CS 168 Fall 2024 — Intro to the Internet: Architecture and Protocols — Midterm

Solutions last updated: October 18, 2024

Name: ______________________________

Student ID: ______________________________

This exam is 110 minutes long.

Question:	1	2	3	4	5	6	7	Total
Points:	21	12	14	12	18	11	12	100

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

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

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

- ■ You can select
- ■ multiple squares (completely filled)

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: ______________________________

Q1 *Quick Questions* **(21 points)**

Each question is 1 point unless otherwise specified.

Q1.1 A packet "on the wire" carries headers for all layers.

● True ● False

Solution: The intended answer was True. Packets on the wire have all layers wrapped around it.

However, after the exam, some students said they interpreted "all layers" to refer to Layers 1, 2, 3, 4, and 7. In this case, the answer would be False, because some packets never have higher-layer headers to begin with. For example, traceroute packets are UDP-over-IP, so they never have a Layer 7 packet at any point.

Because of this alternate interpretation, we decided to grant points to both True and False.

Q1.2 Logical ports are needed to demultiplex a packet at the receiving end host.

● True ○ False

Solution: True. Logical ports tell us which application to pass the packet up to.

Q1.3 The Protocol field in the IPv4 header is needed to demultiplex a packet at the receiving end host.

● True ○ False

Solution: True. The Protocol field tells us which Layer 4 protocol to pass the packet up to.

Q1.4 A TCP sender can send a packet to a UDP receiver, as long as the sender uses the correct Protocol field in the packet's IPv4 header.

● True ● False

Solution: The intended answer was False. TCP and UDP are fundamentally different protocols, deeming them not compatible to speak to each other. If you build a TCP packet and put the UDP protocol number, the receiver cannot parse the TCP packet using the UDP protocol.

However, after the exam, some students said they interpreted this question as: "The TCP sender can send the packet, regardless of whether the UDP receiver can actually receive it." In this case, nothing is really stopping the sender from sending the packet (even if it will never be received by anybody).

Because of this alternate interpretation, we decided to grant points to both True and False.

Q1.5 A reservation-based network architecture violates fate sharing.

● True ○ False

Solution: True. If one router fails, other routers are affected as well.

Q1.6 Reducing the MTU (Maximum Transmission Unit) on a link makes it easier to build high-speed router linecards.

○ True ● False

Solution:

False. The same amount of data would need to be broken into more packets, which means that there are more headers that need to be processed for the same data. For example, consider an extreme example where the MTU is set to 1 byte. Then the router has to process far more headers to send the same amount of data.

In other words: Any linecard must be able to forward packets at the full capacity of the link (and the link capacity stays unchanged whether my MTU is small or large). Hence a smaller MTU means more packets per second to be processed (meaning more parsing, more lookups, etc).

Q1.7 As of 2024, IPv6 is universally adopted in the Internet.

○ True ● False

Solution: False. As seen in lecture, less than half of the Internet has adopted IPv6.

Q1.8 AS 205 is connected to a single other AS, via a single link. AS 205 needs to configure MEDs.

○ True ● False

Solution: False. The MED is used to advertise to your neighbors which entrance you prefer for a particular destination. However, AS 205 only has a single outgoing link to other ASes, so there's no need to advertise a preference between two entrances.

Q1.9 The count-to-infinity problem exists in BGP.

○ True ● False

Solution: False. BGP uses path-vector, so advertisements will not travel in a loop. If an AS sees an advertisement with itself along the path, it will not further propagate that advertisement.

Q1.10 In BGP, policy oscillations will always lead to packet loops.

- ○ True
- ● False

> **Solution:** False. Consider the example in lecture: policy oscillations cause the AS to change which path it is selecting but not necessarily to loops in the path.

Q1.11 How does the OS know if an **IPv6** packet is a UDP packet or a TCP packet?

- ○ Protocol field
- ○ TTL field
- ● Next Header field
- ○ Port number

> **Solution:** In IPv6, we use the Next Header field to decide the Layer 4 protocol.
>
> (IPv4 uses the Protocol field, while IPv6 uses the Next Header field.)

Q1.12 IPv6 eliminates the checksum that was in the IPv4 header. Which design principle best describes this change?

- ● End-to-end principle
- ○ Layering
- ○ Fate sharing
- ○ Statistical multiplexing

> **Solution:** By removing checksums at the IP layer, we're leaving reliability to the end hosts. This is what the end-to-end principle says.

Q1.13 Which best describes the Ethernet CSMA/CD protocol?

- ● Talk only if you hear no one else talking, but then stop talking if you hear someone else currently talking.
- ○ Pass a token around and only talk if you are holding the token.
- ○ Raise your hand, and wait until a coordinator gives you permission to talk.
- ○ Every person is scheduled a time to talk.

> **Solution:** CSMA is described in the first option.

Q1.14 (2 points) Host A is transmitting multiple packets to Host B.

We notice that each packet's delay is different. Select all possible causes.

- ■ Different packets take different routes.
- ■ Packets encounter different queuing delays.
- □ BGP leads to asymmetric paths between hosts.
- ■ Some packets include IPv4 options, which requires additional processing at routers.
- □ None of the above

Solution:

A: True. A packet taking a different route could experience longer delay.

