CS 168 Spring 2025

Introduction to the Internet 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:	20	15	10	11	10	12	14	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 *Potpourri* **(20 points)**

For the next two subparts: Recall that in Project 3 (Transport), the starter code sets `self.snd.wnd = self.TX_DATA_MAX`, where `self.TX_DATA_MAX` is some large hard-coded constant. In Stage 5, you updated this code to change how the window size is set.

Q1.1 (3 points) If we didn't implement Stage 5, and instead left the starter code unchanged, what would happen in the resulting TCP implementation? Select all that apply.

- ☐ The receiver might be overwhelmed with too many out-of-order packets.
- ☐ The network might be overwhelmed with too many packets.
- ☐ The sender would always have to wait for packet i to be acked before sending packet $i + 1$.
- ◯ None of the above

Q1.2 (2 points) In Stage 5, what did you set `self.snd.wnd` to, and why?

- ◯ A value reported from the other side, for flow control.
- ◯ A value reported from the other side, for congestion control.
- ◯ A value computed by your code, for flow control.
- ◯ A value computed by your code, for congestion control.

Q1.3 (2 points) What does a TCP receiver do when it receives 2 identical duplicate packets, both with sequence number 50?

- ◯ Send one packet with ack number 50.
- ◯ Send one packet with ack number 51.
- ◯ Send two packets with ack number 50.
- ◯ Send two packets with ack number 51.

Q1.4 (1 point) When deploying CDNs, each CDN server must contain a complete copy of all the resources on the origin server.

- ◯ True
- ◯ False

Q1.5 (1 point) Which HTTP header helps ensure that the cache keeps up-to-date copies of cached resources?

- ◯ Cache-Control
- ◯ Content-Type
- ◯ Version
- ◯ User-Agent

Q1.6 (1 point) To access a **private** HTTP cache, the user must send at least one packet through the network.

- ◯ True
- ◯ False

Q1.7 (1 point) To access a **proxy** HTTP cache, the user must send at least one packet through the network.

- ◯ True
- ◯ False

Q1.8 (1 point) How many physical DNS root name servers exist?

◯ 1 ◯ 2 ◯ 13 ◯ More than 13

Q1.9 (1 point) Consider a router hashing each connection's 5-tuple to decide how to forward the packet.

If there are two shortest paths to a given destination, this router guarantees that exactly half of the packets take one path, and half of the packets take the other path.

◯ True ◯ False

In the next two subparts, consider a datacenter using overlay and underlay networks to forward packets. Server Y is some arbitrary server somewhere in the datacenter.

Q1.10 (2 points) What IP addresses can be found in the forwarding table(s) of a **virtual switch** on Server Y? Select all that apply.

☐ The physical IP of Server Y. ☐ Virtual IPs of VMs on some other servers.

☐ Physical IPs of some other servers. ◯ None of the above

☐ The virtual IP of all VMs on Server Y.

Q1.11 (2 points) What IP addresses can be found in the forwarding table(s) of a **physical router** directly connected to Server Y? Select all that apply.

☐ The physical IP of Server Y. ☐ Virtual IPs of VMs on some other servers.

☐ Physical IPs of some other servers. ◯ None of the above

☐ The virtual IP of all VMs on Server Y.

Q1.12 (1 point) In software-defined networking, an OpenFlow table can be used to program destination-based forwarding rules into a router.

◯ True ◯ False

Q1.13 (1 point) In a network with low-bandwidth links, offloading operations to the network interface card (NIC) will result in significant performance benefits.

◯ True ◯ False

Q1.14 (1 point) RDMA uses the operating system's TCP implementation to ensure reliability.

◯ True ◯ False

Q2 *TCP Congestion Control* **(15 points)**

In this question, you are sending data using the TCP congestion control algorithm seen in lecture.

Assumptions for the entire question:
- All TCP values are measured in packets (not bytes), unless otherwise specified.
- The sender always has new data to send, and RWND is very large.

