CS 168 — Introduction to the Internet
Spring 2026 — Final Exam

Print Your Name: ______________________________

Print Your Student ID: ______________________________

You have 170 minutes. There are 8 questions of varying credit. (100 points total)

Question:	1	2	3	4	5	6	7	8	Total
Points:	25	15	10	13	11	7	11	8	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: ______________________________

Q1 *TCP Congestion Control* **(25 points)**

For the entire question, all TCP values are measured in packets, not bytes, unless otherwise specified.

In Q1.1 to Q1.7, consider two modifications to the TCP implementation from lecture:
- In slow start, CWND is multiplied by 3 on each RTT (instead of 2).
- If loss is detected from duplicate ACKs, CWND is divided by 3 (instead of 2).

The rest of TCP is unchanged, except where these two modifications require corresponding updates.

Q1.1 (2 points) To implement the modified slow start with event-driven updates, CWND should increase by ___ on each new ACK received.

◯ 0.5 ◯ 1 ◯ 1.5 ◯ 2 ◯ 2.5 ◯ 3

Q1.2 (1 point) Consider sending an endless stream of data on a link with some constant, unknown bandwidth. No other users are using the link.

We run two separate, independent experiments on the same link:
- Run unmodified slow start (doubling CWND), and record the first non-infinite SSTHRESH value.
- Run modified slow start (tripling CWND), and record the first non-infinite SSTHRESH value.

Which algorithm has the higher initial finite value of SSTHRESH?

◯ Unmodified (doubling) ◯ Modified (tripling) ◯ Not enough information

Q1.3 (1 point) When we enter fast recovery (after receiving 3 duplicate ACKs), what should CWND be?

◯ $\frac{\text{CWND}}{2} + 2$ ◯ $\frac{\text{CWND}}{2} + 3$ ◯ $\frac{\text{CWND}}{3} + 2$ ◯ $\frac{\text{CWND}}{3} + 3$

Q1.4 (1 point) During fast recovery, CWND should increase by ___ on each duplicate ACK received.

◯ 0.5 ◯ 1 ◯ 1.5 ◯ 2 ◯ 2.5 ◯ 3

(Question 1 continued...)

The two modifications, reprinted for your convenience:
- In slow start, CWND is multiplied by 3 on each RTT (instead of 2).
- If loss is detected from duplicate ACKs, CWND is divided by 3 (instead of 2).

Next, we will derive a throughput equation for the modified TCP, using the same steady-state assumptions as the original equation:
- A single connection is sending an endless stream of data on a link.
- The link has fixed maximum bandwidth W_{max}, and packet loss occurs when CWND reaches this value.
- CWND is adjusted only by AIMD updates (i.e. ignore slow start and fast recovery).
- RTT is fixed.

Q1.5 (1 point) Over one steady-state AIMD cycle, what is the average value of CWND?

○ $\frac{1}{6}W_{\text{max}}$ ○ $\frac{1}{3}W_{\text{max}}$ ○ $\frac{1}{2}W_{\text{max}}$ ○ $\frac{2}{3}W_{\text{max}}$ ○ $\frac{5}{6}W_{\text{max}}$ ○ W_{max}

Q1.6 (2 points) In steady state, how many packets are sent between consecutive loss events?

○ $\frac{2}{9}W_{\text{max}}$ ○ $\frac{1}{3}W_{\text{max}}$ ○ $\frac{4}{9}W_{\text{max}}$ ○ $\frac{5}{9}W_{\text{max}}$ ○ $\frac{2}{3}W_{\text{max}}$ ○ $\frac{7}{9}W_{\text{max}}$

○ $\frac{2}{9}W_{\text{max}}^2$ ○ $\frac{1}{3}W_{\text{max}}^2$ ○ $\frac{4}{9}W_{\text{max}}^2$ ○ $\frac{5}{9}W_{\text{max}}^2$ ○ $\frac{2}{3}W_{\text{max}}^2$ ○ $\frac{7}{9}W_{\text{max}}^2$

Q1.7 (2 points) In this subpart only, ignore your answers to the previous two subparts and instead assume:
- The average value of CWND is $\frac{1}{3}W_{\text{max}}$.
- The number of packets sent between two losses is $\frac{3}{4}W_{\text{max}}^2$.

Based on these values, what is the value of k in the throughput equation?

$$\text{throughput} = k \cdot \frac{\text{MSS}}{\text{RTT}\sqrt{p}}$$

○ $\frac{\sqrt{3}}{9}$ ○ $\frac{\sqrt{6}}{9}$ ○ $\frac{2\sqrt{3}}{9}$ ○ $\frac{2\sqrt{6}}{9}$ ○ $\frac{2\sqrt{3}}{3}$ ○ $\frac{\sqrt{3}}{3}$

In Q1.8 to Q1.9, consider the two modifications from earlier, with one additional modification:
We enter fast recovery when we receive 2 duplicate ACKs (instead of 3 duplicate ACKs).

Q1.8 (1 point) When we enter fast recovery, what should CWND be set to?

○ $\frac{\text{CWND}}{2} + 2$ ○ $\frac{\text{CWND}}{3} + 2$ ○ $\frac{\text{CWND}}{2} + 3$ ○ $\frac{\text{CWND}}{3} + 3$

Q1.9 (2 points) How does this new modification affect TCP (compared to the version with only two modifications)? Select all that apply.

☐ It has no effect on sending rate in networks where packet reordering occurs.

☐ It can reduce the average sending rate, because TCP may cut its window more often.

☐ During fast recovery, CWND should now increase by 2 on each duplicate ACK received.

☐ It affects the throughput equation (which models TCP using only AIMD).

☐ It makes TCP more likely to reduce CWND, even when no packet was dropped.

○ None of the above

Q1.10 to Q1.11 are independent of earlier subparts (i.e. ignore the modifications above). For these subparts, assume:

- TCP transfers occur on a path with fixed RTT and fixed available bandwidth.
- Packet loss only occurs when CWND exceeds the available bandwidth.

Q1.10 (2 points) In this subpart only, consider unmodified TCP from lecture.

Two long, ongoing TCP connections, X and Y, are identical, except for one difference:

- Connection X currently has a **lower** SSTHRESH value.
- Connection Y currently has a **higher** SSTHRESH value.

How does Connection Y compare to Connection X? Select all that apply.

☐ Y may stay in slow start for more RTTs before switching to congestion avoidance.

☐ Y may reach a high sending rate sooner.

☐ Y may overshoot the available bandwidth by more before switching to congestion avoidance.

○ None of the above

Q1.11 (2 points) In this subpart only, consider unmodified TCP from lecture.

A TCP connection is sending data on a path with a much larger RTT than a typical network (e.g. an Earth-to-Moon link).

How does this connection compare to a connection with a smaller RTT? Select all that apply.

☐ It takes longer in real time for slow start to grow CWND to a given value.

☐ At a given CWND value, the sending rate in packets per second is lower.

☐ If a timeout occurs, recovering to the previous sending rate may take longer in real time.

○ None of the above

Q1.12 (2 points) In this subpart only, we modify TCP by making the initial CWND larger (with no other modifications).

For a **short** transfer, how does this TCP version compare to unmodified TCP? Select all that apply.

☐ The sender can transmit more data before the first ACK returns.

☐ The transfer may finish in fewer RTTs.

☐ The sender may trigger packet loss sooner if the path cannot support the larger initial CWND.

☐ Application completion time is unchanged, because additive increase is unchanged.

☐ The change affects the throughput equation (which models TCP using only AIMD).

○ None of the above

In Q1.13 to Q1.15 (which are independent of all earlier subparts), consider this modified version of TCP:

- In slow start, CWND is multiplied by 3 each RTT.
- In congestion avoidance, CWND increases by 1 packet per RTT.
- When loss is detected from duplicate ACKs, CWND is divided by 3.
- TCP enters fast recovery after 2 duplicate ACKs.

An ongoing TCP connection currently has these parameters:

- Current mode = slow start.
- Current CWND = 3 packets.
- Current SSTHRESH = 27 packets.
- RTT = 0.5 seconds.
- Packet loss occurs after CWND first reaches 31 packets.
- After TCP enters fast recovery, the retransmitted packet is acknowledged exactly 1 RTT later.

Assume packet transmission time and processing time are negligible.

Q1.13 (2 points) Starting from the given parameters, how many seconds pass before TCP first enters congestion avoidance?

seconds

Q1.14 (2 points) Starting from when TCP first enters congestion avoidance, how many additional seconds pass before TCP first enters fast recovery?

seconds

Q1.15 (2 points) Starting from when TCP first enters fast recovery, how many additional seconds pass before TCP exits fast recovery?

seconds

Q2 *DNS: Aliasing* **(15 points)**

In DNS, the `CNAME` record type holds a name-value pair, where the name is a domain, and the value is another domain. A `CNAME` record indicates the name domain is an *alias* of the value domain.

"Alias" means that the name domain is another way to refer to the value domain. When a user queries for the name domain, they will find the `CNAME` record, which tells them to query for the value domain instead.

For example, `cs168.berkeley.edu   CNAME   sp26.cs168.io` is a record that indicates that `cs168.berkeley.edu` is another way to refer to `sp26.cs168.io`.

If a user queries for the IP address of `cs168.berkeley.edu`, they will find this `CNAME` record, which tells the user to look up the IP address of `sp26.cs168.io` instead.

Consider the DNS hierarchy below. Each box contains a zone, and the domain and IP of the authoritative name server for that zone. There are no other zones and name servers besides the ones shown.

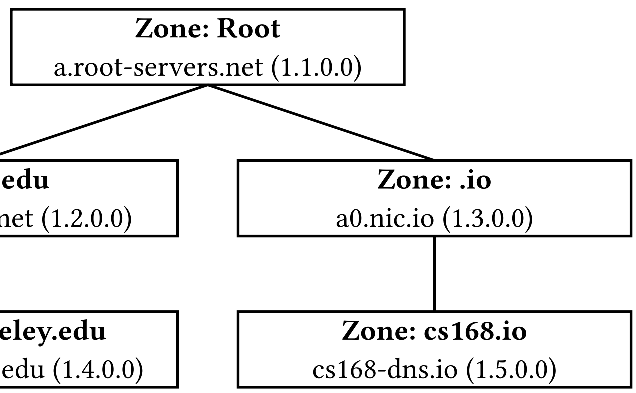

In the berkeley.edu zone, we add the record: `cs168.berkeley.edu   CNAME   sp26.cs168.io`

In Q2.1 to Q2.2, a recursive resolver starts with an empty cache, and first makes a query to fill up the cache. Assume no DNS or network errors occur.

Then, some time later, without any cached records expiring, the resolver wants to know the IP address of `cs168.berkeley.edu`. How many DNS queries does the resolver send in order to resolve this query?

Q2.1 (1 point) The resolver starts with an empty cache.

How many queries are needed to learn the IP address of `cs168.berkeley.edu`?

○ 0 ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 ○ 6 ○ 7

Q2.2 (1 point) The resolver first queries for the IP address of `sp26.cs168.io`, and caches all records from this query.

Now, how many queries are needed to learn the IP address of `cs168.berkeley.edu`?

○ 0 ○ 1 ○ 2 ○ 3 ○ 4 ○ 5 ○ 6 ○ 7

EvanBot suggests storing the record `sp26.cs168.io    A    113.26.9.12` in the adns1.berkeley.edu name server (and not in any other name server).

Q2.3 (1 point) If no additional changes are made, EvanBot's suggestion is not allowed, because the ____ (i) domain is not in the ____ (ii) that the name server is ____ (iii) for.

- ◯ (i): `sp26.cs168.io` (ii) cache (iii) the recursive resolver
- ◯ (i): `sp26.cs168.io` (ii) zone (iii) authoritative
- ◯ (i): `113.26.9.12` (ii) zone (iii) the recursive resolver

Q2.4 (2 points) Which record should be added in order for EvanBot's change to be allowed?

- ◯ `berkeley.edu  NS  adns1.berkeley.edu`
- ◯ `berkeley.edu  A  1.4.0.0`
- ◯ `berkeley.edu  NS  cs168-dns.io`
- ◯ `berkeley.edu  A  1.5.0.0`
- ◯ `sp26.cs168.io  NS  adns1.berkeley.edu`
- ◯ `sp26.cs168.io  A  1.4.0.0`
- ◯ `sp26.cs168.io  NS  cs168-dns.io`
- ◯ `sp26.cs168.io  A  1.5.0.0`

Q2.5 (1 point) Where should the record in the previous subpart be added?

- ◯ In the adns1.berkeley.edu name server.
- ◯ In the cs168-dns.io name server.

In Q2.6 and Q2.7, suppose we implement EvanBot's suggestion and add the record in Q2.4.

Hint: If a name server is authoritative for both domains in a `CNAME` record, then a `CNAME` and `A` record can both be returned in a single DNS response.

Q2.6 (1 point) If EvanBot's suggestion is implemented, and a resolver starts with an empty cache, how many queries are needed to learn the IP address of `cs168.berkeley.edu`?