B: True. Queuing delays can affect the overall packet delay.

C: False. All packets are being transmitted from Host A to Host B, and not from Host B to Host A.

D: True. IPv4 options can increase the time it takes for the router to process the packet.

SR: given our concern with part D how about we switch it to: "Some ASes use hot-potato routing and others don't" (F).

Q1.15 (2 points) Select all true statements about time-division multiplexing.

- ■ There could be times when a channel is idle, even a sender has data that they'd like to send on the channel.
- ■ The channel requires the sender and receiver's clocks to be closely synchronized.
- □ Data in the channel could experience variable delays due to queuing in the network.
- □ In times of high utilization, a sender could be completely denied access to the channel.
- □ None of the above

Solution:

A: True. If a sender has data to send and it's not their turn, the channel could end up idle.

B: True. Strict Synchronization is required so that the sender and receiver both use the correct time slots.

C: False. When it's the sender's turn, it has full access to the channel, so packets don't get queued.

D: False. A sender can still use their dedicated time slot to send data.

Q1.16 (2 points) Select all disadvantages of CIDR compared to classful addressing.

- ☐ Forwarding tables are larger.
- ■ Forwarding table lookups are more complicated.
- ■ Address assignment is more complicated.
- ☐ Routers are more likely to drop packets.
- ☐ None of the above

Solution:

A: False. CIDR allows you to aggregate hosts in a way that classful addressing does not.

B: True. CIDR requires longest prefix matching, because multiple prefixes in the forwarding table might match. In classful addressing, at most one prefix in the forwarding table will patch.

C: True. CIDR enables multi-level hierarchical address assignment. In classful addressing, there is only one level in the address hierarchy (networks, then hosts).

D: False. CIDR does not directly lead to packets being dropped more often, as this is more related to network congestion or other routing issues.

Q1.17 (2 points) Select all header fields that an IP router must modify before forwarding an IPv4 packet to its next hop.

- ■ TTL
- ☐ Destination IP
- ☐ Protocol
- ☐ Source IP
- ■ Checksum
- ☐ None of the above

Solution: The router must decrement the TTL and update the header checksum.

Q2 *Ethernet* (12 points)

Consider running the Spanning Tree Protocol (STP) for the following network topology:

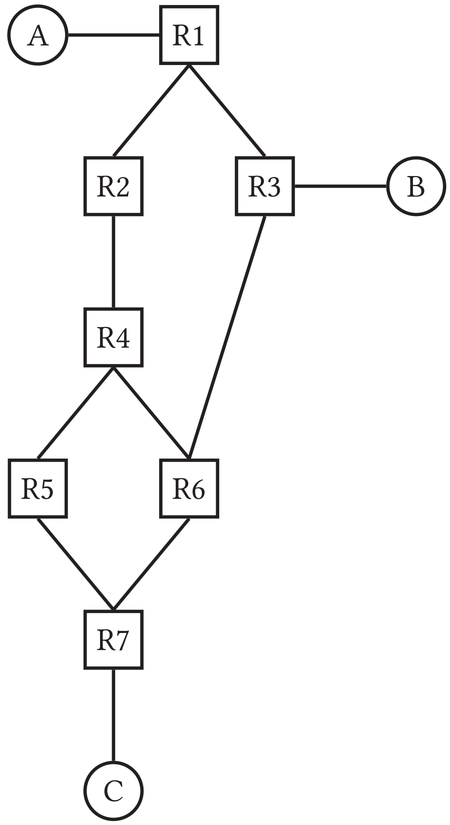

Assume the IDs are ordered according to the router labels. For example, R4 has a lower ID than R5.

Solution: Note (Oct 18, 2024): We didn't assign a lot of partial credit on this question for two reasons:

1. We're offering a redemption STP question on the final exam.

2. A lot of incorrect answers we saw looked like random guesses, so it was hard to distinguish any notion of a "conceptual error" that would warrant partial credit.

Q2.1 (2 points) Select all of the links that are **enabled** after running the STP protocol.

- ■ R1-to-R2
- ■ R2-to-R4
- ■ R4-to-R5
- ☐ R5-to-R7
- ■ R1-to-R3
- ■ R3-to-R6
- ☐ R4-to-R6
- ■ R6-to-R7

Solution: The switch with the smallest ID is R1, so R1 will end up being the root.

Therefore, we have to find the spanning tree that includes all shortest paths to R1.

Q2.2 (2 points) Select all routers that disable at least one of their links.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ■ R6 ■ R7

Solution: R7 disables its link to R5.

R6 disables its link to R4.

Suppose STP has converged (i.e. only the links you chose above are enabled). Switches R1-R7 are all learning switches, and their forwarding tables start out empty.

In each of the next three subparts, select all switches that will receive the given packet.

The packets are sent one after the other. In other words, forwarding table entries created in one subpart carry over to later subparts.

Q2.3 (2 points) B sends a packet to C.

■ R1 ■ R2 ■ R3 ■ R4 ■ R5 ■ R6 ■ R7

Solution: All the forwarding tables are empty at this point, so every router floods the packet.