Q2.1 (1 point) Immediately after the TCP handshake, suppose you set CWND = 5 packets.

In general (not necessarily for this flow), why might senders initialize CWND to 5, instead of 1?

◯ To prevent cheating.
◯ To improve performance for short flows.
◯ To distinguish congestion and corruption.
◯ To improve performance for long flows.

Immediately after the TCP handshake, you are in Slow Start mode, with the following settings:
- CWND = 5, SSTHRESH = ∞
- The sender's first data packet has sequence number #20.

Q2.2 (2 points) Suppose you receive an ack for packet #20 (i.e. the receiver has received packet #20, and sends you an ack).

At this point, which packets are allowed to be in flight? Select all that apply.

☐ #20 ☐ #22 ☐ #24 ☐ #26 ☐ #28 ☐ #30
☐ #21 ☐ #23 ☐ #25 ☐ #27 ☐ #29 ☐ #31

Q2.3 (2 points) Suppose you then receive an ack for packet #21.

At this point, which packets are allowed to be in flight? Select all that apply.

☐ #20 ☐ #22 ☐ #24 ☐ #26 ☐ #28 ☐ #30
☐ #21 ☐ #23 ☐ #25 ☐ #27 ☐ #29 ☐ #31

The rest of this question is independent of earlier subparts. Each subpart continues on from the previous one.

Some time later, you are in Slow Start mode, with the following settings:

- CWND = 14, SSTHRESH = ∞
- All packets up to, and including packet #30, have been sent and acked.
- Packets #31 through #44, inclusive, have been sent, but not acked.
- All packets #45 and later have not been sent.

Suppose that packet #31 is dropped in transit. All other packets are successfully sent, and the receiver sends you acks for packets #32, #33, #34, etc., in that order, with no timeouts.

Q2.4 (1 point) In this scenario, TCP will switch from Slow Start mode to Fast Recovery mode.

TCP switches to Fast Recovery mode immediately after you receive the ack corresponding to which packet?

◯ #32 ◯ #33 ◯ #34 ◯ #35 ◯ #36 ◯ #37

Q2.5 (2 points) What is the value of CWND the instant **after** TCP switches from Slow Start mode to Fast Recovery mode?

Reminder: During Slow Start mode, duplicate acks do not change the value of CWND.

◯ 7 ◯ 8 ◯ 9 ◯ 10 ◯ 11 ◯ 12

Reminder: You are now in Fast Recovery mode.

Q2.6 (2 points) The receiver receives packets in this order: #32, #33, #34, ..., #43, #44, #31 (retransmitted).

The receiver sends acks for these packets, and you receive these acks in order, with no timeouts.

What is the value of CWND the instant **before** TCP switches out of Fast Recovery mode?

◯ 7 ◯ 10 ◯ 13 ◯ 15 ◯ 17 ◯ 20

Q2.7 (2 points) What is the value of CWND the instant **after** TCP switches out of Fast Recovery mode?

◯ 7 ◯ 10 ◯ 13 ◯ 15 ◯ 17 ◯ 20

Q2.8 (2 points) At the instant **before** TCP switches out of Fast Recovery mode, which packets have been sent out?

All packets up to, and including _____, have been sent out.

◯ #38 ◯ #41 ◯ #44 ◯ #47 ◯ #50 ◯ #53

Q2.9 (1 point) After TCP switches out of Fast Recovery mode, what mode is TCP in?

◯ Slow Start ◯ Congestion Avoidance

Q3 *DNS* **(10 points)**

Consider the following DNS name server hierarchy. Each box represents a zone. Each box contains the domains and corresponding IP addresses for all name servers that are authoritative for that zone. Assume that no other zones exist, besides the ones shown.

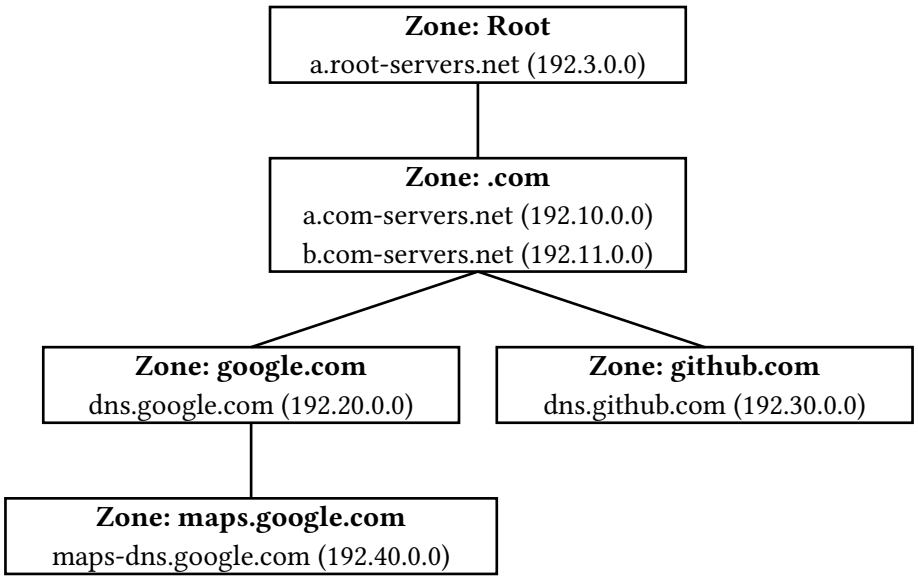

For each of the following records (name, type, value), select the zone that could provide the given record.

If the given record is invalid, or no zone would provide the given record, select "None."

Q3.1 (1 point) `classroom.github.com A 142.250.9.138`

- ○ Root
- ○ google.com
- ○ github.com
- ○ .com
- ○ maps.google.com
- ○ None

Q3.2 (1 point) `google.com NS dns.google.com`

- ○ Root
- ○ google.com
- ○ github.com
- ○ .com
- ○ maps.google.com
- ○ None

Q3.3 (1 point) `google.com NS maps-dns.google.com`

- ○ Root
- ○ google.com
- ○ github.com
- ○ .com
- ○ maps.google.com
- ○ None

Q3.4 (1 point) `maps-dns.google.com A 192.40.0.0`

- ○ Root
- ○ google.com
- ○ github.com
- ○ .com
- ○ maps.google.com
- ○ None

In the next three subparts, a modification is given. Select the zone that needs to update its record(s) to process this modification. If no records need to be updated, select "None."

Q3.5 (2 points) The IP address of images.google.com is changed from 142.250.189.14 to 142.250.189.15.

◯ Root	◯ google.com	◯ github.com
◯ .com	◯ maps.google.com	◯ None

Q3.6 (2 points) A new name server, c.com-servers.net (192.12.0.0), is added. This name server is authoritative for the .com zone.

◯ Root	◯ google.com	◯ github.com
◯ .com	◯ maps.google.com	◯ None

Q3.7 (2 points) A new mirror name server for the .com zone is installed. Using anycast, this new name server has IP address 192.10.0.0.

◯ Root	◯ google.com	◯ github.com
◯ .com	◯ maps.google.com	◯ None

Q4 *End-to-End* **(11 points)**

Consider a subnet with router R1, and hosts A and B, all connected on a single shared medium (also called a bus).

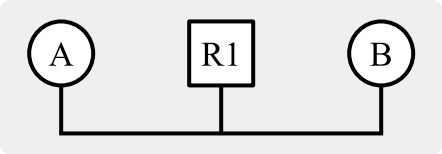

In this question, each subpart continues on from the previous one.

R1 is using NAT with Port Address Translation, as shown in lecture, with the following addresses:
- R1's public address is 143.3.4.5, and its private address is 192.168.1.1.
- R1 can allocate private addresses 192.168.1.2, 192.168.1.3, and so on, to hosts.
- The DNS recursive resolver has IP 8.8.8.8.

