Implementing incremental updates.

[D] Determining which routes have expired.

Ⓔ None of the above

Solution: The history data structure helps record the most recent advertisement sent out of each port for each destination so we can implement incremental updates.

This content is protected and may not be shared, uploaded, or distributed.

Q1.4 (2 points) A 10 Mbps link uses circuit switching.

- Flow A reserves 6 Mbps but is currently sending at 1 Mbps.
- Flow B reserves 4 Mbps but is currently trying to send at 6 Mbps.

What is the best explanation of what happens?

Ⓐ Flow B can automatically use Flow A's unused 5 Mbps as the link has enough total capacity.

Ⓑ Flow B is limited by its reserved bandwidth and therefore cannot send at 6 Mbps.

Ⓒ Flow A's reservation is automatically lowered to its current sending rate of 1 Mbps.

Ⓓ The capacity of the link automatically increases to accommodate both flows.

Solution: Circuit switching allocates bandwidth per flow reservation, not based on the exact packet demand at each moment. Flow A's unused reserved bandwidth may sit unused while Flow B is still limited by its own reservation.

Q1.5 (1 point) In a statistically multiplexed network, the sum of the peak rates of all flows is generally a better estimate of the required link capacity than the peak of the aggregate traffic.

Ⓐ True Ⓑ False

Solution: Statistical multiplexing by definition allocates capacity corresponding the peak of the aggregate traffic.

Q1.6 (1 point) Routers that receive IP packets must inspect the protocol number to forward the packet to the next hop.

Ⓐ True Ⓑ False

Solution: Only end hosts need to read the protocol number to decide which Layer 4 protocol should be used to process the packet.

Q1.7 (1 point) A router running eBGP advertises routes to all routers within its network.

Ⓐ True Ⓑ False

Solution: eBGP only advertises route updates to BGP routers in neighboring ASes, not to every router in the network. Internal communication within an AS uses iBGP or an internal routing protocol.

This content is protected and may not be shared, uploaded, or distributed.

Q1.8 (2 points) Which is **not** an advantage of link-state protocols, compared to distance-vector protocols?

Ⓐ Routers know the full network topology, so they have more information when choosing paths.

Ⓑ Link-state protocols tend to converge faster than distance-vector protocols.

Ⓒ A router can detect if neighboring routers have a different view of the network topology.

Ⓓ Routers can compute shortest paths without learning the full network topology.

Solution: Link-state works because routers learn the full network topology first then compute shortest paths.

Q1.9 (2 points) Recall the classful addressing scheme from lecture:

Class Name	Class A	Class B	Class C
Number of Network Bits	1	2	3

Which of the following are true about classful addressing? Select all that apply.

[A] The various classes are used to support different numbers of hosts in an organization.

[B] One Class B address space is large enough to support most organizations.

[C] Class A network addresses allow hosts to use more network bandwidth than Class C.

[D] Classful addressing is used for small organizations in the modern Internet.

Ⓔ None of the above

Solution:

Choice A: Class A is used to support large organizations with with 16 million hosts per network; Class B supports 65,000 hosts per network; Class C supports 256 hosts per network.

Choice B: Most organizations have too few hosts to use Class A addressing and too many hosts to use Class C addressing, so most would need to use Class B.

Choice C: IP Addressing is not related identification and not bandwidth usage in a network.

Choice D: Classful addressing is obsolete in the modern Internet.

This content is protected and may not be shared, uploaded, or distributed.

Q1.10 (2 points) Which of the following statements about router hardware are true? Select all that apply.

[A] The data plane computes routes and installs forwarding table entries.

[B] The fabric allows packets to travel between linecards.

[C] The control plane is invoked every time a user packet arrives.

[D] Operators typically interact with routers through the management plane.

(E) None of the above

Solution: The control plane runs routing protocols. The fabric is used to connect linecards together and allow communication. The data plane forwards most packets when they arrive. Operators use the management plane to configure and monitor routers.

Q1.11 (1 point) For TCP timer implementations, every time the sender receives an ACK for a packet that has not yet been ACKed, it will recalculate ETO (expected timeout).

(A) True (B) False

Solution: Only clean samples are used. For example, ACKs on a packet that have been retransmitted are not used since the sender cannot be sure which version the ACK is from.

Q1.12 (2 points) Which of these best explains why TCP often sends a standalone ACK instead of waiting to piggyback the ACK with data?

(A) The operating system does not know when the application will generate data to send.

(B) Piggybacking is only permitted during the TCP three-way handshake.

(C) Piggybacking violates reliability, because ACKs must always be sent in separate packets.

(D) TCP does not support full-duplex communication, so ACKs and data cannot travel together.

Solution: The operating system does not know when the application will have data to send, so waiting to piggyback an ACK could unnecessarily delay it.

This content is protected and may not be shared, uploaded, or distributed.

Q2 *Links* **(14 points)**

Consider the topology to the right, where each link is labeled with its bandwidth and propagation delay.

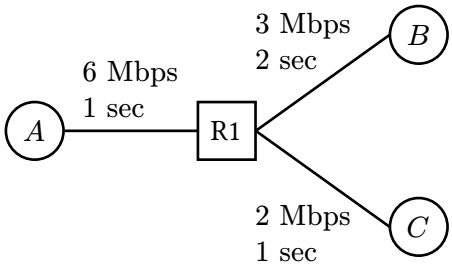

Assumptions:

- A router must receive all bits of a payload before sending that payload along the next link.
- Queue sizes are infinite, and queues are processed in FIFO order.
- A packet is considered to "arrive" at a router/host when the last byte of the packet is received by that router/host.

A sends an endless number of 6 Mbit packets to B, and an endless number of 12 Mbit packets to C. A alternates sending packets to $B, C, B, C, ...$, with each packet transmitted one after the other. R1 only forwards packets out of one link at a time.

Q2.1 (2 points) A sends the first packet to B at $t = 0$ sec. What time does the packet arrive at B?

Ⓐ 3 Ⓑ 4 Ⓒ 5 Ⓓ 6 Ⓔ 7 Ⓕ 8 Ⓖ 9 Ⓗ 10

Solution: The first packet:

- Starts leaving A at $t = 0$.
- Finishes leaving A at $t = 1$. (6 Mbit at 6 Mbps)
- Arrives at R1 at $t = 2$. (Propagation delay 1 sec)
- Starts leaving R1 at $t = 2$.
- Finishes leaving R1 at $t = 4$. (6 Mbit at 3 Mbps = 2 sec)
- Arrives at B at $t = 6$. (Propagation delay 2 sec)

Q2.2 (2 points) What time does the first packet with destination C arrive at C?

Ⓐ 6 Ⓑ 7 Ⓒ 8 Ⓓ 9 Ⓔ 10 Ⓕ 11 Ⓖ 12 Ⓗ 13

Solution: The first packet finishes leaving A at $t = 1$, so the second packet:

- Starts leaving A at $t = 1$.
- Finishes leaving A at $t = 3$. (12 Mbit at 6 Mbps = 2 sec)
- Arrives at R1 at $t = 4$. (Propagation delay 1 sec)
- Starts leaving R1 at $t = 4$.
- Finishes leaving R1 at $t = 10$. (12 Mbit at 2 Mbps = 6 sec)
- Arrives at C at $t = 11$. (Propagation delay 1 sec)

This content is protected and may not be shared, uploaded, or distributed.

Q2.3 (2 points) Which of these best describes the pattern of packets arriving at B and C?

(A) One packet arrives at B, then one packet arrives at C, then this repeats.

(B) One packet arrives at B, then two packets arrive at C, then this repeats.

(C) Two packets arrive at B, then one packet arrives at C, then this repeats.

(D) Not enough information.

Solution: R1 runs a FIFO queue, so the packets arrive at R1 with destinations B, C, B, C, and R1 sends out the packets in the same order.

Even though the overall packet delay to B and C is different, this does not affect the relative ordering of packets.

Q2.4 (2 points) After a long time, which of these best describes the pattern of the queue size at R1?

(A) It is always 0.

(B) It is always between 0 and some non-zero maximum value.

(C) It is always between some non-zero minimum value and some non-zero maximum value.

(D) It grows with no maximum value.

Solution: A bottleneck occurs at R1, because A is constantly sending 6 Mbps into R1, and at most 3 Mbps is flowing out of R1. Since the in-flow is greater than the out-flow, the queue only ever increases.

This content is protected and may not be shared, uploaded, or distributed.

Q2.5 (3 points) After a long time, let b be the total number of **bits** received by B, and let c be the total number of **bits** received by C.

Which of these best represents the approximate relationship between b and c?

Ⓐ $b \approx 1.5c$ Ⓒ $b \approx 2c$ Ⓔ $b \approx 3c$ Ⓖ $b \approx 4c$ Ⓘ $b \approx c$

Ⓑ $c \approx 1.5b$ Ⓓ $c \approx 2b$ Ⓕ $c \approx 3b$ Ⓗ $c \approx 4b$

Ⓙ Not enough information

Solution: First, note that the propagation delays don't factor into this computation because they only shift the receiving of every packet down by a few seconds. In the long run, the total amount of data received is mainly determined by the bandwidth.

If we look at R1 specifically, it alternates between:
- Sending a 6 Mbit packet at 3 Mbps, for 2 sec, towards B.
- Sending a 12 Mbit packet at 2 Mbps, for 6 sec, towards C.

Therefore, within each 8 sec block, twice as much data is sent to C as B. This same 8 sec block repeats over and over, so after a long time, the ratio of data sent to B and C is still $c \approx 2b$.

It is less intuitive, but it turns out the bandwidth also doesn't matter in this answer, because we only care about the total data sent, not how fast it was sent.

Q2.6 (3 points) After a long time, let b' be the amount of **time** that R1 spends forwarding packets to B, and let c' be the amount of **time** that R1 spends forwarding packets to C.

Which of these best represents the approximate relationship between b' and c'?

Ⓐ $b' \approx 1.5c'$ Ⓒ $b' \approx 2c'$ Ⓔ $b' \approx 3c'$ Ⓖ $b' \approx 4c'$ Ⓘ $b' \approx c'$

Ⓑ $c' \approx 1.5b'$ Ⓓ $c' \approx 2b'$ Ⓕ $c' \approx 3b'$ Ⓗ $c' \approx 4b'$

Ⓙ Not enough information

Solution: As seen in the previous subpart, out of each 8 sec block, we have 2 sec toward B and 6 sec toward C. Therefore, $b' = 2/8$ and $c' = 6/8$, and we have $c' \approx 3b'$.

This content is protected and may not be shared, uploaded, or distributed.

Q3 *Distance-Vector* **(14 points)**

Consider running the distance-vector algorithm from lecture on this topology. All unlabeled links cost 1.

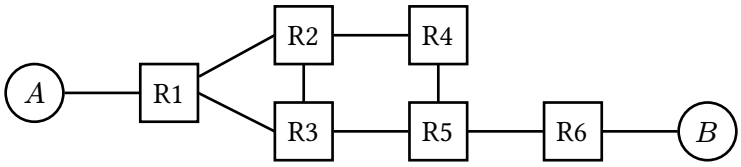

Assume the following unless otherwise stated:

- At $t = 0$, all forwarding tables are empty.
- At $t = 1$, static routes are installed.
- At every time step starting at $t = 2$, each router expires routes, then processes advertisements and updates its table, then sends out advertisements, which all happen instantly.
- All propagation delays are 2 seconds, i.e. an advertisement sent out at $t = 2$ can be processed at $t = 4$.
- Routes expire after 1 second of not hearing advertisements.
- Periodic advertisements are sent every 3 seconds.
- All costs 16 or greater are considered infinite.
- Split horizon, poison reverse, and route poisoning are disabled.

Q3.1 (2 points) What is the first time step where R4 adds A to its forwarding table?

Ⓐ 1 Ⓑ 2 Ⓒ 3 Ⓓ 4 Ⓔ 5 **Ⓕ 6** Ⓖ 7 Ⓗ None

Solution:

- At $t = 0$ all forwarding tables are empty.
- At $t = 1$ static routes are established.
- At $t = 2$ R1 sends out its advertisement for A.
- At $t = 4$ R2 receives the advertisement, updates its table, and sends out its advertisement for A
- At $t = 6$ the advertisement reaches R4 and R4 updates is routing table with an entry for A

Q3.2 (2 points) What is the first time step where all routers simultaneously have an entry for B?

Ⓐ 2 Ⓑ 3 Ⓒ 4 Ⓓ 5 Ⓔ 6 Ⓕ 7 Ⓖ 8 **Ⓗ None**

Solution: The route to B is installed by R6 at $t = 1$.

- At $t = 2$, R6 advertises the route to R5.
- At $t = 4$, R5 receives the advertisement, adds B to its table, and advertises to R4 and R3.
- At $t = 6$, R4 and R3 can learn the route, and advertise it further.
- At $t = 8$, R1 and R2 may receive advertisements.

However, routes expire after 1 second while advertisements are only resent every 3 seconds. Therefore, learned routes disappear before they can propagate through the entire network.

Since R6 is the only router with a permanent static route to B, there is no time where all routers simultaneously have an entry for B.

This content is protected and may not be shared, uploaded, or distributed.

Q3.3 (2 points) What is true about this network after a long time?

[A] The network reaches steady-state and all packets are routed through the least-cost path.

[B] All routers' routing tables will repeatedly lose and repopulate all entries in their tables.

[C] R1 and R6 will always have at least one entry in their routing table.

[D] Reducing all propagation delays to 1 second will help the scheme reach steady state.

(E) None of the above

Solution:

Choice A: Learned routes continuously expire before advertisements can be resent, preventing the network from reaching steady-state.

Choice B: R1 and R6 always retain their static routes, so not all routers loose all their entries

Choice C: Static routes never expire.

Choice D: Triggered updates only help with making advertisements propagate after a change, they do not actually fix the continuous timeout issue.

This content is protected and may not be shared, uploaded, or distributed.

Now consider rerunning the algorithm, starting over at $t = 0$ with empty forwarding tables.

We modify the following assumptions:
- All propagation delays are 1 second, i.e. an advertisement sent out at $t = 2$ can be processed at $t = 3$.
- Routes expire after 5 seconds of not hearing advertisements.

Q3.4 (2 points) What is the first time step that R3's routing table reaches steady-state?

$t = 4$

Solution:

- At $t = 1$, static routes are installed at R1 and R6.
- At $t = 2$, R1 and R6 advertise their static routes to their neighbors.
- At $t = 3$, R2 and R5 receive the advertisements and update their tables.
- At $t = 4$, R3 receives advertisements from R1, R2, and R5.

At this point, R3 has received the shortest paths to all destinations:
- A through R1 with cost 2.
- B through R5 with cost 2.

Since no future advertisement improves any of these costs, R3 has reached steady-state.

A long time after convergence, the link between A and R1 goes down. All routers expire their routes to A except R3.

In each of Q3.5–Q3.7, select the routers that will eventually have a cost-16 entry for A. Each subpart is independent.

Q3.5 (2 points) Assume split horizon and poison reverse are disabled.

[A] R1 [B] R2 [C] R3 [D] R4 [E] R5 [F] R6 Ⓖ None

Solution:
- R3 advertises to R1, R2 and R5, which accept this as they don't have an entry for A.
- R1 then sends its entry to A. R3 accepts the advertisement, even though it would be a worse cost, as R1 is its next hop to reach A.
- R3 updates its cost, and then advertises this information to its neighbors again.
- This causes increasing costs to propagate through the network, causing all routers to have 16 cost to A.

This content is protected and may not be shared, uploaded, or distributed.

Q3.6 (2 points) Assume only split horizon is enabled.

A R1 B R2 C R3 D R4 E R5 F R6 G None

Solution:
- R3 still has a route to A, so it advertises the route to R2 and R5.
- R2 and R5 accept the route because their routes to A have expired.
- Since split horizon is enabled, R2 and R5 do not advertise the route back to R3.
- R2 and R5 advertise the route to their other neighbors, causing R1, R4, and R6 to learn routes through them.
- R1 advertises its route back to R3. Since R1 is R3's current next hop to A, R3 accepts the update even though the cost is higher.
- This creates a count-to-infinity loop, eventually causing all routers to reach cost 16.

Q3.7 (2 points) Assume only poison reverse is enabled.

A R1 B R2 C R3 D R4 E R5 F R6 G None

Solution:
- We explicitly send poison with cost 16 or inf to our next hop 1.
- R3 sends its cost to the other routers.
- R1 sends its inf cost to its neighbors, including R3
- R3 accepts an advertisement from its next hop, updating its cost to 16.
- R3 advertises inf cost to its neighbors, causing it to propagate through the network.

This content is protected and may not be shared, uploaded, or distributed.

Q4 *Link-State* **(11 points)**

Consider running the link-state algorithm from lecture on this topology. Each subpart is independent.

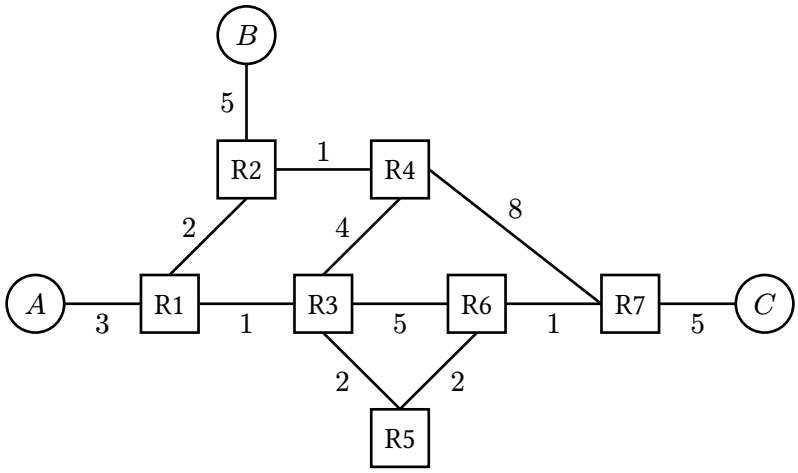

Q4.1 (2 points) Suppose that the link between R5 and R6 goes down. R5 and R6 have recomputed their routes, but they have **not** yet sent any updates.

What is the total cost of a packet's route from A to C? If the packet does not reach C, write "inf."

inf.

Solution: If link R5-R6 goes down, the packet will take the route A → R1 → R3 → R5 → R3 (loop).

Q4.2 (2 points) Suppose that the link between R2 and R4 goes down. R2 and R4 have recomputed their routes, but they have **not** yet sent any updates.

What is the total cost of a packet's route from B to C? If the packet does not reach C, write "inf."

18

Solution: If link R2-R4 goes down, the packet will take the route B → R2 → R1 → R3 → R6 → R7 → C.

For subparts Q4.3 to Q4.4, suppose that our routers failed to track sequence numbers or timestamps in link-state advertisements.

This content is protected and may not be shared, uploaded, or distributed.

Q4.3 (2 points) In 10 words or fewer, explain why this causes problems.

This will cause infinite flooding in a loop or between two routers.

Q4.4 (2 points) As a quick fix, you delete router R3 from the network.

True or False: This removes the need to track sequence numbers or timestamps.

Ⓐ True Ⓑ False

Solution: Deleting R3 prevents infinite flooding in loops but does not solve infinite flooding between two routers.

Q4.5 (3 points) Suppose two routers know the whole topology. However, one router computes shortest paths using Dijkstra's (minimum total link cost), while the other router uses Breadth First Search (minimum number of hops, ignoring link costs).

In general (not necessarily for the provided topology), which of the following statements are true? Select all that apply.

[A] Packets may enter a forwarding loop.

[B] Packets may take non-shortest paths but still reach their destinations.

[C] Every packet is guaranteed to follow a globally shortest path.

[D] One router's shortest-path computations can affect another router's forwarding table.

Ⓔ None of the above

Solution:

Choice A: Different metrics can cause neighboring routers to forward packets to each other, creating a loop.

Choice B: Packets may still arrive, but via a path that is not globally shortest.

Choice C: Different routing metrics do not guarantee globally shortest paths.

Choice D: Each router computes its own forwarding table independently.

This content is protected and may not be shared, uploaded, or distributed.

Q5 *Addressing* **(12 points)**

Q5.1 (2 points) Given an IP address and a subnet mask, how would a host calculate its network prefix?

Ⓐ Bitwise OR Ⓑ Bitwise AND Ⓒ Bitwise XOR Ⓓ None of these

Solution: The host calculates the network prefix by bitwise ANDing the IP address and subnet mask.

Q5.2 (2 points) A network has address `10.50.0.0` with a subnet mask of `255.255.248.0`. What is the CIDR (slash) notation for this network?

Hint: 248 in binary is `11111000`.

Ⓐ `10.50.0.0/20`
Ⓒ `10.50.0.0/22`
Ⓑ `10.50.0.0/21`
Ⓓ `10.50.0.0/24`

Solution: CIDR notation represents the subnet mask as the number of bits of ones at the start of the mask. Since the mask is `255.255.248.0`, we know there are 16 bytes from the two `255`s, and 5 bytes from `248` (which is `11111000`, as-per the hint). $16 + 5 = 21$.

Q5.3 (2 points) A provider allocates the block `172.20.0.0/19` for a new office. What is the total number of usable host IP addresses available within this network block?

Ⓐ 2^{11} Ⓑ 2^{12} Ⓒ 2^{13} Ⓓ 2^{14}

Solution: There are 32 bits in the address, and 19 prefix bits. Therefore there are $32 - 19 = 13$ host bits available.

A router has the forwarding table below, and represents it as a trie to make forwarding decisions:

CIDR Notation	Fixed Bits of Prefix	Output Port
`10.1.0.0/16`	`00001010 00000001`	Port 1
`10.1.64.0/18`	`00001010 00000001 01`	Port 2
`10.1.80.0/20`	`00001010 00000001 0101`	Port 3
`10.1.82.0/23`	`00001010 00000001 0101001`	Port 4
`0.0.0.0/0`	None	Port 5

Q5.4 (2 points) A packet arrives with the destination IP `10.1.82.1`. Which output port will the router forward the packet to?

Ⓐ Port 1 Ⓑ Port 2 Ⓒ Port 3 Ⓓ Port 4 Ⓔ Port 5

Solution: IP 10.1.82.1 converted to binary is `00001010 00000001 01010010 00000001` which matches Port 4 in the table.

This content is protected and may not be shared, uploaded, or distributed.

Q5.5 (2 points) A packet arrives with the destination IP `10.1.95.1`. How many bits of the destination address does the router check before determining that the address maps to Port 3?

Hint: `10.1.95.1` in binary is `00001010 00000001 01011111 00000001`.

Ⓐ 19 bits Ⓑ 20 bits Ⓒ 21 bits Ⓓ 22 bits Ⓔ 23 bits Ⓕ 24 bits

Solution: The first 20 bits of the packet's binary IP address match port 3's prefix and port 4's prefix. The packet's IP address' 21st bit is 1, while port 4's 21st IP address bit is 0. This diverges from port 4's prefix, confirming the packet destined for port 3.

Q5.6 (2 points) A packet arrives with the destination IP `10.1.81.1`. How many bits of the destination address does the router check before determining that the address maps to Port 3?

Hint: `10.1.81.1` in binary is `00001010 00000001 01010001 00000001`.

Ⓐ 19 bits Ⓑ 20 bits Ⓒ 21 bits Ⓓ 22 bits Ⓔ 23 bits Ⓕ 24 bits

Solution: The first 20 bits of the packet's binary IP address correspond to port 3's prefix. The packet's IP address' 23rd bit is 0, while port 4's 23rd address bit is 1. This diverges from port 4's prefix, confirming the packet destined for port 3.

This content is protected and may not be shared, uploaded, or distributed.

This page left intentionally (mostly) blank

The exam continues on the next page.

This content is protected and may not be shared, uploaded, or distributed.

Q6 *BGP* **(15 points)**

Consider the AS graph below, where all ASes follow Gao-Rexford import and export policies.

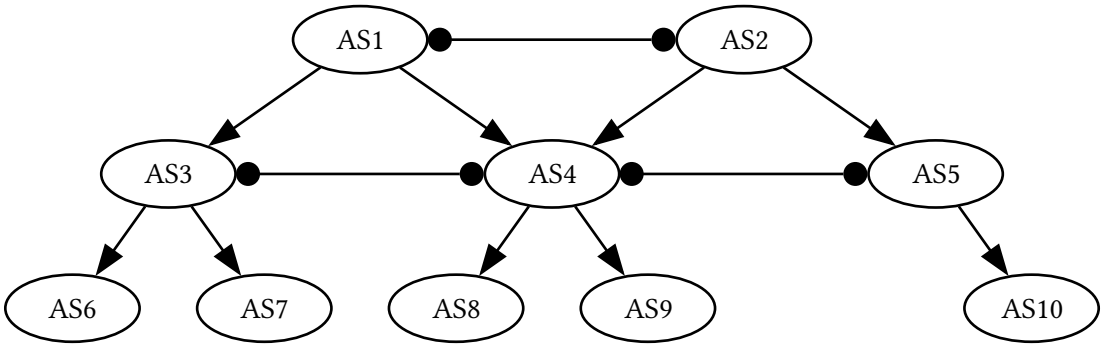

Q6.1 (2 points) AS9 sends a packet to AS7. Which of these paths does the packet take?

(A) AS9 ⟶ AS4 ⟶ AS1 ⟶ AS3 ⟶ AS7

(B) AS9 ⟶ AS4 ⟶ AS3 ⟶ AS7

(C) AS9 ⟶ AS4 ⟶ AS5 ⟶ AS2 ⟶ AS1 ⟶ AS3 ⟶ AS7

(D) AS9 ⟶ AS4 ⟶ AS2 ⟶ AS1 ⟶ AS3 ⟶ AS7

(E) The packet cannot reach AS7.

Solution: This is the only choice that follows the Gao-Rexford "valley-free" rule out of all the possible choices. From AS9 to AS4, traffic flows from a customer to a provider. Then from AS4 to AS3, it then crosses a peer link. Then from AS3 to AS7, it goes from provider to customer. This is allowed because a valid path can go up to a provider, across a peer, and down to a customer. The other answer choices traced out create valleys.

Q6.2 (2 points) AS4 receives a route advertisement from its peer AS5. Which neighbors will AS4 forward that advertisement to?

(A) AS1 and AS2

(B) AS1, AS2, AS8, and AS9

(C) AS1, AS2, AS3, AS8, and AS9

(D) AS8 and AS9

(E) No neighbors

Solution: Under Gao-Rexford export rules, a route learned from a peer can only be advertised to AS4's customers, which in this case are AS8 and AS9.

This content is protected and may not be shared, uploaded, or distributed.

Q6.3 (3 points) For this subpart only, suppose the peer link between AS1 and AS2 fails.

Which of these ASes are no longer able to exchange traffic? Select all that apply.

- [A] AS1 and AS2
- [B] AS2 and AS4
- [C] AS1 and AS10
- [D] AS2 and AS3
- [E] AS7 and AS1
- (F) None of the above

Solution: Once the AS1-AS2 peer link fails, AS1 and AS2 lose their direct peer route.

AS1 also cannot reach AS10 through AS4 and AS5, because AS4 would have to export a peer-learned route from AS5 up to AS1, which is not allowed.

Similarly, AS2 cannot reach AS3 through AS4, because AS4 would have to export a peer-learned route from AS3 up to AS2, which is not allowed.

Q6.4 (2 points) A router in AS2 hears about the two advertisements below. Which route does AS2 prefer?

Advertisement I: A router in AS1 sends: "Destination AS7 with ASPATH [AS1, AS3, AS7]."

Advertisement II: A router in AS4 sends: "Destination AS7 with ASPATH [AS4, AS3, AS7]."

- (A) Advertisement I
- (B) Advertisement II
- (C) Both are equally preferred.
- (D) Not enough information.

Solution: Advertisement 1 comes from AS1, a peer of AS2 while Advertisement 2 comes from AS4, a customer of AS2. AS2 will prefer the advertisement that allows it to make money so Advertisement 2 is preferred.

Q6.5 (2 points) For this subpart only, suppose the peer link between AS4 and AS5 fails.

A router in AS4 hears about the two advertisements below. Which route does AS4 prefer?

Advertisement I:
A router in AS2 sends: "Destination AS10 with ASPATH [AS2, AS5, AS10], and MED 10."

Advertisement II:
A different router in AS2 sends: "Destination AS10 with ASPATH [AS2, AS5, AS10], and MED 1."

- (A) Advertisement I
- (B) Advertisement II
- (C) Both are equally preferred.
- (D) Not enough information.

Solution: Both routes come from AS2 and have the same AS path, so they tie on local preference and AS path length.

Since the advertisements come from different AS2 routers, AS4 may choose based on internal egress cost, which is not shown. So there is not enough information.

This content is protected and may not be shared, uploaded, or distributed.

Q6.6 (2 points) Which statement best explains why the route selected by BGP is not necessarily the shortest AS path in the graph?

(A) BGP ignores AS path length whenever Gao-Rexford rules are used.

(B) Some shorter paths may never be advertised because export policy removes them before route selection happens.

(C) MED values are always compared before AS path length, so a longer path with lower MED may be imported.

(D) AS path length is only used to prevent loops, not to import routes.

Solution: BGP does not choose from every path that appears possible in the graph.

Export policy determines which routes are actually advertised, so a shorter path may never be available if some AS is not allowed to export it. BGP only selects the best route among the advertisements it actually receives.

Q6.7 (2 points) Suppose AS4 learns a route from its peer AS3 and then accidentally advertises that route to AS2, even though this violates Gao-Rexford export policy. What happens?

(A) AS2 may accept the route, because AS2 does not care that AS4 violated its export policy.

(B) AS2 rejects the route, because AS4 would not make money from carrying AS2's traffic.

(C) AS2 rejects the route, because the advertisement comes from AS2's customer.

(D) AS2 accepts the route only if the destination is AS4's customer.

Solution: Gao-Rexford says AS4 should not export a peer-learned route to AS2.

However, if AS4 sends the advertisement anyway, BGP does not automatically know that the route was learned from a peer. AS2 may accept the route according to its own import policy.

This content is protected and may not be shared, uploaded, or distributed.

Q7 *Transport* **(15 points)**

Host A is sending data to Host B using TCP, as described in lecture.

Assume the following, unless otherwise stated:
- Congestion control is not in use.
- All TCP sequence numbers, ACK numbers, and window sizes are measured in bytes.
- Each data segment holds 4 bytes.

Q7.1 (2 points) A's initial sequence number is #700, and B's initial sequence number is #900.

What is the sequence number of the first data byte that A sends to B?

Ⓐ #700	Ⓒ #702	Ⓔ #900	Ⓖ #902
Ⓑ #701	Ⓓ #703	Ⓕ #901	Ⓗ #903

Solution: The sequence number of the first data byte is the one after A's initial sequence number of #700.

This content is protected and may not be shared, uploaded, or distributed.

Q7.2 (5 points) Before receiving any segments, B's advertised window is 16 bytes.

B receives segments S1, S3, S2 (in that order), and immediately sends an ACK for each segment. Fill in the table below.

Hint: B buffers out-of-order packets that cannot be delivered to the application.

Segment received	**ACK number sent by B**	**Segment(s) B can deliver to the application**	**B's advertised window after processing segment(s)**
#701–#704 (S1)	Ⓐ #701 Ⓑ #704 **Ⓒ #705**	**A** S1 B S2 C None	Ⓐ 8 Ⓑ 12 **Ⓒ 16**
#709–#712 (S3)	**Ⓐ #705** Ⓑ #709 Ⓒ #712	A S2 B S3 **C None**	Ⓐ 8 **Ⓑ 12** Ⓒ 16
#705–#708 (S2)	Ⓐ #708 Ⓑ #712 **Ⓒ #713**	**A** S2 **B** S3 C None	Ⓐ 8 Ⓑ 12 **Ⓒ 16**

Solution: For segment 1, B sends an ack for the next smallest unacked sequence number, which is 705. We can send Segment 1 to the application because it has been acked, and the advertised window remains the same because B's buffer is empty, as we send Segment 1.

For segment 3, B sends an ack for the next smallest unacked sequence number, which is still 705. We can not send segment 2 or segment 3 because they have not been acked. We store 4 bytes from segment 3 in our buffer, so A can only advertise 12 more bytes to B.

For segment 2, B sends an sneds an ack for the next smallest unacked sequence number which is 713. We can now send both segment 3 and segment 2 which have been acked. Because all of the segments have been sent, B's buffer is empty, and can store 16 new bytes.

This content is protected and may not be shared, uploaded, or distributed.

For Q7.3 to Q7.6, congestion control is in use, as described in lecture.

Assumptions for Q7.3 to Q7.6:
- CWND and SSTHRESH are measured in segments, where each data segment holds 4 bytes.
- A's usable send window is $\min(\text{RWND}, \text{CWND})$, and A always has more data to send.

Q7.3 (2 points) Suppose A receives ACK #713. This ACK advertises $\text{RWND} = 2$ segments.

Right before receiving this ACK, $\text{CWND} = 5$ segments, and A is in Slow Start.

After processing ACK #713, which bytes is A allowed to send?

Ⓐ #713 through #716
Ⓒ #717 through #720
Ⓑ #713 through #720
Ⓓ #721 through #724

Solution: After processing ACK #713, CWND increases by 1 to 6 segments as we are in Slow Start. However, the usable send window is min(RWND, CWND) or 2 segments. Therefore A sends 2 segments or 8 bytes i.e. #713 through #720.

Q7.4 (2 points) Suppose instead that ACK #713 advertises a very large RWND.

Right before receiving this ACK, CWND = 4 segments, SSTHRESH = 4 segments, and A is in Slow Start.

After processing ACK #713, what mode is A in?

Ⓐ Slow Start
Ⓑ Fast Recovery
Ⓒ Congestion Avoidance

Solution: After processing ACK #713, CWND increases by 1 to 5 segments as we are in Slow Start. The usable send window is min(RWND, CWND) or 5 segments since RWND is very large. Since the usable segment window is greater than SSTHRESH, we enter Congestion Avoidance.

This content is protected and may not be shared, uploaded, or distributed.

Q7.5 (2 points) Some time later, A is in Congestion Avoidance with these parameters:

- CWND $= 4$ segments, SSTHRESH $= 8$ segments.
- Bytes #713-#728 have been sent but not yet ACKed.
- Bytes #713-#716 are lost.
- Bytes #717-#728 arrive at B.

After A receives 3 duplicate ACKs for ACK #713, what are the new values of SSTHRESH and CWND?

Ⓐ SSTHRESH $= 1$ segment, CWND $= 5$ segments

Ⓑ SSTHRESH $= 2$ segments, CWND $= 2$ segments

Ⓒ SSTHRESH = 2 segments, CWND $= 5$ segments

Ⓓ SSTHRESH $= 8$ segments, CWND $= 4$ segments

Solution: Upon receiving 3 duplicate ACKs, TCP triggers Fast Recovery and sets SSTHRESH to CWND / 2, or 2 segments. It also sets CWND to SSTHRESH + 3, or 5 segments.

Q7.6 (2 points) Continuing from Q7.5, A retransmits bytes #713-#716. Then, A receives ACK #729. What happens next?

Ⓐ A stays in Fast Recovery and CWND stays 5 segments.

Ⓑ A enters Congestion Avoidance and CWND becomes 5 segments.

Ⓒ A enters Congestion Avoidance and CWND becomes 2 segments.

Ⓓ A enters Slow Start and CWND becomes 1 segment.

Solution: When A receives a new ACK #729, TCP exits Fast Recovery and returns to Congestion Avoidance. At this point, CWND is set to SSTHRESH or 2 segments.

This content is protected and may not be shared, uploaded, or distributed.

Comment Box

Congrats for making it to the end of the exam! Leave any thoughts, comments, feedback, or doodles here.

Nothing in the comment box will affect your grade.

Ambiguities

If you feel like there was an ambiguity on the exam, you can put it in the box below.

For ambiguities, you must qualify your answer and provide an answer for both interpretations. For example, "if the question is asking about A, then my answer is X, but if the question is asking about B, then my answer is Y." You will only receive credit if it is a genuine ambiguity and both of your answers are correct. We will only look at this box if you request a regrade.

This content is protected and may not be shared, uploaded, or distributed.