Q2.4 (2 points) C sends a packet to B.

☐ R1 ☐ R2 ■ R3 ☐ R4 ☐ R5 ■ R6 ■ R7

Solution: After the previous subpart, all routers now have a next-hop to B, so the packet is forwarded from C, to R7, to R6, to R3, to B.

Q2.5 (2 points) C sends a packet to A.

- [x] R1
- [x] R2
- [x] R3
- [x] R4
- [x] R5
- [x] R6
- [x] R7

> **Solution:** None of the switches know how to reach A (because A has never sent out a packet), so the packet must be flooded to reach A.

Q2.6 (1 point) What is the earliest point at which **at least one** switch has a forwarding table entry for destination C?

- ○ After the packet in Q2.3 is sent.
- ● After the packet in Q2.4 is sent.
- ○ After the packet in Q2.5 is sent.
- ○ Not all switches have an entry after the three packets are sent.

> **Solution:** See the solution to the next subpart.

Q2.7 (1 point) What is the earliest point at which **all** switches have a forwarding table entry for destination C?

- ○ After the packet in Q2.3 is sent.
- ○ After the packet in Q2.4 is sent.
- ● After the packet in Q2.5 is sent.
- ○ Not all switches have an entry after the three packets are sent.

> **Solution:** After the first packet (Q2.2), every router has an entry for B, but nobody has an entry for C.
>
> After the second packet (Q2.3), only routers R3, R6, and R7 see an outgoing packet from C, so only those 3 routers have an entry for C.
>
> After the third packet (Q2.4), every router sees an outgoing packet from C, so every router is able to install a table entry for C.

Q3 *Routing* **(14 points)**

In this question, consider finding least-cost routes in the following network topology:

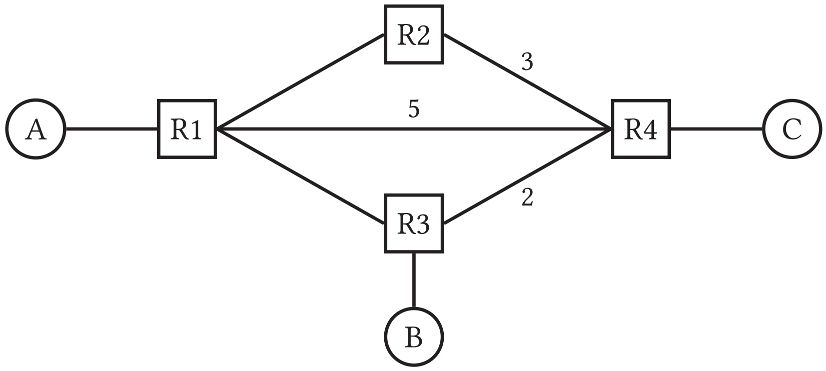

Q3.1 (2 points) Select all directed edges that are included in the directed delivery tree for destination A.

☐ A → R1	■ R1 → A	☐ R1 → R2	■ R2 → R1
☐ R1 → R4	☐ R4 → R1	☐ R1 → R3	■ R3 → R1
☐ R2 → R4	☐ R4 → R2	☐ R3 → R4	■ R4 → R3

Solution:

Tree 1 has one arrow out of each router, and all arrows point toward A, and the paths toward A are the shortest paths.

Tree 2 is wrong because R4's shortest path is via R3, not via R1.

Tree 3 is wrong because it has two arrows coming out of R4.

Tree 4 is wrong due to the arrows pointing in the opposite direction. In a delivery tree, arrows point toward the destination.

Q3.2 (2 points) Suppose there are two flows on this network:

Flow 1: A sends a large amount of traffic to B.

Flow 2: A sends a large amount of traffic to C.

The network administrator wants to change the cost along the R2-to-R4 link, such that the two flows use different paths.

Select all costs we can assign to the R2-to-R4 link such that no link is used by both flows (except the A-to-R1 link).

■ 1　　☐ 4　　☐ 5　　☐ 6　　☐ 7　　☐ None

Solution: The current shortest path from A to C is the "lower path": A, R1, R3, R4, C. This has cost 5.

One other alternative path from A to C is the "upper path": A, R1, R2, R4, C. This currently has cost 6.

If we decrease the R2-to-R4 cost from 3 to 1, the cost of the upper path is now 4. This causes A-to-C traffic to take the upper path.

All other answer choices would cause the upper path to cost more, which would cause the A-to-C traffic to continue using the lower path (sharing links with the A-to-B traffic).

Suppose we use the distance-vector protocol from lecture to compute least-cost routes in this topology, and all routes have converged.

Regardless of your earlier answers, suppose the R2-to-R4 link cost changes from 3 to 1.

Q3.3 (2 points) R2 learns about the link cost changing.

Select all entries in R2's forwarding table that are changed as a direct result of this.

- ☐ Next-hop to A
- ☐ Next-hop to B
- ☐ Next-hop to C
- ☐ Cost to A
- ☐ Cost to B
- ■ Cost to C
- ☐ None of the above

Solution:

At convergence, R2's forwarding table looks like this:

Destination, Next-hop, cost

A, R1, 2

B, R1, 3

C, R4, 4

The only path that uses the R2-to-R4 link is the path to C. Since the link to the next-hop along this path changes, we have to update the cost to C from 4 to 2.

R2 sends advertisements to all its neighbors after updating its forwarding table. In the next 2 subparts, fill in the advertisements that R2 sends.

Assume that **poison reverse** is enabled.

Q3.4 (2 points) To R1: "I can reach ___ with cost ___."

Solution:

Solution:

Solution: I can reach **C** with cost **2**.

R2 changed its cost to C, so it needs to tell its neighbors about its new cost to C. (In the previous part, we showed that this new cost is 2.)

Q3.5 (2 points) To R4: "I can reach ___ with cost ___."

Solution:	**Solution:**

Solution: I can reach **C** with cost **INFINITY**.

Since R4 is the next hop for R2 to destination C, poison reverse tells us to send back to R4 that we can reach C with cost infinity in order to prevent any loops from occurring.

Q3.6 (1 point) Which router(s) will update their forwarding table(s) based on the advertisements sent by R2?

- ● R1 only
- ○ Both R1 and R4
- ○ R4 only
- ○ Neither R1 nor R4

Solution: R1 changes its forwarding table.

The new route has cost 2 (advertised by R2) + 1 (link cost to R2) = 3.

The original best route to C has cost 4 (R1 to R3 to R4 to C).

Since the new cost is better, the advertisement is accepted.

R4 does not accept the advertisement because its current best cost to C is 1 (direct), and infinity via R2 is not better than 1.

The remaining subparts use the same topology, but they are independent of the earlier subparts.

Now, consider using the link-state protocol from lecture to compute least-cost routes in this topology.

Q3.7 (1 point) After the link-state protocol converges, the resulting routing tables are _____ compared to the routing tables computed by the distance-vector protocol.

- ○ smaller
- ● the same size
- ○ larger

Solution: In both routing protocols, each router needs one entry per destination, so the size of the routing tables stays the same.

Q3.8 (2 points) Suppose that after the link-state protocol converges, the R2-to-R4 link cost changes from 3 to 1 again.

Select all routers that would need to re-compute paths through the network (even if the actual paths don't change).

- ■ R1
- ■ R3
- ☐ None
- ■ R2
- ■ R4

Solution: In the link-state protocol, the updated link cost would get flooded to all routers, and each router has to independently check whether or not the updated link cost changes any of their shortest paths.

Q4 *Addressing* **(12 points)**

Consider the following address hierarchy:

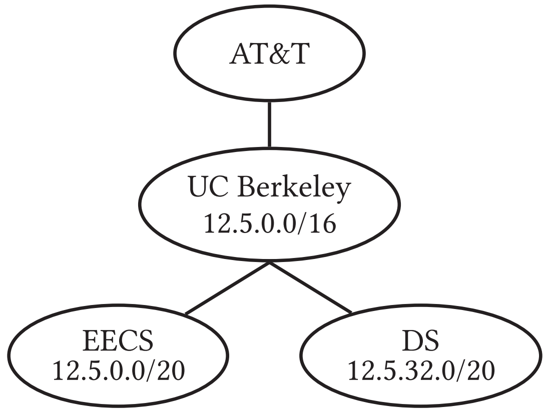

The forwarding table for UC Berkeley looks like this:

Prefix	**Next-Hop**
12.5.0.0/20	EECS
12.5.32.0/20	DS
0.0.0.0/0	AT&T

Q4.1 (1 point) How many IP addresses are in EECS's prefix?

○ 2^8 ● 2^{12} ○ 2^{16} ○ 2^{20} ○ 2^{24} ○ 2^{32}

Solution: The prefix is 20 bits long, so 20 bits are fixed, and the remaining 12 bits can change. This gives us 2^{12} addresses in this range.

Q4.2 (1 point) Assuming each department gets its own separate /20 prefix, how many total departments can UC Berkeley allocate prefixes to?

● 2^4 ○ 2^8 ○ 2^{12} ○ 2^{16} ○ 2^{20} ○ 2^{32}

Solution: UC Berkeley has a prefix with 16 bits fixed.

It fixes 4 more bits and allocates that to a department.

There are 2^4 ways to fix those 4 bits.

Q4.3 (2 points) Given this forwarding table, does UC Berkeley need to run longest prefix matching (LPM) in order to forward incoming packets to the correct department?

- ○ Yes, because EECS and DS are part of the same /16 prefix.
- ○ Yes, because EECS and DS have overlapping prefixes, and UC Berkeley needs to pick the most specific one.
- ● No, because EECS and DS have non-overlapping prefixes.
- ○ No, because LPM is used for sending advertisements, not forwarding packets.

Solution: In this specific case, LPM is not needed because EECS and DS have non-overlapping prefixes. We could just check bits 16-20 and forward to the appropriate department.

Note that in general, LPM is needed because we can't guarantee that all the entries in the forwarding table have non-overlapping prefixes.

Suppose that someone outside of UC Berkeley sends a packet to 12.5.255.1. AT&T forwards this packet to UC Berkeley.