Q4.1 (2 points) Host A joins the network, with all caches empty and no active connections.

In Host A's DHCP Discover request, what is the destination IP address?

◯ 8.8.8.8 ◯ 255.255.255.255 ◯ 143.3.4.5 ◯ 192.168.1.1

Q4.2 (2 points) In R1's DHCP Offer to Host A, what is the subnet in the offer?

◯ 0.0.0.0/0 ◯ 192.168.1.0/24 ◯ 192.168.1.1/32 ◯ 143.3.4.0/24

At this point, Host A has completed the DHCP handshake, and has been assigned IP address 192.168.1.2. Host A has not initiated any connections yet.

Host A types www.berkeley.edu in their browser. In the next 3 subparts, fill out the fields of the first IP packet that Host A creates and sends as a result.

Q4.3 (1 point) Source IP:

◯ 8.8.8.8 ◯ 143.3.4.5 ◯ 192.168.1.1 ◯ 192.168.1.2

Q4.4 (1 point) Destination IP:

◯ 8.8.8.8 ◯ 143.3.4.5 ◯ 192.168.1.1 ◯ 192.168.1.2

Q4.5 (2 points) Destination MAC:

◯ A's MAC ◯ R1's MAC ◯ berkeley.edu's MAC

◯ B's MAC ◯ DNS server's MAC ◯ `ff:ff:ff:ff:ff:ff`

The first IP packet that Host A creates eventually reaches R1. At this point, R1 possibly rewrites some of the header fields, and then forwards the packet. In the next two subparts, fill out the fields of the packet that R1 sends out (possibly with rewritten headers).

Q4.6 (1 point) Source IP:

◯ 8.8.8.8 ◯ 143.3.4.5 ◯ 192.168.1.1 ◯ 192.168.1.2

Q4.7 (1 point) Destination IP:

◯ 8.8.8.8 ◯ 143.3.4.5 ◯ 192.168.1.1 ◯ 192.168.1.2

Q4.8 (1 point) Will the IP address 192.168.1.2 always be associated with Host A?

◯ Yes ◯ No

Q5 *STP* **(10 points)**

Consider running the Spanning Tree Protocol (STP) for the network topology to the right.

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

Assume the links with no label have a cost of 1.

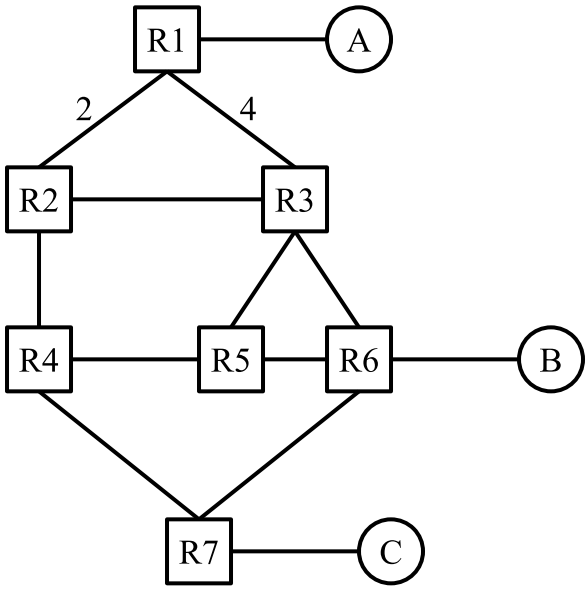

For each of the next six subparts, select the link disabled by the given router. Option "R1" means "the link to R1," and likewise for other options.

If the router does not disable any link, select "None."

Q5.1 (1 point) Which link (if any) does R2 disable?

○ None ○ R1 ○ R3 ○ R4

Q5.2 (1 point) Which link (if any) does R3 disable?

○ None ○ R1 ○ R2 ○ R5 ○ R6

Q5.3 (1 point) Which link (if any) does R4 disable?

○ None ○ R2 ○ R5 ○ R7

Q5.4 (1 point) Which link (if any) does R5 disable?

○ None ○ R3 ○ R4 ○ R6

Q5.5 (1 point) Which link (if any) does R6 disable?

○ None ○ R3 ○ R5 ○ R7

Q5.6 (1 point) Which link (if any) does R7 disable?

○ None ○ R4 ○ R6

Suppose STP has converged. Regardless of your answers to the previous subparts, assume that the following 4 links are disabled (also shown in the diagram to the right):

- R1-to-R3
- R4-to-R5
- R5-to-R6
- R4-to-R7

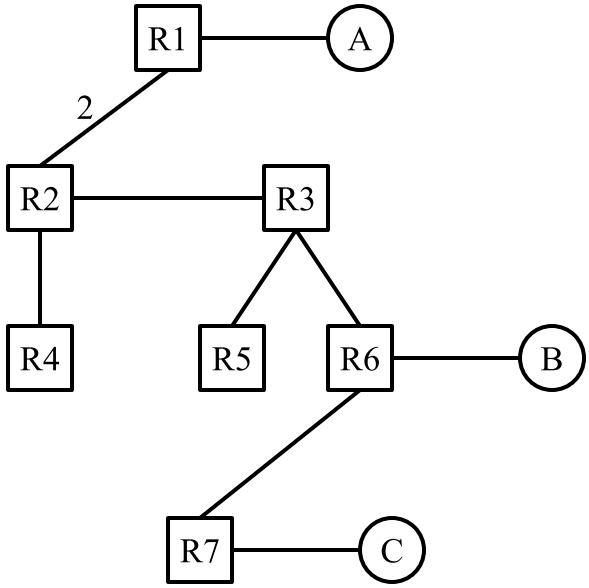

Switches R1 to R7 are all learning switches.

All forwarding tables start out empty.

In each of the next two 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.

Q5.7 (2 points) A sends a packet to C.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6 ☐ R7

Q5.8 (2 points) C sends a packet to A.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6 ☐ R7

Q6 *Datacenters* **(12 points)**

Consider the Clos-like topology below.

- Each rack has 2 servers, and each server has its own link to the adjacent router. In other words, each rack has two links to its adjacent router.
- The bandwidth of each link, and the line rate of each server, are all equal.
- Racks 1, 2, 3 (with 6 servers in total) are on the left side. Racks 4, 5, 6 (with 6 servers in total) are on the right side.

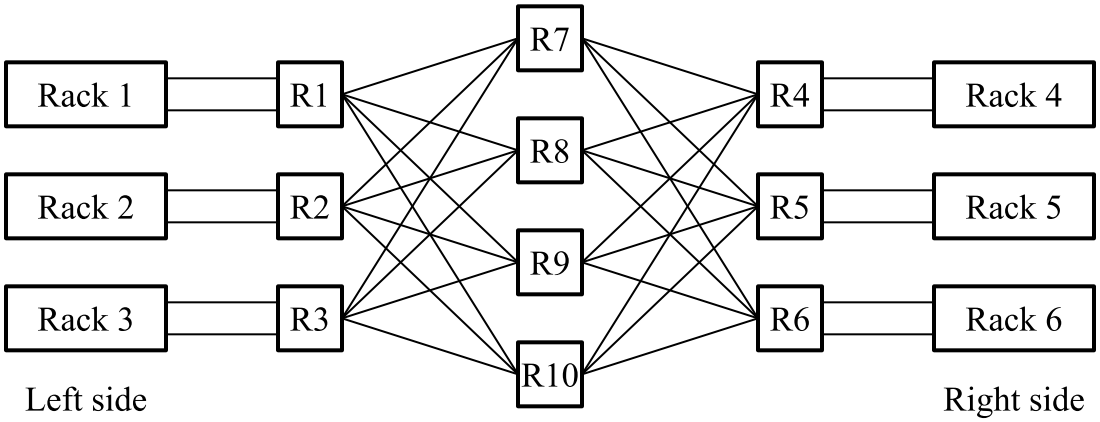

Q6.1 (1 point) Each server on the left side can send data to a corresponding server on the right side, at full line rate.

(In other words: We can create 6 connections, each sending at full line rate, where the 6 left-side servers each participate in one connection, and the 6 right-side servers each participate in one connection.)

◯ True ◯ False

Q6.2 (1 point) For this subpart only, suppose each rack has 3 servers, instead of 2 servers. Each of the 3 servers still has its own link to the adjacent router.

Each server on the left side can send data to a corresponding server on the right side, at full line rate.

◯ True ◯ False

Q6.3 (2 points) What is the maximum number of servers per rack, such that each server on the left side can send data to a corresponding server on the right side, at full line rate?

Assume that each server still has its own link to the adjacent router.

Your answer should be a single integer.

For the rest of the question, suppose we generalize the topology:

- There are R racks on the left side, and R racks on the right side. ($R = 3$ in the example.)
- There are S servers per rack. ($S = 2$ in the example.)
- There are k servers in the middle layer. ($k = 4$ in the example.)

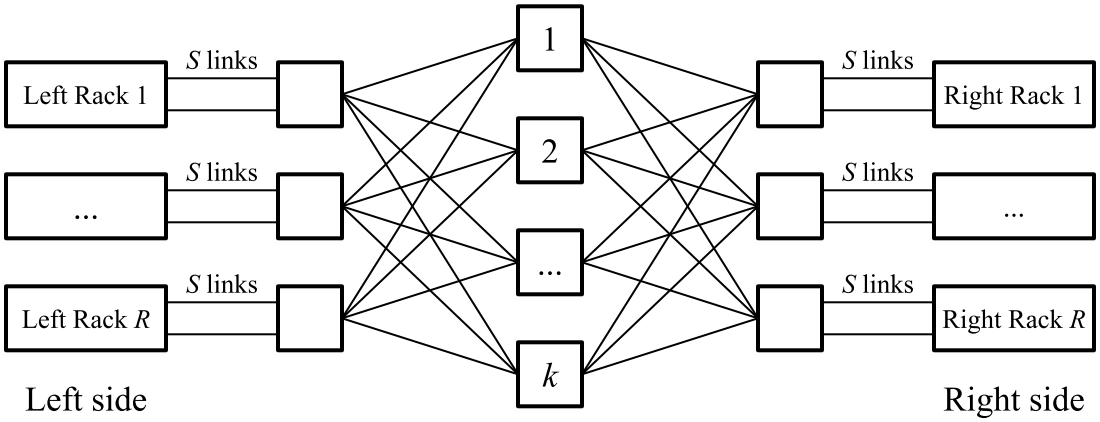

Q6.4 (2 points) What values of k allow each server on the left side to send data to a corresponding server on the right side, at full line rate?

Fill in the inequality. Your answer could be in terms of R, S, and/or k.

$k \geq$

Q6.5 (3 points) How many total links are used to build this topology?

Your answer could be in terms of R, S, and/or k.

Reminder: The radix of a switch is the number of ports that switch has.

Assume that each switch has exactly the number of ports needed for the topology (i.e. no additional unused ports).

Q6.6 (1 point) The radix of every switch in this topology is the same.

◯ True ◯ False

Q6.7 (2 points) What value of k causes every switch in the topology to have the same radix?

Fill in the expression. Your answer could be in terms of R, S, and/or k.

$k =$

Q7 *Multicast* **(14 points)**

Consider the topology below. All unlabeled links cost 1. Each host belongs to either group G1, or group G2, or both.

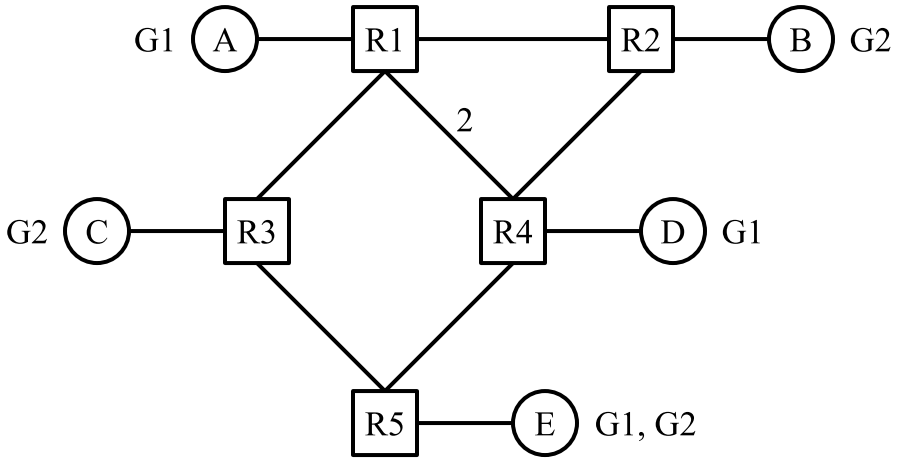

In the next four subparts, E wants to send a multicast packet to all other members of G1, using DVMRP.

Q7.1 (3 points) Before any pruning is performed, which links will be used to forward this packet? Select all that apply.

- ☐ R1—R2
- ☐ R1—R4
- ☐ R3—R5
- ☐ R1—R3
- ☐ R2—R4
- ☐ R4—R5

When building the E-to-G1 DVMRP tree, _____ (i) sends a prune message to _____ (ii).

Q7.2 (1 point) Blank (i): Who sends the prune message?

- ○ R1
- ○ R2
- ○ R3
- ○ R4
- ○ R5
- ○ R6

Q7.3 (1 point) Blank (ii): Who is the prune message sent to?

- ○ R1
- ○ R2
- ○ R3
- ○ R4
- ○ R5
- ○ R6

Q7.4 (2 points) At convergence, is R3 part of the E-to-G1 DVMRP tree? In other words, is R3 forwarding packets from E to G1?

- ○ Yes, because R3 has a child who is not pruned.
- ○ Yes, because R3 is directly-connected to a G1 member.
- ○ No, because all of R3's children have been pruned.
- ○ No, because R3 is not directly-connected to any G1 member.

The diagram, reprinted for your convenience:

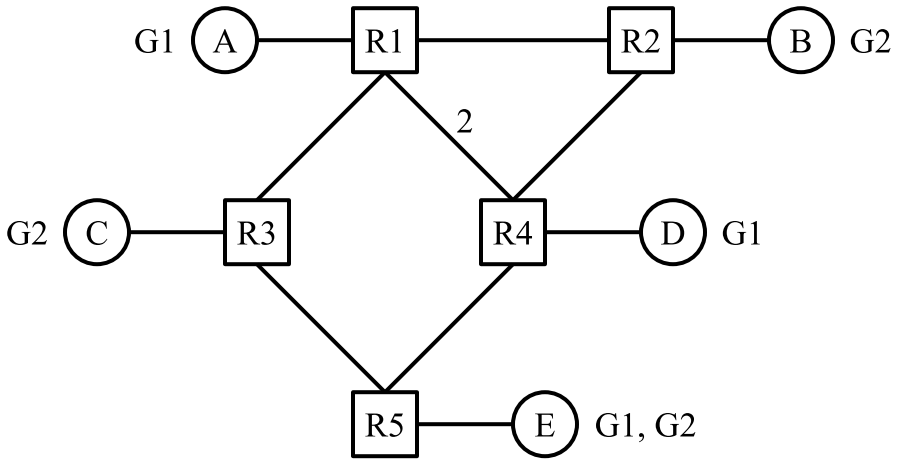

Q7.5 (2 points) In this subpart only: Suppose that some time later, B decides to join group G1.

According to the DVMRP protocol from lecture, when will B start to receive multicast packets sent to G1?

- ◯ Immediately after B joins G1.
- ◯ Immediately after R2 learns that B joined G1.
- ◯ Immediately after pruning state is cleared at all the routers.
- ◯ B will never receive multicast packets sent to G1.

In the next two subparts, C wants to send a multicast packet to all other members of G2, using CBT.

Suppose R4 is chosen as the core, and the routing state has converged (i.e. all G2 members have sent Join messages to the core).

Q7.6 (3 points) Which links will be used to forward this packet? Select all that apply.

☐ R1—R2	☐ R1—R4	☐ R3—R5
☐ R1—R3	☐ R2—R4	☐ R4—R5

Q7.7 (2 points) When R4 is the core, which statement is true about the paths the packet takes from C to all other G2 members?

- ◯ The packet takes the shortest paths to all G2 members.
- ◯ The packet takes the shortest paths to some, but not all, G2 members.
- ◯ The packet does not take the shortest path to any G2 members.

Q8 *Collectives* **(8 points)**

Recall the Broadcast and Reduce collective operations from lecture, with Node 1 as the root node:

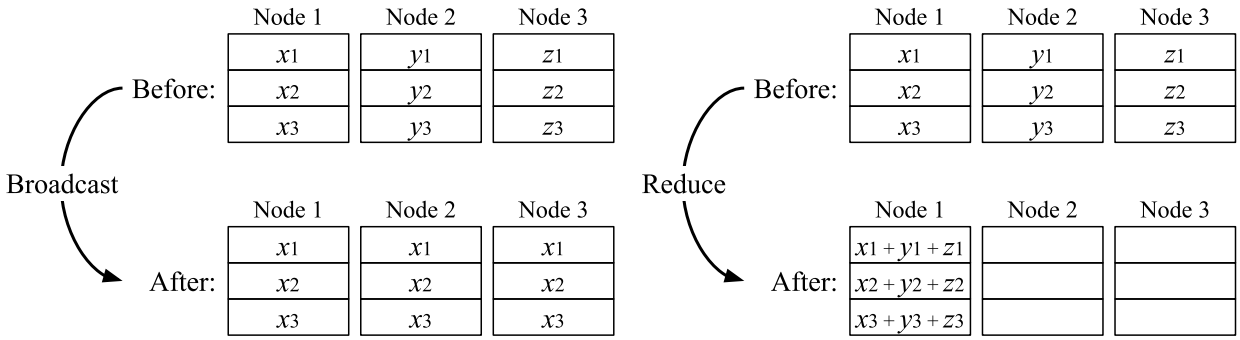

Q8.1 (3 points) Suppose we start with the "Before" state shown above. We run a Broadcast operation, immediately followed by a Reduce operation (using the Broadcast output as the input to Reduce).

What are the resulting values at Node 1?

Write one expression per box. Your expression could be in terms of: $x_1, x_2, x_3, y_1, y_2, y_3, z_1, z_2, z_3$.

Node 1:

Q8.2 (1 point) Are the Broadcast and Reduce operations duals of each other?

◯ Yes ◯ No

In the next two subparts, we connect the 3 nodes in a ring topology, and implement Broadcast and Reduce using a similar approach as naive ring-based AllReduce. Assume each vector (e.g. $[x_1, x_2, x_3]$) is D bytes.

Your answers below can be an expression, possibly in terms of D. Give exact answers (not big-O bounds). Count all data the node sends across all time steps, but don't count data received.

Q8.3 (2 points) To implement the **Broadcast** operation on 3 nodes, what is the maximum amount of data sent by any single node?

Q8.4 (2 points) To implement the **Reduce** operation on 3 nodes, what is the maximum amount of data sent by any single node?

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.