◯ 0 ◯ 1 ◯ 2 ◯ 3 ◯ 4 ◯ 5 ◯ 6 ◯ 7

Q2.7 (1 point) If EvanBot's suggestion is implemented, and a resolver starts with an empty cache, how many queries are needed to learn the IP address of `sp26.cs168.io`?

◯ 0 ◯ 1 ◯ 2 ◯ 3 ◯ 4 ◯ 5 ◯ 6 ◯ 7

Q2.8 and Q2.9 are independent of the earlier subparts.

Q2.8 (3 points) Provide two records that could cause a DNS query to encounter an infinite loop.
In each empty box, write a domain (e.g. a.com), or an IP address (e.g. 1.1.1.1), or a record type (e.g. `A`).

	Name	**Type**	**Value**
Record 1:		`CNAME`	
Record 2:			

Q2.9 (2 points) The DNS specification requires that if a `CNAME` record exists for a domain, then no other records should exist with that same domain.

For example, if `a.com CNAME b.com` exists, then `a.com A 1.2.3.4` cannot exist.

Which option best describes the problem that would occur if somebody broke this rule?

- ◯ One domain could have two IP addresses, which is not allowed.
- ◯ Two domains could have the same IP address, which is not allowed.
- ◯ Two domains could be aliases, but map to different sets of IP addresses, which is not allowed.
- ◯ Two domains could be aliases, but map to the same IP addresses, which is not allowed.

CodaBot runs the **`www.berkeley.edu`** website, and wants to load-balance requests across 9 CDN servers with domains **`cdn-1.berkeley.edu`**, **`cdn-2.berkeley.edu`**, ..., **`cdn-9.berkeley.edu`**.

For the remaining subparts, assume the user directly makes DNS queries, i.e. ignore stub and recursive resolvers.

Q2.10 (1 point) When a user makes a DNS query for **`www.berkeley.edu`**, what DNS record should be returned to achieve CodaBot's load-balancing goal?

Assume **`i`** is a random number between 1 and 9, generated once per DNS request.

- ◯ `cdn-i.berkeley.edu CNAME www.berkeley.edu`
- ◯ `cdn-i.berkeley.edu NS www.berkeley.edu`
- ◯ `www.berkeley.edu CNAME cdn-i.berkeley.edu`
- ◯ `www.berkeley.edu NS cdn-i.berkeley.edu`

Q2.11 (1 point) Is it possible to include an additional **`A`** record in the DNS response, so that the user only needs to make a single DNS query to access the website?

- ◯ Yes, the response can include an **`A`** record for **`www.berkeley.edu`**.
- ◯ Yes, the response can include an **`A`** record for **`cdn-i.berkeley.edu`**.
- ◯ No, because DNS queries cannot include additional records.
- ◯ No, because the name server sending the response is not authoritative for the **`A`** record.
- ◯ No, because an additional query would still be needed, even if the **`A`** record is included.

Q3 *STP* (10 points)

(0.5 points per blank) Consider running the Spanning Tree Protocol (STP) on the topology below.

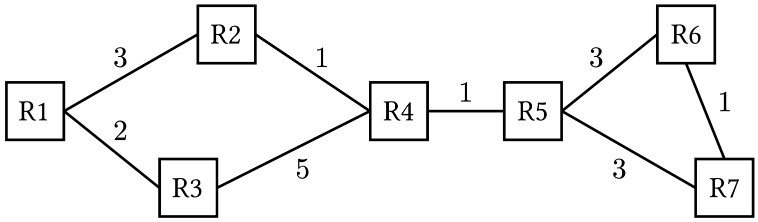

Suppose R1 is elected as the root. After running STP, in the left cycle, _____ (Q3.1) disables its link to _____ (Q3.2), and in the right cycle, _____ (Q3.3) disables its link to _____ (Q3.4).

Q3.1: ○ R1 ○ R2 ○ R3 ○ R4

Q3.3: ○ R5 ○ R6 ○ R7

Q3.2: ○ R1 ○ R2 ○ R3 ○ R4

Q3.4: ○ R5 ○ R6 ○ R7

Starting from the original topology again, suppose R7 is elected as the root instead. After running STP, in the left cycle, _____ (Q3.5) disables its link to _____ (Q3.6), and in the right cycle, _____ (Q3.7) disables its link to _____ (Q3.8).

Q3.5: ○ R1 ○ R2 ○ R3 ○ R4

Q3.7: ○ R5 ○ R6 ○ R7

Q3.6: ○ R1 ○ R2 ○ R3 ○ R4

Q3.8: ○ R5 ○ R6 ○ R7

Regardless of your previous answers, suppose the network has converged on this spanning tree:

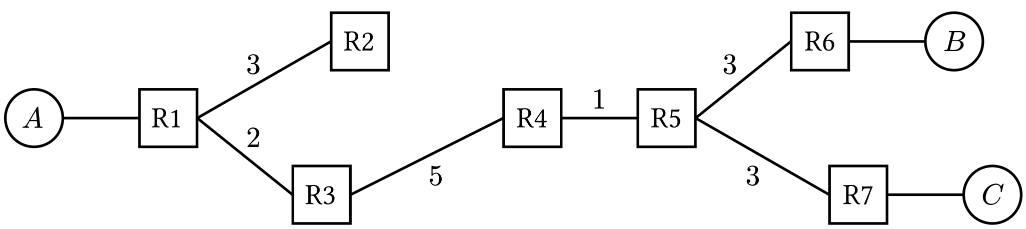

Switches R1 to R7 are all learning switches. All forwarding tables start out empty, except R5, whose table has a hard-coded entry with the least-cost routes to B.

In Q3.9 to Q3.11, 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.

Q3.9 (2 points) A sends a packet to B.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6 ☐ R7

Q3.10 (2 points) C sends a packet to A.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6 ☐ R7

Q3.11 (2 points) B sends a packet to C.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6 ☐ R7

Q4 *End-to-End* **(13 points)**

In this question, consider the network topology below, with A, B, R1 in a subnet connected with a single shared medium, and R3, C, D, in another subnet connected with a single shared medium. Everyone uses the same DNS resolver.

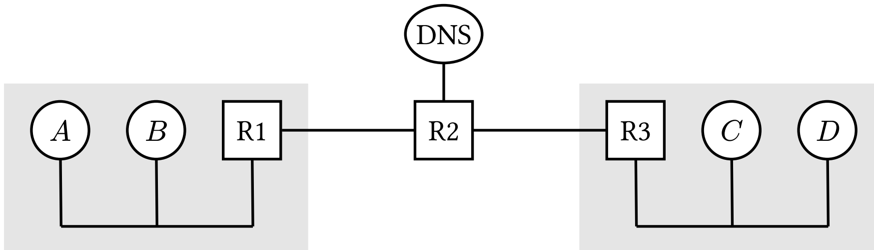

A joins the network and types **`http://www.d.com`** in its browser to load a webpage from D.

Q4.1 (2 points) Host A runs DHCP and learns that its own IP is 10.1.1.3, the router IP is 10.1.1.2, and the subnet mask is /24.

If Host B later joins the network and runs DHCP (while Host A is still on the same DHCP lease), what could be the IP address assigned to Host B? Select all that apply.

☐ 10.1.1.0 ☐ 10.1.4.0 ☐ 10.1.1.3 ☐ 10.4.1.0 ○ None

In Q4.2 to Q4.4, consider a packet sent by A during the HTTP connection.

Q4.2 (1 point) Which protocol was used to learn the value of the Source MAC field?

○ DHCP ○ ARP ○ DNS ○ TCP ○ HTTP ○ None

Q4.3 (1 point) Which protocol was used to learn the value of the Destination MAC field?

○ DHCP ○ ARP ○ DNS ○ TCP ○ HTTP ○ None

Q4.4 (1 point) Which layer's header is used in computing the sequence number of the next packet that A sends out?

○ Layer 7 ○ Layer 4 ○ Layer 3 ○ Layer 2 ○ None

Q4.5 (1 point) In total, how many HTTP requests were sent by A in the process of loading the webpage?

○ Exactly 1, because A typed only one URL in its browser.

○ Exactly 2, because the DNS request added one more HTTP request.

○ 1 or more, because the webpage HTML could trigger more requests.

○ 2 or more, because of both the DNS request and the webpage HTML triggering more requests.

Q4.6 (1 point) True or false: If HTTP pipelining is **disabled**, then A initiates exactly 1 TCP handshake in the process of loading the webpage.

◯ True ◯ False

Q4.7 (1 point) True or false: If HTTP pipelining is **enabled**, then A initiates exactly 1 TCP handshake in the process of loading the webpage.

◯ True ◯ False

Q4.8 (2 points) For this subpart only, suppose R1 is using NAT (as shown in lecture) in its subnet.

Recall that NAT assigns private addresses to hosts inside the network. Which device(s) can deduce that NAT is in use by seeing a private address in the header of a packet (possibly a packet it is not supposed to be reading)? Select all that apply.

☐ A ☐ B ☐ C ☐ DNS ☐ D ◯ None

The rest of the subparts are independent of earlier subparts, i.e. we start over with empty caches.

Now, A joins the network and enters `https://www.d.com` to form an HTTPS connection to D.

Here are some reminders about TLS from lecture:
- HTTPS is a Layer 7 protocol running on top of TLS.
- TLS is a Layer 4.5 protocol running on top of TCP.
- TLS encrypts its entire payload, and adds an unencrypted header.

Q4.9 (1 point) When A visits `https://www.d.com`, in what order does A run these protocols?

◯ TCP handshake, then TLS handshake, then DNS lookup.

◯ DNS lookup, then TCP handshake, then TLS handshake.

◯ DNS lookup, then TLS handshake, then TCP handshake.

◯ TLS handshake, then TCP handshake, then DNS lookup.

Q4.10 (2 points) During the TLS connection, what values can R2 learn, possibly by reading packets it is not supposed to be reading? Select all that apply.

☐ R2 can read the records in the DNS response sent before the TLS connection.

☐ R2 can learn A's IP address.

☐ R2 can learn D's IP address.

☐ R2 can learn the TCP initial sequence numbers.

☐ R2 can read the HTTPS payloads.

☐ R2 can read the status codes in the HTTP responses (e.g. 404 File Not Found).

◯ None of the above

Q5 *Datacenters* **(11 points)**

Consider the topology below. All links have a bandwidth of 1 Gbps.

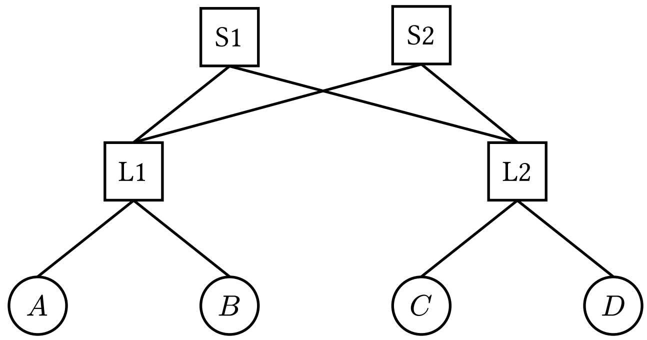

Q5.1 (3 points) Server A hosts VM 1 (Overlay: 10.0.0.1, ID: 50) and Server D hosts VM 2 (Overlay: 10.0.0.2, ID: 50).

Suppose an attacker compromises S1 and runs a packet sniffer. If VM 1 sends a packet to VM 2, which of the following addresses will the attacker observe in the captured packet's headers? Select all that apply.

- ☐ Underlay IP of A
- ☐ Underlay IP of D
- ☐ Overlay IP of VM 1 (10.0.0.1)
- ☐ Overlay IP of VM 2 (10.0.0.2)
- ☐ Underlay MAC address of A
- ☐ Underlay IP of L1
- ◯ None of the above

Q5.2 (3 points) VM 1 (ID: 50) is moved from A to C.

Assuming an SDN-managed overlay network, which of the following *must* occur for VM 1 to successfully resume receiving packets at its new location? Select all that apply.

- ☐ VM 1 must broadcast a DHCP request to obtain a new overlay IP address.
- ☐ Underlay routers (L1, L2, S1, S2) must recalculate routing tables to locate VM 1's overlay IP.
- ☐ The virtual switch on C must be updated with a new forwarding rule for VM 1.
- ☐ The SDN controller must update the mapping of VM 1's overlay IP to C's underlay IP.
- ☐ The virtual switch on A must forward all future datacenter traffic to C.
- ◯ None of the above

Q5.3 (2 points) Suppose the switches in the network use 5-tuple ECMP for load balancing.

VM 1 on A needs to transfer a massive dataset to VM 2 on D. The application opens **two** concurrent TCP connections to maximize throughput.

Assuming no other traffic in the network, what is the throughput the application will achieve, and why?

- ◯ 2 Gbps, because ECMP will route one connection through S1 and the other through S2.
- ◯ 1 Gbps, because ECMP hashes both connections to the same path since they are between the same source-destination pair.
- ◯ It has a 50% chance of achieving 2 Gbps and a 50% chance of achieving 1 Gbps, depending on the ECMP hash results.
- ◯ None of the above

Q5.4 (2 points) Suppose we want to run a collective operation among four VMs, one located on each of the four servers (A, B, C, D).

To optimize performance, the VMs will logically arrange themselves into a ring topology to pass data. Which logical ring minimizes traffic across the network?

If multiple options are equally optimal, select all the optimal options.

- ☐ $A \to B \to C \to D \to A$
- ☐ $A \to C \to B \to D \to A$
- ☐ $A \to C \to D \to B \to A$
- ◯ None of the above

Q5.5 (1 point) What is the bisection bandwidth of this topology, in Gbps?

◯ 1 ◯ 2 ◯ 3 ◯ 4 ◯ 6 ◯ 8

Q6 *Multicast* **(7 points)**

Consider running DVMRP on the topology below. Each host belongs to group G1, or group G2, or both.

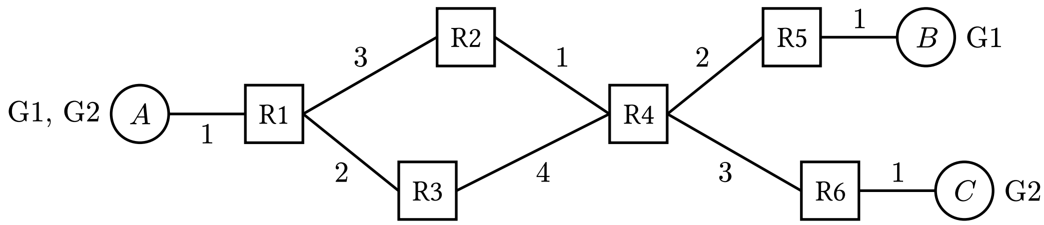

Fill in the DVMRP multicast table at R4, **before any pruning** takes place. Each row is worth 0.5 points.

- For example, if packets from A to G1 are forwarded to R2 and R5, select "R2" and "R5" in the first row.
- If a row should not be included in the table, select "None" for that row.

	R4's Multicast Table Before Pruning						
	Source	**Destination**	**Next-hop(s)**				
Q6.1	A	G1	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.2	A	G2	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.3	B	G1	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.4	B	G2	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.5	C	G1	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.6	C	G2	☐ R2	☐ R3	☐ R5	☐ R6	○ None

Next, fill in the DVMRP multicast table at R4, **at convergence**. Each row is worth 0.5 points.

	R4's Multicast Table at Convergence						
	Source	**Destination**	**Next-hop(s)**				
Q6.7	A	G1	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.8	A	G2	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.9	B	G1	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.10	B	G2	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.11	C	G1	☐ R2	☐ R3	☐ R5	☐ R6	○ None
Q6.12	C	G2	☐ R2	☐ R3	☐ R5	☐ R6	○ None

Q6.13 (1 point) Regardless of your answers to the previous parts, suppose you filled in 20 options in the first table, and 15 options in the second table. Based on these numbers, how many distinct pruning messages were sent? (Ignore periodically resending messages.)

○ 0 ○ 5 ○ 10 ○ 15 ○ 20 ○ Not enough information

Q7 *Collectives* **(11 points)**

Consider the following computation problem on N nodes:

Inputs:

- Node 1 has an array of N elements, where every element is either **`include`** or **`exclude`**.
- The remaining nodes (2 through N) each have an N-integer array.

Outputs:

- If the ith element of Node 1 is **`exclude`**, then the corresponding element in the result should be 0.
- If the ith element of Node 1 is **`include`**, then the corresponding element in the result should be the **product** of the ith elements of every other node.
- After the operation completes, **every node** should have a copy of the N-integer result array.

An example with $N = 4$ is shown below. Your answers should work for any input, not just the example.

Node 1	Node 2	Node 3	Node 4		Result
`exclude`	7	9	5	$\rightarrow$	0
`include`	2	3	10		60
`exclude`	5	1	4		0
`include`	3	4	2		24

To complete this operation, you will first need to convert the values in Node 1 to integers.

Q7.1 (1 point) What integer value should each **`exclude`** be converted to?

○ $-N$ ○ -1 ○ 0 ○ 1 ○ $N-1$ ○ N

Q7.2 (1 point) What integer value should each **`include`** be converted to?

○ $-N$ ○ -1 ○ 0 ○ 1 ○ $N-1$ ○ N

Q7.3 (1 point) After converting the values in Node 1 to integers, which collective operation can be used to complete the operation?

○ Broadcast ○ Gather ○ Reduce ○ ReduceScatter

○ Scatter ○ AllGather ○ AllReduce

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

Consider a k-ary fat-tree (folded Clos) topology, where:

- k is even,
- there are k pods,
- each pod contains $\left(\frac{k}{2}\right)^2$ hosts,
- each host has uplink bandwidth R,
- and the pods are split evenly into a left half and a right half.

As a reminder, the $k = 4$ topology from lecture is shown below:

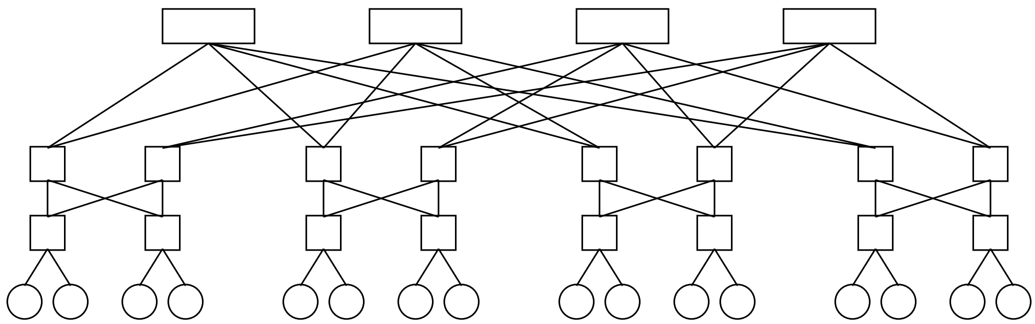

Q7.4 (2 points) How many hosts are in the left half of the datacenter?

○ $\frac{k^2}{4}$ ○ $\frac{k^3}{16}$ ○ $\frac{k^3}{8}$ ○ $\frac{k^3}{4}$ ○ $\frac{k^2}{2}$

Q7.5 (2 points) During one phase of a collective operation, the datacenter runs an all-to-all communication pattern: every host sends one packet to every other host in the datacenter.

In this all-to-all phase, how many packets cross from the left half of the datacenter to the right half?

○ $\frac{k^3}{8}$ ○ $\frac{k^3}{4}$ ○ $\frac{k^6}{128}$ ○ $\frac{k^6}{64}$ ○ $\frac{k^6}{32}$

Q7.6 (2 points) For this subpart only, assume the topology's bisection bandwidth is $\left(\frac{k^3}{8}\right)R$.

If one link of bandwidth R that crosses the bisection cut fails, what is the new worst-case oversubscription ratio of the topology?

○ $1 : 1$ ○ $k^2 : (k^2 - 4)$ ○ $k^3 : (k^3 - 8)$ ○ $k^3 : (k^3 - 4)$

Q7.7 (2 points) Regardless of your previous answers, assume a full k-ary fat-tree can support $\frac{k^3}{4}$ hosts with no oversubscription.

What is the smallest even value of k that can support at least 200 hosts?

○ 4 ○ 6 ○ 8 ○ 10

Q8 *Potpourri* **(8 points)**

Q8.1 (1 point) In Project 3 (Transport), why did we call `set_pending_ack()` instead of immediately sending out the ack?

- ◯ To support piggybacking.
- ◯ To implement flow control.
- ◯ To support pipelining.
- ◯ To pick random initial sequence numbers.

Q8.2 (2 points) In Project 3 (Transport), Stage 3.2 processed any in-order segments from the receive queue. What does this line of code (provided in the starter code as a hint) do?

```
data = packet.app[self.rcv.nxt |MINUS| packet.tcp.seq:]
```

- ◯ Extracts all of the payload in the TCP packet.
- ◯ Extracts the TCP payload, discarding bytes overlapping with already-received bytes.
- ◯ Checks if the next packet in the receive queue is in-order.
- ◯ Removes bytes from the receive queue if they overlap with already-received bytes.

Q8.3 (2 points) If we use Core-Based Trees (CBT) for multicast routing, which columns are needed in a router's multicast routing table? Select all that apply.

- ☐ Source
- ☐ Destination group
- ☐ Next-hops
- ◯ None

Q8.4 (1 point) True or false: When RDMA completes a data transfer, an interrupt is triggered, which tells the recipient's CPU to process the received data.

- ◯ True
- ◯ False

Q8.5 (2 points) Which of the following benefits does RDMA provide? Select all that apply.

- ☐ Lower latency
- ☐ Easier deployment of new protocols
- ☐ Reduced CPU usage for host networking
- ◯ None of the above

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.