Q4.4 (2 points) What happens to this packet?

- ● It gets stuck in a routing loop between UC Berkeley and AT&T.
- ○ It gets stuck in a routing loop between UC Berkeley and EECS.
- ○ It gets dropped by UC Berkeley.
- ○ UC Berkeley forwards the packet to EECS, which drops it.
- ○ UC Berkeley forwards the packet to AT&T, which drops it.

Solution: This packet matches the 0.0.0.0/0 route, so UC Berkeley forwards it back to AT&T.

AT&T then forwards the packet back to UC Berkeley because the destination is in the 12.5.0.0/16 range.

This continues infinitely, in a routing loop.

Q4.5 (2 points) Select all modifications that would fix this problem, without introducing any more dead-ends or loops.

- ☐ Let the routing algorithm run until it converges.
- ☐ Change the default route entry in UC Berkeley's forwarding table.
- ☐ Delete a single entry from UC Berkeley's forwarding table.
- ☐ Delete a single entry from AT&T's forwarding table.
- ■ None of the above

Solution:

A: False. The routing algorithm has already converged, so letting it run on will not fix this issue.

B: False. There's no way to encode the range "all destinations except DS and EECS" within a single prefix.

C: False. Deleting the default route will create a dead-end because all outgoing UC Berkeley packets would get dropped. Likewise, deleting the EECS or DS routes would create dead-ends for packets intended for those destinations.

D: False. AT&T has a single forwarding table entry with range 12.5.0.0/16 and next-hop UC Berkeley. Deleting this entry would stop all UC Berkeley packets from being delivered (i.e. a dead-end).

UC Berkeley wants to add one entry to its forwarding table to prevent this problem.

Q4.6 (2 points) What is the prefix of this new entry?

Solution: The intended answer was 12.5.0.0/16.

This includes all destinations inside UC Berkeley (so these packets should not be forwarded back to AT&T).

For partial credit, we accepted any prefix that avoids loops for the exact IP given, but does not avoid loops in general. Specifically, we accepted any prefix that satisfies these conditions:

1. Is a subset of 12.5.0.0/16 (UC Berkeley).

2. Does not include 12.5.0.0/20 (EECS) and 12.5.32.0/20 (DS).

3. Does include 12.5.255.1 (the packet described in the problem).

The most common answers that fits this rubric item are 12.5.240.0/20 and 12.5.128.0/17 and 12.5.255.1/32 and 12.5.255.0/20.

Q4.7 (1 point) What should UC Berkeley do with packets matching this prefix?

- ○ Forward to EECS.
- ● Drop the packet.
- ○ Forward to DS.
- ○ Forward to AT&T.

Solution: UC Berkeley should drop this packet, because it should not be sent to any other next-hop in this topology.

Q4.8 (1 point) With this new forwarding table entry, does UC Berkeley need to run longest prefix matching to process incoming packets?

- ● Yes
- ○ No

Solution: Yes, because 12.5.0.0/16 and 12.5.0.0/20 and 12.5.32.0/20 are now overlapping prefixes.

Q5 *BGP* **(18 points)**

Consider the following AS graph:

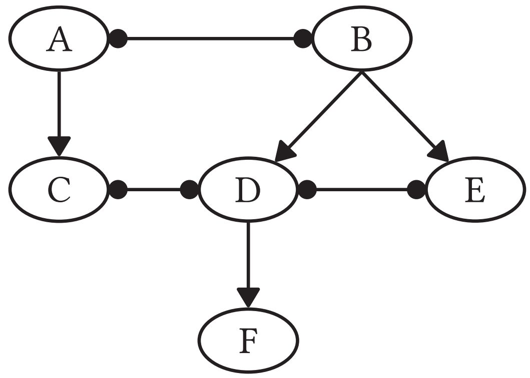

Q5.1 (1 point) Select all Tier 1 ASes in this AS graph.

■ A ■ B □ C □ D □ E □ F

Solution: A and B are Tier 1 ASes since they have no providers.

For the next 2 subparts, select whether each path is possible under the Gao-Rexford rules.

Q5.2 (1 point) C to D to E:

○ Valid ● Invalid

Solution: D is paid by nobody on this path (C and E are both its peers).

Q5.3 (1 point) F to D to B to E:

- ● Valid
- ● Invalid

Solution: The intended answer was Valid, because D is paid by F, and B is paid by both D and E.

However, after the exam, students pointed out that this path is only Valid if you consider BGP export rules (what paths an AS will agree to participate in).

If you additionally consider BGP import rules (prefer a path from customer over peer over provider), then D would actually prefer the route F to D to E, because D prefers a next-hop of E (peer) over a next-hop of B (provider).

Because we didn't specify whether you need to consider import rules (in addition to output rules), we accepted both Valid and Invalid.

On future exams, we will make sure to specify whether we're asking about validity for export rules only, or both export/import rules.

Suppose AS A and AS B each publish a *dump* of all BGP routes that they learn.

Note: Each BGP route includes a destination AS and the AS path to that destination.

Q5.4 (2 points) A researcher extracts all the edges from the two dumps to reconstruct a view of the AS graph. Select all edges that appear in this reconstructed AS graph.

- ■ A-B
- ■ B-D
- □ C-D
- ■ D-F
- ■ A-C
- ■ B-E
- □ D-E
- □ None

Solution:

A's dump looks like this:

Destination	Path
B	A-B
C	A-C
D	A-B-D
E	A-B-E
F	A-B-D-F

And B's dump looks like this:

Destination	Path
A	B-A
C	B-A-C
D	B-D
E	B-E
F	B-D-F

The union of all edges in these tables is the answer.

Note that when finding routes from A and B to all the other ASes, edges C to D and D to E are never traversed due to Gao Rexford Rules.

Q5.5 (3 points) Select the minimum set of ASes that need to publish a dump, so that the researcher is able to reconstruct the entire AS graph (i.e. with all seven edges).

Note: This subpart is independent from the previous one, i.e. A or B will only publish dumps if you select A or B below.

Note: There may be multiple correct answers.

☐ A ☐ B ☐ C ■ D ☐ E ☐ F

Solution: D's dump on its own should be sufficient to learn all edges.

Destination	Path
A	D-B-A
B	D-B
C	D-C
C	D-B-A-C
E	D-E
E	D-B-E
F	D-F

(An older version of this PDF had a different answer, which we now think is incorrect. Specifically, the dump should include any BGP path exported, not just the one that is actually used.)

In the next 3 subparts, suppose you are AS C, and you want to send malicious BGP advertisements, such that packets with destination F are sent to you.

Q5.6 (2 points) Construct an advertisement that AS C could send that would cause some packets with destination F to be forwarded to AS C.

"I can reach _____ via path _____."

Solution:

Solution:

Solution: I can reach **F** via path **C - F**.

This causes others to believe that they can send packets with destination F to AS C.

Q5.7 (1 point) AS C exports this advertisement to AS D.

Will AS D accept this advertisement?

- ○ Yes, because of the Gao-Rexford rules.
- ○ Yes, because the `ASPATH` in the advertisement is shorter.
- ○ No, because the `ASPATH` in the advertisement is longer.
- ● No, because of the Gao-Rexford rules.

Solution: In the advertised path, D does not get paid by C.

In D's current path to F, D gets paid by F.

By the Gao-Rexford rules, D prefers the path via a customer (F), instead of a path via a peer (C).

Note that the `ASPATH` length is irrelevant here because we only use that when there is a tie between two choices from the Gao-Rexford rules.

Q5.8 (1 point) AS C exports this advertisement to AS A.

Will AS A accept this advertisement?

- ● Yes, because of the Gao-Rexford rules.
- ○ Yes, because the `ASPATH` in the advertisement is shorter.
- ○ No, because the `ASPATH` in the advertisement is longer.
- ○ No, because of the Gao-Rexford rules.

Solution: In the advertised path, A gets paid by C.

A's current path to F is A, B, D, F. In this path, A is not paid by anybody.

By the Gao-Rexford rules, A prefers the path via a customer (C), instead of a path via a peer (B).

Note that the `ASPATH` length is irrelevant here because we only use that when there is a tie between two choices from the Gao-Rexford rules.

The rest of this question is independent from the earlier subparts.

Suppose there are two links between AS D and AS F:

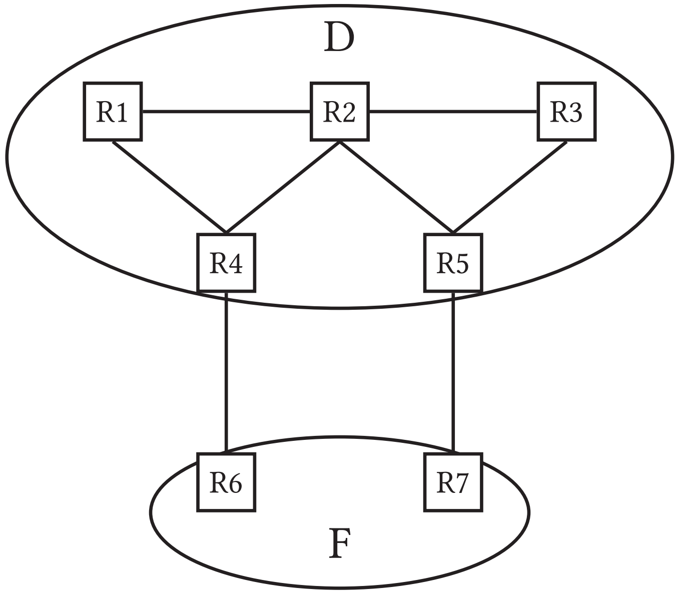

Suppose you are R6, and you want to send malicious BGP advertisements, such that packets entering AS F are sent to you (i.e. R6 instead of R7).

Q5.9 (1 point) Which BGP attribute can you set in R6's advertisements to try and convince incoming packets to be sent to R6 instead of R7?

- ○ Destination
- ○ `LOCAL PREF`
- ○ `ASPATH`
- ● `MED`

Solution: You can use the MED to convince incoming packets to be sent to R6 instead of R7.

Q5.10 (1 point) How does R6 advertise the attribute from the previous subpart?

- ● Using eBGP.
- ○ Using iBGP.
- ○ Using IGP.

Solution: R6 needs to advertise the MED to other ASes, which means the attribute is included in the eBGP message.

Suppose R6 has sent the malicious advertisement from the previous subparts.

Q5.11 (1 point) R1 receives a packet with destination in F.

Does this packet get sent via R6 or R7?

- ● R6, because of hot-potato routing.
- ○ R7, because of hot-potato routing.
- ○ R6, because of `MED`.
- ○ R7, because of `MED`.
- ○ R6, because of `ASPATH`.
- ○ R7, because of `ASPATH`.
- ○ R6, because of `LOCAL PREF`.
- ○ R7, because of `LOCAL PREF`.

Q5.12 (1 point) R2 receives a packet with destination in F.

Does this packet get sent via R6 or R7?

- ○ R6, because of hot-potato routing.
- ○ R7, because of hot-potato routing.
- ● R6, because of `MED`.
- ○ R7, because of `MED`.
- ○ R6, because of `ASPATH`.
- ○ R7, because of `ASPATH`.
- ○ R6, because of `LOCAL PREF`.
- ○ R7, because of `LOCAL PREF`.

Q5.13 (1 point) R3 receives a packet with destination in F.

Does this packet get sent via R6 or R7?

- ○ R6, because of hot-potato routing.
- ● R7, because of hot-potato routing.
- ○ R6, because of `MED`.
- ○ R7, because of `MED`.
- ○ R6, because of `ASPATH`.
- ○ R7, because of `ASPATH`.
- ○ R6, because of `LOCAL PREF`.
- ○ R7, because of `LOCAL PREF`.

Q5.14 (1 point) Which routers in AS D would have forwarded via R7, but are now forwarding via R6 as a result of your advertisement?

- ○ All routers in AS D.
- ● Only routers with equal IGP cost between R4 and R5.
- ○ Only routers with a next-hop of R4.
- ○ Only routers with a next-hop of R5.
- ● None of the routers in AS D.

Solution: The intended answer was "Only routers with equal IGP cost between R4 and R5."

When D sends packets, it uses hot-potato routing before considering MED. When a packet first enters D, if the packet is closer to R5, then that packet will be sent out of R5 (and to R7), and won't pass through R6.

The only packets that will pass through R6 are the ones that enter D equidistant from R4 and R5 (so MED is used to pick R4-R6), or the ones that are enter D closer to R4 (so hot-potato routing sends the packet through R4-R6).

However, after the exam, some students pointed out that given the actual routers in AS D, the only router equidistant from the egresses (that would consider MED) is R2. If the two paths somehow had the same MED, then R2 would tiebreak arbitrarily and choose R4-R6 (lower router ID), so R2 ends up not changing its forwarding decision either. This was not our intended solution, because you can't assume the MED values being advertised before the malicious advertisement were equal. However, we gave students the benefit of the doubt and accepted "None of the routers in AS D" as an answer.

This alternate interpretation seemed to be very rare, as 343/500 students answered with our intended answer, and only 19/500 students answered with the alternate interpretation (which probably also includes some random guesses).

Q6 *Routers* **(11 points)**

Rob proposes that we merge L2 and L3 into a single layer, with addressing and routing done on MAC addresses (replacing IP addressing entirely). Rob also proposes modifying routing protocols to use MAC addresses as the destination.

Let's compare Rob's proposal to the current IP-based Internet architecture. Each true/false is 1 point.

Q6.1 Because MAC addresses are burned into hardware, once a router installs a forwarding table entry for a specific destination, that entry will never change.

○ True ● False

> **Solution:** False. Although it is true that MAC addresses are burned into hardware, the computer with that MAC could move elsewhere, which would require the forwarding table entry to change.

Q6.2 It is possible for multiple destination MAC addresses to be associated with the same physical port.

● True ○ False

> **Solution:** True. If multiple destinations have the same next-hop, they would be sent out of the same physical port.

Q6.3 Because the forwarding table is used in the data plane (not the control plane), Rob's proposal results in the same load on a router's control plane CPU.

○ True ● False

> **Solution:** False. The control plane is needed to build the actual forwarding table (by running a routing protocol). Because MAC addresses can't be aggregated, the routing protocol will have to send far more advertisements, so the control plane CPU will have increased load.

Q6.4 Longest prefix matching is the most efficient method for MAC destination-based forwarding.

○ True ● False

> **Solution:** False. Because MAC addresses can't be aggregated, it would be more efficient to just look up an exact match.

Q6.5 Recall that MAC addresses are 48 bits long. If the router does not aggregate any forwarding table entries, the forwarding table must contain exactly 2^{48} entries.

○ True ● False

> **Solution:** False. Not all MAC addresses have been allocated. Also, some devices with an allocated MAC address might not be connected to the Internet.

Sylvia modifies Rob's proposal: All routers and hosts inside an AS must buy their Ethernet equipment from a single vendor, and no vendor can sell their equipment to more than one AS.

Let's compare Sylvia's proposal to **the current IP-based Internet architecture** (not to Rob's proposal).

Q6.6 Sylvia's proposal results in less load on a router's control plane CPU.

○ True ● False

> **Solution:** False. IP-based routing lets you aggregate multiple ASes, but in Sylvia's proposal, no aggregation between ASes can occur.

Q6.7 Sylvia's proposal reduces churn in forwarding tables, because hosts can leave one AS and join another AS without causing any forwarding tables to change.

○ True ● False

> **Solution:** False. If a host moves to a different AS, their MAC address stays the same, so the forwarding tables must change to map the same host to a new next-hop.

In the next subpart, assume that:

- There are W ASes in total.
- Each AS has exactly X hosts.
- Each AS has exactly Y border routers.
- Each AS is connected to exactly Z other ASes.

Your answer in the next subpart should be a big-O bound in terms of W, X, Y, and Z. You can drop constants, and you may not need all four variables.

Q6.8 (4 points) After convergence, how many forwarding table entries does each router have in...

...Rob's proposal?

> **Solution:**

...Sylvia's proposal? (Assume no multi-homing.)

> **Solution:**

> **Solution:** Rob's proposal: $O(WX)$
>
> We need one entry for each destination, and no aggregation is possible. There are W ASes, and each has X hosts, for a total of WX entries.

> **Solution:** Sylvia's proposal: $O(X + W)$
>
> We need one entry for each local destination (X entries). Since we can aggregate entries for an AS, we only need one entry for each other AS ($W - 1$ entries).

Q7 *Traceroute* (12 points)

The standard MTU on the Internet is 1500 bytes. Suppose one buggy router sets its MTU to 1000 bytes, and drops any packets larger than 1000 bytes.

Suppose there is a single path between you and a destination: you, R1, R2 ..., destination. The buggy router is along this path, and you want to discover the IP address of the buggy router.

Q7.1 (1 point) Why would running traceroute from the project fail to reveal the buggy router?

- ○ The buggy router still processes TTLs correctly.
- ● Packets sent by traceroute are usually very small.
- ○ Other routers besides the buggy router will drop traceroute probes.
- ○ The buggy router will drop traceroute probes.

> **Solution:** Traceroute packets are pretty small, meaning that they would be unaffected by this buggy router.

In order to discover the buggy router, you need to perform two separate runs of traceroute.

In the next 3 subparts, design an algorithm for discovering the buggy router. You can assume there are no network errors (e.g. no drops besides the buggy router, no duplicates, no invalid packets, etc.).

Write each modification in 10 words or fewer. If no modification is needed, write "Standard Traceroute."

Q7.2 (3 points) What modification(s) will you make in your **first** run of traceroute?

> **Solution:** Standard Traceroute
>
> We use standard traceroute first to discover all the routers along the path to the destination.

Q7.3 (3 points) What modification(s) will you make in your **second** run of traceroute?

> **Solution:** Set payload to 1001 bytes.
>
> This can help us find the router that drops packets larger than 1000 bytes.
>
> This run of traceroute will return all routers up to, but not including, the buggy one. This is because the buggy router will drop your size-1001 probe, which means that the buggy router and all subsequent routers will not receive probes and not send replies.

Q7.4 (3 points) From your two runs of traceroute, you receive two lists of routers.

The first run returns the following list of routers: R1, R2, ..., R27.

The second run returns the following list of routers: R1, R2, ..., R13.

From these lists, which router is the buggy router? Answer with the router number, e.g. "R81."

Solution: The intended answer was **R14**.

The first is a list of all routers along the path to the destination.

The first is a list of all routers up to, but not including, the buggy router.

Therefore, the buggy router is the one immediately after the last router in the second list.

However, after the exam, some students said that they interpreted the buggy router as replying with a TTL Exceeded message (but not forwarding packets). Under that interpretation, R13 is the buggy router, since it replied to the TTL Exceeded message but did not forward the large packet to R14 (so R14 did not reply). Therefore, we accepted both **R13** and **R14** as correct answers.

Q7.5 (2 points) If router _____ drops all probes on the _____ run of traceroute, then your modified algorithm would return the wrong answer.

○ R1, first ○ R13, first ○ R14, first ○ R27, first

● R1, second ● R13, second ○ R14, second ○ R27, second

Solution: The intended answer was R13, second.

This would cause the second list to be R1, R2, ..., R12, which would incorrectly cause your algorithm to return R13 (not R14) as the buggy router.

However, after the exam, some students said that they interpreted "a router dropping probes" as the router dropping all probes, even ones not meant for that router. Under this interpretation, if R1 drops all probes on the second run, then the second run of traceroute would simply return an empty list, and you would mistakenly think that R1 is the buggy router.

This was not our intended interpretation, since we were using the wording from Project 1 (where a router dropping probes refers to probes meant for that router). However, we gave students the benefit of the doubt and accepted the alternate answer.

Nothing on this page will affect your grade.

Post-Exam Activity

CS 168 is looking for a new course mascot! Suggest any designs or doodles below:

Comment Box

Congratulations for making it to the end of the exam! Feel free to leave any thoughts, comments, feedback, or doodles here: