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

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

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

- ◯ True
- ◯ False

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

- ◯ True
- ◯ False

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

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

- ◯ True
- ◯ False

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

- ◯ True
- ◯ False

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

- ◯ True
- ◯ False

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

- ◯ True
- ◯ False

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

- ◯ True
- ◯ False

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

- ◯ True
- ◯ False

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

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

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.

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

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

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

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

Q2 *Ethernet* **(12 points)**

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

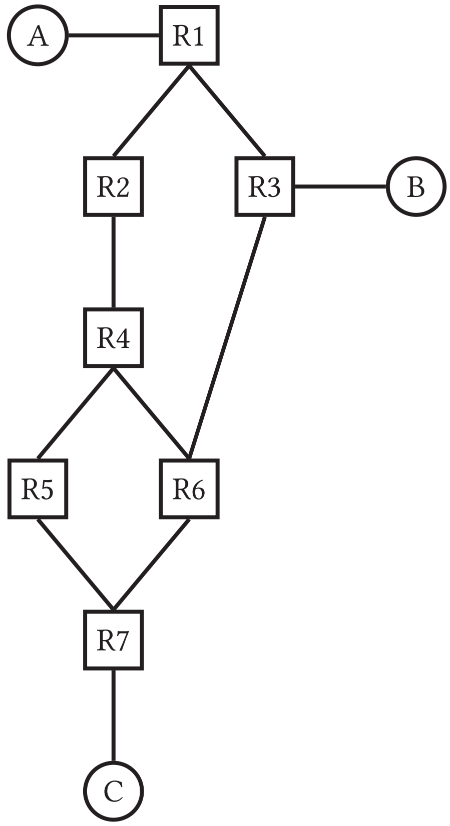

Assume the IDs are ordered according to the router labels. For example, R4 has a lower ID than R5.

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

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

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6 ☐ R7

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

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

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6 ☐ R7

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

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6 ☐ R7

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.

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.

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

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

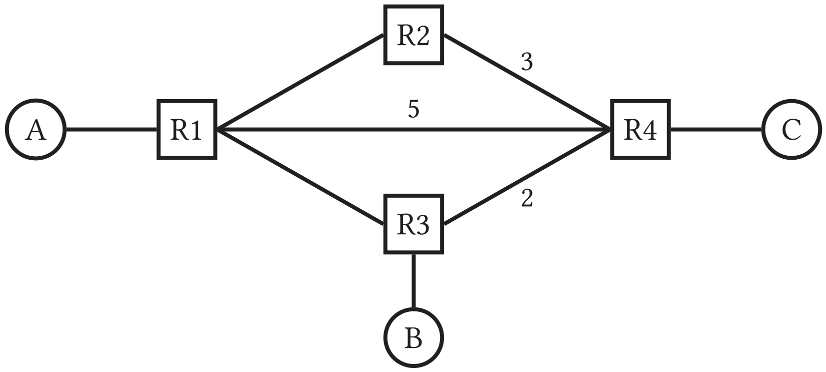

Q3.1 (2 points) Select all directed edges that are included in the directed delivery tree for destination A.

- ☐ A → R1
- ☐ R1 → A
- ☐ R1 → R2
- ☐ R2 → R1
- ☐ R1 → R4
- ☐ R4 → R1
- ☐ R1 → R3
- ☐ R3 → R1
- ☐ R2 → R4
- ☐ R4 → R2
- ☐ R3 → R4
- ☐ R4 → R3

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

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

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 ___."

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

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

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

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

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

Consider the following address hierarchy:

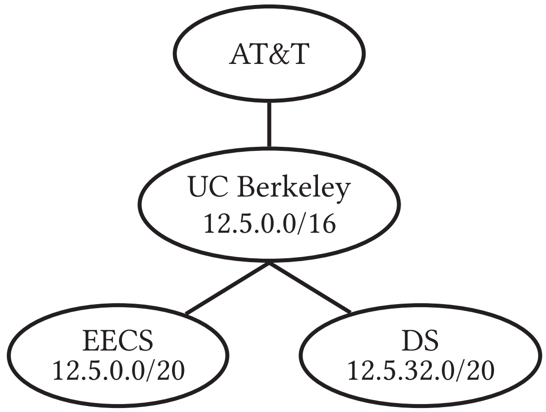

The forwarding table for UC Berkeley looks like this:

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}

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}

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.

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.

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

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?

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.

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

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

Consider the following AS graph:

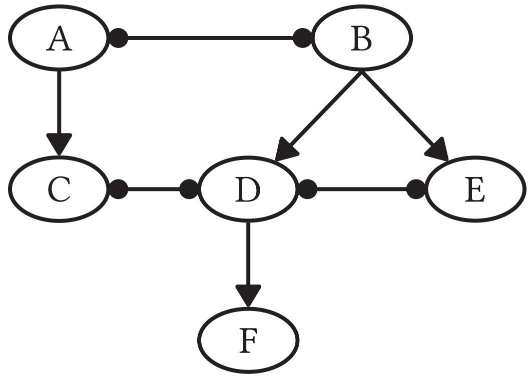

Q5.1 (1 point) Select all Tier 1 ASes in this AS graph.

☐ A ☐ B ☐ C ☐ D ☐ E ☐ F

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

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

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

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

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 _____."

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.

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.

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

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

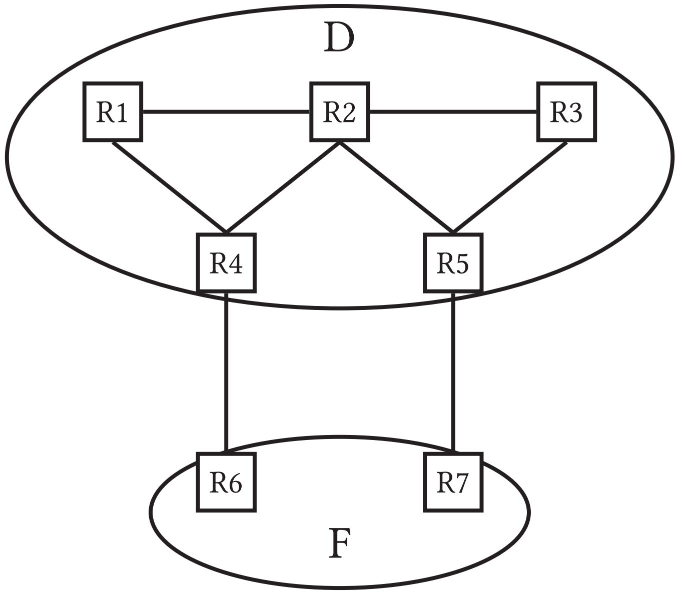

Suppose you are R6, and you want to send malicious BGP advertisements, such that packets entering AS F are sent to you (i.e. R6 instead of R7).

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`

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

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

Suppose R6 has sent your 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.

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

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

○ True ○ False

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

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

○ True ○ False

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

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

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

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?

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

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.

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?

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

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."

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

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: