CS 168
Fall 2024

Intro to the Internet: Architecture and Protocols

Final Exam

Name: ______________________________

Student ID: ______________________________

This exam is 170 minutes long. There are 9 questions, worth 100 points.

In addition, Question 10 is a redemption question from the midterm (see Q10 for more details).

Question:	1	2	3	4	5	6	7	8	9	Total
Points:	18	14	11	13	16	11	5	7	5	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)
- ☑ 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 *Quick Questions* (18 points)

Each subpart is 1 point, unless otherwise specified.

Q1.1 When an error occurs, HTTP requires servers to send the exact correct error code.

- ○ True
- ○ False

In the next 3 subparts, consider TCP with no congestion control, as implemented in the project.

Q1.2 Both active close and passive close require you to send a FIN packet.

- ○ True
- ○ False

Q1.3 Suppose you set RWND (receiver window, `snd.wnd` in the project) to a very large value, instead of the value reported by the other side. What issue could occur?

- ○ You may run out of buffer space.
- ○ The other side may run out of buffer space.
- ○ The other side will have fewer in-flight packets (from the other side to you).
- ○ The other side will have more in-flight packets (from the other side to you).

Q1.4 Suppose you report a very large value of RWND to the other side. This causes the other side to set its RWND (`snd.wnd`) to a very large value. What issue could occur?

- ○ You may run out of buffer space.
- ○ The other side may run out of buffer space.
- ○ You will have fewer in-flight packets (from you to the other side).
- ○ You will have more in-flight packets (from you to the other side).

In the next two subparts, our goal is to add a cache to reduce the total amount of bandwidth used on the network (compared to when we had no cache).

Q1.5 A **private** cache only achieves our goal when a single client tries to access a resource multiple times.

- ○ True
- ○ False

Q1.6 A **proxy** cache only achieves our goal when a single client tries to access a resource multiple times.

- ○ True
- ○ False

Q1.7 In TLS, the TLS header is included as part of the TCP payload.

- ○ True
- ○ False

Q1.8 Consider an application provider using a CDN. The provider uses application-level mapping to redirect clients to CDN servers.

The provider can always identify a single CDN server that is closest to the client, and redirect the client to that closest CDN server.

○ True ○ False

Q1.9 In host networking for datacenter servers, basic operations like checksum verification and fragmentation have been offloaded from software to hardware.

○ True ○ False

Q1.10 In host networking, offloading is done for improved performance, not improved correctness.

○ True ○ False

Q1.11 In RDMA, when a job is scheduled and a WQE (Work Queue Element) is created, the NIC (Network Interface Card) must stop what it is doing and process the job immediately.

○ True ○ False

Q1.12 In RDMA, when a job is completed and a CQE (Completion Queue Element) is created, the operating system must stop what it is doing and process the job immediately.

○ True ○ False

Q1.13 With SDN, individual routers no longer need the **data** plane.

○ True ○ False

Q1.14 With SDN, individual routers no longer need the **control** plane.

○ True ○ False

Q1.15 In a cellular network, two devices can have the same IMEI (device ID).

○ True ○ False

In the next two subparts, consider a new collective operation called GatherReduce. Using the representation from lecture, the GatherReduce operation is depicted as follows:

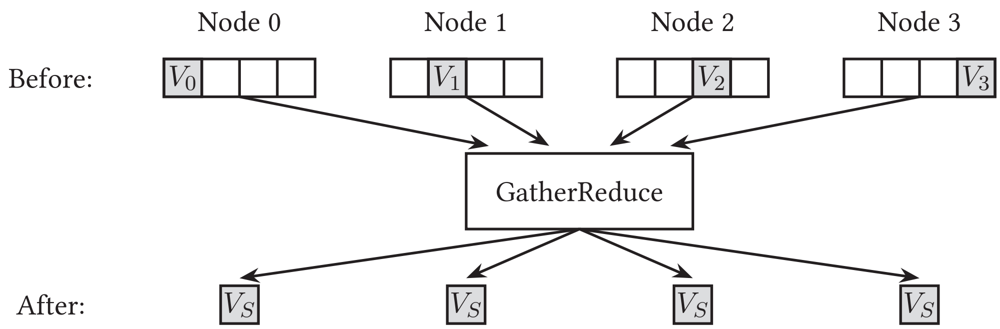

GatherReduce should compute $\boxed{V_S} = \boxed{V_0} + \boxed{V_1} + \boxed{V_2} + \boxed{V_3}$ and send this sum to each node.

Note: Each individual square, i.e. $\boxed{V_S}$, $\boxed{V_0}$, $\boxed{V_1}$, $\boxed{V_2}$, $\boxed{V_3}$, represents a vector of size D bytes.

We have three topologies to choose from:

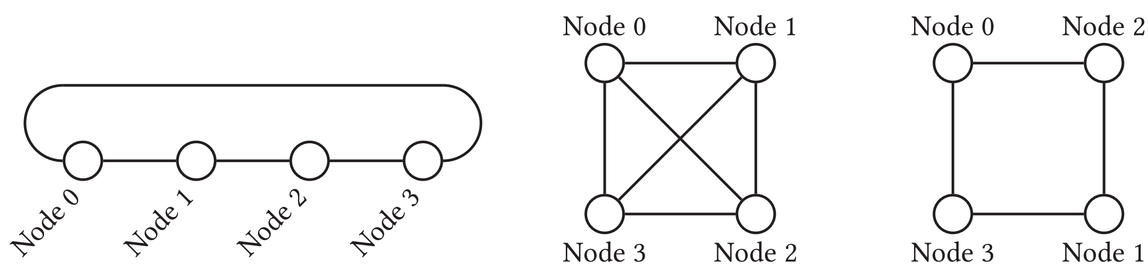

Topology A **Topology B** **Topology C**

We want to implement this collective, while satisfying the following efficiency constraint:
For each link X–Y, at most D bytes are sent from X to Y on this link, and at most D bytes are sent from Y to X on this link.

Q1.16 (2 points) Which topologies would allow us to implement the GatherReduce collective, while satisfying the efficiency constraint above? Select all that apply.

☐ Topology A ☐ Topology B ☐ Topology C ☐ None

Q1.17 Out of the topologies you chose in the previous subpart (i.e. the topologies satisfying the constraint), which topology allows GatherReduce to complete in the fastest time?

Assume all links are identical (i.e. same bandwidth, same propagation delay).

○ Topology A ○ Topology B ○ Topology C

○ Multiple topologies achieve the same completion time.

Q2 *TCP Congestion Control* **(14 points)**

Consider two hosts communicating using the TCP congestion control algorithm seen in class. For the entire question, assume all TCP values are measured in packets (not bytes), unless otherwise specified.

Reminders:

- The hosts are using fast retransmit and fast recovery.
- The sender always has new data to send, and RWND is very large.
- The algorithm moves from Slow Start to Congestion Avoidance when $\text{CWND} > \text{SSTHRESH}$.
- During Congestion Avoidance, for each new acked packet, $\text{CWND} \leftarrow \text{CWND} + \dfrac{1}{\text{CWND}}$.

Q2.1 (1 point) Immediately after the TCP handshake, what mode of TCP is active?

- ○ Slow Start
- ○ Congestion Avoidance
- ○ Fast Recovery

Q2.2 (1 point) Immediately after the TCP handshake, CWND = 1 packet.

After the 3-way TCP handshake, 3 distinct packets have been sent and 3 distinct, in-order acks have been received. At this point, what is the new value of CWND?

- ○ 1
- ○ 2
- ○ 3
- ○ 4
- ○ 8
- ○ 16

The following subparts are independent of the earlier subparts.

Some time later, you are in Congestion Avoidance mode, with the following settings:

- $\text{CWND} = 12$, $\text{SSTHRESH} = 16$
- All packets up to, and including packet #60, have been sent and acked.
- Packets #61 through #72, inclusive, have been sent, but not acked.
- All packets #73 and later have not been sent.

Q2.3 (1 point) For this subpart only, assume there is infinite bandwidth available in the network.

Suppose you receive distinct, in-order acks for packets #61, #62, #63, etc., and you send packets #73, #74, #75, etc. This continues for a long time, with no timeouts or duplicate acks.

As you continue to receive acks and send out more packets indefinitely, what happens?

- ○ As soon as CWND exceeds SSTHRESH, you switch to Slow Start mode.
- ○ CWND reaches SSTHRESH = 16, and stops increasing after that.
- ○ CWND increases indefinitely, with no limit, but SSTHRESH doesn't change.
- ○ As soon as CWND exceeds SSTHRESH, both CWND and SSTHRESH increase indefinitely.

Now, suppose that only packet #66 is dropped in transit. All other packets are successfully sent, and you receive distinct, in-order acks for all other packets, with no timeouts.

Q2.4 (1 point) In this scenario, TCP would eventually switch out of Congestion Avoidance mode, and into which mode?

- ○ Slow Start
- ○ Fast Recovery

Q2.5 (2 points) What is the approximate value of CWND the instant **before** TCP switches out of Congestion Avoidance mode? (Hint: You don't need to do any complicated arithmetic here.)

- ○ 1
- ○ $12\frac{5}{12}$
- ○ 16
- ○ 18
- ○ 384
- ○ $12\frac{4}{12}$
- ○ $12\frac{6}{12}$
- ○ 17
- ○ 192
- ○ 768

Q2.6 (2 points) What is the value of CWND the instant **after** TCP switches out of Congestion Avoidance mode, and into a different mode? Round your answer to the nearest integer.

- ○ 1
- ○ 6
- ○ 8
- ○ 9
- ○ 11
- ○ 12

Q2.7 (1 point) Which of these ack numbers, if received, would cause you to leave the different mode that you switched to?

- ○ An ack number less than 65
- ○ 66
- ○ 65
- ○ An ack number greater than 66

The following subparts are independent of the earlier subparts.

Some time later, you are in Congestion Avoidance mode, with the following settings:

- CWND = 6, SSTHRESH = 8
- All packets up to, and including packet #202, have been sent and acked.
- Packets #203 through #208, inclusive, have been sent, but not acked.
- All packets #209 and later have not been sent.

At this point, the sender stops receiving any further acks.

Q2.8 (1 point) In this scenario, TCP would eventually switch out of Congestion Avoidance mode, and into which mode?

○ Slow Start ○ Fast Recovery

Q2.9 (1 point) What is the value of CWND the instant after TCP switches out of Congestion Avoidance mode, and into a different mode?

○ 1 ○ 2 ○ 3 ○ 4 ○ 6 ○ 8

Q2.10 (1 point) What is the packet sent as a result of switching to a different mode?

○ #202 ○ #203 ○ #208 ○ #209

The following subparts are independent of the earlier subparts.

Some time later, the TCP connection has the following settings:

- CWND = 4, SSTHRESH = 8
- All packets up to, and including packet #411, have been sent and acked.
- Packets #412 through #415, inclusive, have been sent, but not acked.
- All packets #416 and later have not been sent.

For the next two subparts, you can assume CWND is rounded to the nearest integer when deciding which packets to send out.

The next two subparts are independent (i.e. Q2.12 is not continuing from Q2.11).

Q2.11 (1 point) If we are in **Congestion Avoidance** mode, and we receive an ack for packet #412, which packet(s) do we send out as a result? Select all that apply.

☐ #416 ☐ #417 ☐ #418 ☐ #419 ☐ None

Q2.12 (1 point) If we are in **Slow Start** mode, and we receive an ack for packet #412, which packet(s) do we send out as a result? Select all that apply.

☐ #416 ☐ #417 ☐ #418 ☐ #419 ☐ None

Q3 *TCP Performance* **(11 points)**

Consider accessing an application using a CDN server, as shown below:

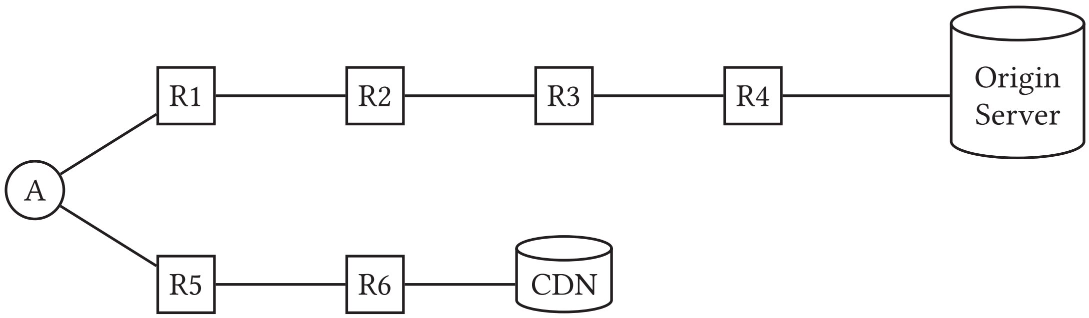

Unless otherwise specified, assume that:

- Each link has the same propagation delay. For example, the delay from A to the origin server is 5/3 times higher than the delay from A to the CDN server.
- Each link uses the standard MTU of 1500 bytes.
- Each link has a packet drop rate of 0.

In the next three subparts, consider the following two scenarios for downloading the same large file:

- **Scenario 1:** Host A downloads the entire file from the origin server.
- **Scenario 2:** Host A downloads the entire file from the CDN server.

Let S_1 be the throughput for Scenario 1, and S_2 be the throughput for Scenario 2.

According to the TCP throughput equation (on the reference card) and the assumptions in each subpart, select the equation that correctly relates the throughput of S_1 and S_2.

Q3.1 (2 points) For this subpart only: The A-to-R1 and A-to-R5 links each have a drop rate of 0.25, and all other links have a drop rate of 0. For example, a connection from A to the CDN has an overall packet drop rate of 0.25.

- ○ $S_1 = \frac{5}{3} \cdot S_2$
- ○ $S_1 = 2 \cdot S_2$
- ○ $S_2 = \sqrt{\frac{5}{3}} \cdot S_1$
- ○ $S_1 = \sqrt{\frac{5}{3}} \cdot S_2$
- ○ $S_2 = \frac{5}{3} \cdot S_1$
- ○ $S_1 = S_2$

Q3.2 (2 points) For this subpart only: The A-to-R1 and A-to-R5 links each have a drop rate of 0.25, and all other links have a drop rate of 0.

Also, for this subpart only: the MTU for all links between A and the origin server is increased to 3000 bytes. All other links have an MTU of 1500 bytes.

Assume the header size is negligible, i.e. you can use the MTU as an approximation for the MSS.

- ○ $S_1 = \frac{3}{2} \cdot S_2$
- ○ $S_1 = \frac{6}{5} \cdot S_2$
- ○ $S_2 = \frac{3}{2} \cdot S_1$
- ○ $S_1 = 2 \cdot S_2$
- ○ $S_2 = \frac{6}{5} \cdot S_1$
- ○ $S_1 = S_2$

Q3.3 (2 points) For this subpart only: The A-to-R1 link has a drop rate of 0.04, the A-to-R5 link has a drop rate of 0.25, and all other links have a drop rate of 0.

- ○ $S_1 = \sqrt{\frac{3}{2}} \cdot S_2$
- ○ $S_1 = \frac{8}{3} \cdot S_2$
- ○ $S_2 = \frac{8}{3} \cdot S_1$
- ○ $S_1 = \frac{3}{2} \cdot S_2$
- ○ $S_2 = \frac{3}{2} \cdot S_1$
- ○ $S_1 = S_2$

Q3.4 (2 points) For this subpart only: The A-to-R1 link has a drop rate that you will determine, the A-to-R5 link has a drop rate of 0.25, and all other links have a drop rate of 0.

What should the A-to-R1 link drop rate be, so that the throughput in the two scenarios ends up equaling the same?

- ○ 0.3
- ○ 0.09
- ○ 0.25
- ○ 0.5
- ○ 0.64
- ○ 0.81

Q3.5 (2 points) Consider a scenario where, for every user, downloading from the origin server is always faster than downloading from any CDN server.

In this scenario, are there any advantages to deploying the CDN?

If you say Yes, name an advantage. If you say No, explain why not. Answer in 10 words or fewer.

- ○ Yes
- ○ No

Q3.6 (1 point) Suppose there are many CDN servers in the network, and the application provider uses anycast to redirect users to the closest CDN server.

Why might anycast cause issues in this scenario?

- ○ Most links have a drop rate of 0.
- ○ Links use a standard MTU of 1500 bytes.
- ○ Each link has the same propagation delay.
- ○ The file we're downloading is large.

Q4 *DNS* (13 points)

Consider the following DNS name server hierarchy. Each node represents a zone. Each node contains the domains and corresponding IP addresses for all name servers that are authoritative for that zone.

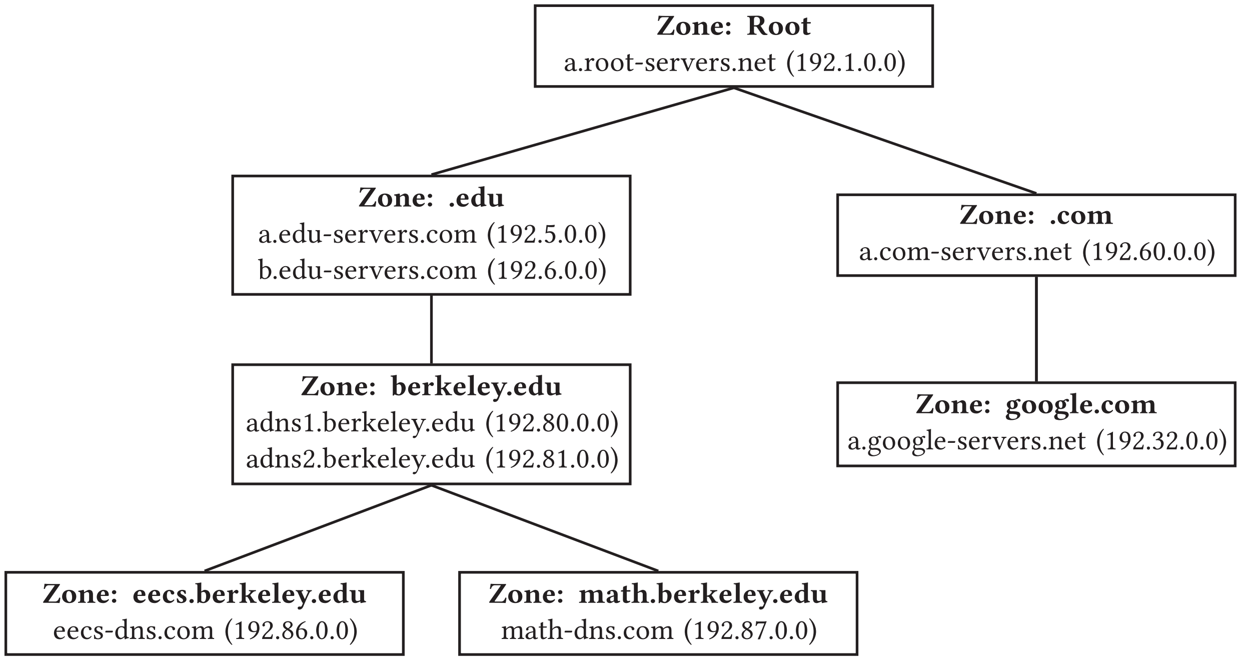

Host A and Host B both use the recursive resolver with IP address 8.8.8.8.

Q4.1 (1 point) What protocol helps Host A and Host B learn the IP address of the recursive resolver?

○ ARP ○ DHCP ○ NAT ○ TLS

For the next 3 subparts, assume all DNS caches start empty, and no packets are dropped.

Also, assume that the recursive resolver has hard-coded the domain and IP address of the root name server. (These hard-coded values do not count as a record in the cache.)

Q4.2 (2 points) Host A wants to learn the IP address of people.eecs.berkeley.edu.

In total, how many DNS packets are sent over the network to answer this query?

○ 0 ○ 2 ○ 6 ○ 8 ○ 10 ○ 12

Q4.3 (2 points) Continuing from the previous part (i.e. caches are not cleared): Host B wants to learn the IP address of people.math.berkeley.edu.

In total, how many DNS packets are sent over the network to answer this query?

○ 0 ○ 2 ○ 6 ○ 8 ○ 10 ○ 12

Q4.4 (2 points) For this subpart only, assume that name servers include information about *all* name servers in the next zone. For example, root will reply with information about both `.edu` name servers.

Also, assume that people.eecs.berkeley.edu and people.math.berkeley.edu each have one IP address.

Remember that it takes two records (an NS record and an A record) to provide information about a name server.

After both queries from earlier are performed, how many records are in the recursive resolver's cache?

○ 5 ○ 7 ○ 10 ○ 14 ○ 22

Q4.5 (1 point) Suppose that people.eecs.berkeley.edu has two IP addresses, an IPv4 address and an IPv6 address. Does this change your answer to Q4.4?

○ Yes, the cache has an additional A record.

○ Yes, the cache has an additional AAAA record.

○ No, because IPv6 is newer, so we can discard the IPv4 record.

○ No, because IPv4 is more widely-adopted, so we can discard the IPv6 record.

Q4.6 (3 points) Which of the following name servers, if directly queried, can give the A record with the IP address of classes.berkeley.edu? Select all that apply.

☐ a.root-servers.net ☐ adns1.berkeley.edu ☐ math-dns.com

☐ a.edu-servers.com ☐ adns2.berkeley.edu ☐ None of the above

☐ b.edu-servers.com ☐ eecs-dns.com

Q4.7 (1 point) There are two IP addresses (192.5.0.0 and 192.6.0.0) for the .edu zone's name server(s).

True or False: This means that the .edu zone **must not** be using anycast.

○ True ○ False

Q4.8 (1 point) There is only a single IP address (192.87.0.0) for the math.berkeley.edu zone's name server(s).

True or False: This means that the math.berkeley.edu zone **must** be using anycast.

○ True ○ False

Q5 *End-to-End* **(16 points)**

In this question, consider the network topology to the right.

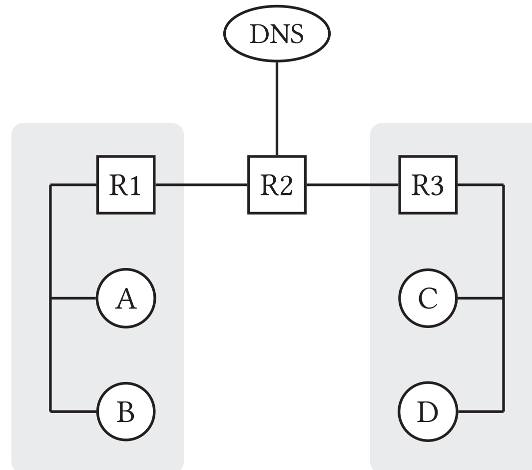

A, B, and R1 are all in the same subnet, and are connected on a single shared medium (also called a bus).

Likewise, C, D, and R3 are all in the same subnet (different from the above subnet), and are connected on a single shared medium.

There is one DNS resolver in the network, labeled DNS. Assume everyone uses the same DNS resolver.

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

Q5.1 (2 points) Host A joins the network for the first time, with all caches empty and no active connections.

Host A broadcasts a DHCP Discover request. R1 sends a DHCP Offer reply containing some IP addresses. Host A accepts this offer.

Which IP addresses were contained in R1's DHCP Offer? Select all that apply.

☐ R1's IP ☐ R2's IP ☐ A's IP ☐ B's IP ☐ DNS's IP

Q5.2 (2 points) Host A now wants to send an HTTP request to Host C.

Before any further steps, select all fields that Host A currently knows how to fill in.

☐ Source IP
☐ Destination IP
☐ Source MAC
☐ Destination MAC
☐ Source port
☐ Destination port
☐ None of the above

Q5.3 (2 points) From the previous subpart, there are some fields that Host A currently does not know how to fill in.

Select all additional protocols that Host A needs to run in order to fill in those missing fields.

☐ ARP ☐ DHCP ☐ DNS ☐ TLS ☐ None

Eventually, A is able to fill in all fields, and sends a packet to C. The packet arrives at C.

Q5.4 (2 points) Select all devices that can deduce (possibly by reading packets they are not supposed to be reading) that Host A's IP address and Host C's IP address are communicating.

☐ B ☐ R1 ☐ R2 ☐ R3 ☐ D ☐ None

Q5.5 (2 points) At this point, which of these IP-to-MAC mappings exists in Host A's cache? Assume no entries have expired, and nobody has sent an unsolicited ARP message. Select all that apply.

- ☐ C's IP to C's MAC
- ☐ DNS's IP to DNS's MAC
- ☐ R2's IP to R1's MAC
- ☐ C's IP to R1's MAC
- ☐ DNS's IP to R1's MAC
- ☐ B's IP to B's MAC

At this point, Host A leaves and re-joins the network, which means that all of Host A's caches have been cleared.

Q5.6 (1 point) After Host A re-joins the network, does Host A still have the same IP address?

- ○ Always
- ○ Sometimes
- ○ Never

For the rest of the question, assume that R1 is using NAT with Port Address Translation, as shown in lecture. Everyone in R1's subnet is using R1's public IP address, and A and B are assigned private IP addresses.

Q5.7 (1 point) Host A wants to communicate with Host C again.

At some point, Host A performs a DNS lookup for Host C's IP address.

Does R1 need to perform NAT translation on the DNS request packet?

- ○ Yes, because Host A's IP address is private.
- ○ Yes, because in NAT, the border router always modifies the source port for outgoing packets.
- ○ No, because DNS messages fit in a single UDP packet, so no long-term connection is formed.
- ○ No, because the DNS server's IP address is public.

Q5.8 (2 points) If Host A wants to send a packet to Host C, what additional steps does **Host A** need to do, now that NAT is enabled? Select all that apply.

- ☐ Choose a source port number that nobody else in the subnet is using.
- ☐ Choose a destination port number that nobody else in the subnet is using.
- ☐ Write R1's IP address in the source IP field.
- ☐ Write R1's IP address in the destination IP field.
- ☐ None of the above

Q5.9 (2 points) Now that NAT is enabled, select all devices that can deduce (possibly by reading packets they are not supposed to be reading) that Host A's IP address and Host C's IP address are communicating.

- ☐ B
- ☐ R1
- ☐ R2
- ☐ R3
- ☐ DNS
- ☐ None

Q6 *Datacenters* **(11 points)**

Q6.1 (2 points) What are some advantages of using a Fat-tree Clos topology (i.e. a topology with pods, edge/aggregation layers, and core switches) instead of a Fat-tree topology? Select all that apply.

- ☐ Links can be lower bandwidth.
- ☐ Bisection bandwidth is higher.
- ☐ Avoids single point of failure.
- ☐ Can be built out of commodity switches.
- ☐ None of the above

In the rest of this question, we'll extend Equal Cost Multi-Path (ECMP) routing so that a single elephant TCP flow can travel along two different paths in the network.

Q6.2 (1 point) Does the unmodified ECMP protocol from lecture allow packets in a single TCP flow to be split across two different paths?

- ◯ Yes, because we include destination IP as an input to the hash.
- ◯ Yes, because we include source IP as an input to the hash.
- ◯ No, because all packets in a single TCP flow all have the same 5-tuple.
- ◯ No, because all packets in a single TCP flow have the same source/destination IP (but not necessarily the same source/destination port).

Q6.3 (1 point) What is a disadvantage of splitting a single elephant TCP flow across different paths, compared to sending the entire flow along a single path?

- ◯ Packets that travel along different paths might experience different delays, and TCP performs poorly when packets arrive out-of-order.
- ◯ There is significant additional overhead in finding all routes to a destination. ECMP does have not this overhead.
- ◯ There is significant additional overhead in performing an additional TCP handshake to start the connection.
- ◯ There is significant additional overhead for routers to forward packets along two different paths. ECMP does not have this overhead.

To implement the splitting up of an elephant TCP flow onto two paths, we identify the two paths with unique IDs (e.g. "Red" and "Blue").

Q6.4 (1 point) What packet header field could be used to tag packets within a single TCP flow with different path IDs?

For this subpart only, assume we are not adding any additional headers to the TCP-over-IP packets.

- ◯ Flow label (IPv6)
- ◯ Protocol (IPv4)
- ◯ Next header (IPv6)
- ◯ Source port (TCP)

In the rest of the question, instead of using any existing headers to tag packets with a path ID, we add a new step to the encapsulation process:

1. The virtual machine (VM) creates a packet with overlay IP addresses in the header.
2. **New Step:** The virtual switch adds another header with the ID of the path that the packet should take. Within a TCP flow, half the packets are tagged with the "Red" header, and the other half are tagged with the "Blue" header.
3. The virtual switch adds another header with the underlay IP addresses.

The ECMP hashing process is extended so that in addition to the usual 5-tuple, routers also use the path ID to decide where to forward the packet.

Q6.5 (2 points) What headers are present in the packet as it travels through intermediate routers in the underlay network? Select all that apply.

- ☐ Headers with the overlay IP addresses.
- ☐ Headers with the underlay IP addresses.
- ☐ Headers with the path IDs.
- ☐ None of the above

Q6.6 (2 points) Consider modifying the ECMP hash so that it uses the 5-tuple and the path ID as input.

With this modified hash function, which sets of packets will always get forwarded to the same next-hop? Select all that apply. Each choice is independent.

- ☐ All packets in all TCP flows to the same destination IP.
- ☐ All packets in all TCP flows from the same source IP.
- ☐ All packets in all TCP flows between the same pair of hosts.
- ☐ All packets within a single TCP flow.
- ☐ None of the above

Suppose we want to send a TCP flow across a single path, but we want the ability to switch to a different path if we detect congestion.

For example, we start by assigning Path ID “Red” to all packets, but when we detect congestion, we switch to assigning Path ID “Blue” to all subsequent packets.

Q6.7 (2 points) What are some ways that the host can detect congestion? Select all that apply.

Assume the network uses unmodified commodity routers, e.g. ECN (Explicit Congestion Notification) is not in use.

- ☐ The host detects that the delay between sending a packet and receiving an ack is getting longer.
- ☐ The host receives many duplicate acks.
- ☐ The host detects that the queues at the routers are getting longer.
- ☐ None of the above

Q7 *Multicast* **(5 points)**

In this question, consider the network topology to the right.

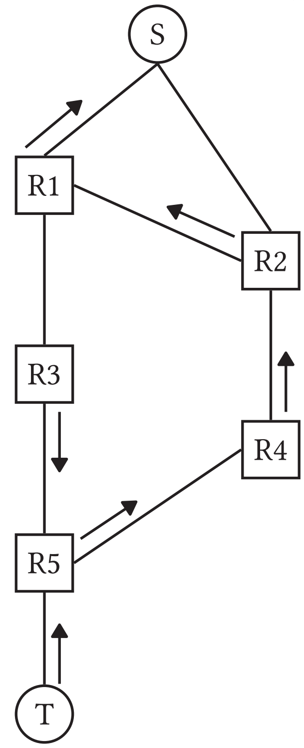

The link costs are omitted (i.e. not all links necessarily cost 1, but their costs don't matter in this question).

The arrows show the directed delivery tree for destination S.

S and T are part of a multicast group.

In the next three subparts, S wants to send a multicast message to this group, using DVMRP. You can assume no pruning takes place.

Q7.1 (1 point) S forwards the packet to R2.

R2 takes this packet and ______ (1), because the incoming link ______ (2) the next-hop to S.

- ◯ (1) Forwards it to R1 and R4, (2) is
- ◯ (1) Drops it, (2) is
- ◯ (1) Forwards it to R1 and R4, (2) is not
- ◯ (1) Drops it, (2) is not

Q7.2 (1 point) S also forwards the packet to R1.

R1 takes this packet and ______ (1), because the incoming link ______ (2) the next-hop to S.

- ◯ (1) Forwards it to R2 and R3, (2) is
- ◯ (1) Drops it, (2) is
- ◯ (1) Forwards it to R2 and R3, (2) is not
- ◯ (1) Drops it, (2) is not

Q7.3 (1 point) What path does the packet take from S to T?

- ◯ S, R1, R2, R4, R5, T
- ◯ S, R1, R3, R5, T
- ◯ S, R2, R4, R5, T
- ◯ S, R2, R1, R3, R5, T

Now, consider the same topology from before.

The arrows show the directed delivery tree for destination R4.

As before, S and T are part of a multicast group.

In the next two subparts, S wants to send a multicast message to this group, using CBT, with a core router of R4.

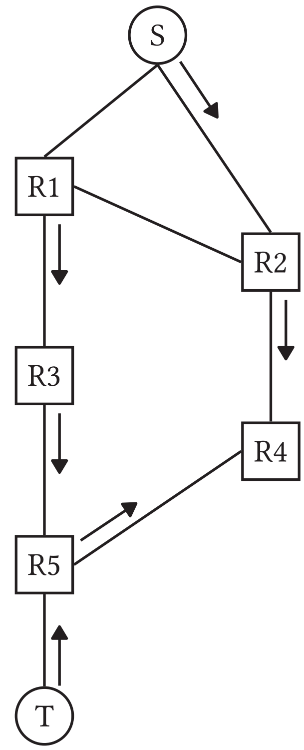

Q7.4 (1 point) First, S and T join the group by sending join messages to R4.

At this point, which routers are on the tree? Select all that apply.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5

Q7.5 (1 point) Once the tree is built, S can now send a multicast message to the group.

What path does the packet take from S to T?

○ S, R1, R2, R4, R5, T

○ S, R1, R3, R5, T

○ S, R2, R4, R5, T

○ S, R2, R1, R3, R5, T

Q8 *Wireless* **(7 points)**

Consider four hosts communicating over a single frequency channel, using the MACA protocol from class (not MACAW, i.e. no ack, no DS, no RRTS). Each subpart is independent.

In the diagrams below, when a message is transmitted, the line segments show who receives the transmission.

Q8.1 (1 point) A transmits a Request to Send (RTS) packet, and the RTS is received by B.

Then, B transmits a Clear to Send (CTS) packet, and the CTS is received by A and C.

No other packets are transmitted.

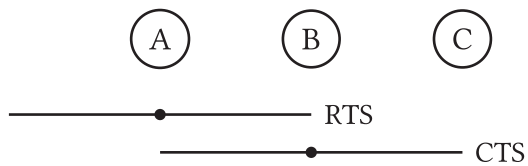

At this point, who is allowed to transmit data packets?

○ A ○ B ○ C ○ D

Q8.2 (2 points) Desired communications:

- A wants to talk to B.
- C wants to talk to B.

Suppose A and C each simultaneously transmit an RTS packet with equal signal strength, and these transmissions arrive at B at the same time.

No other packets are transmitted.

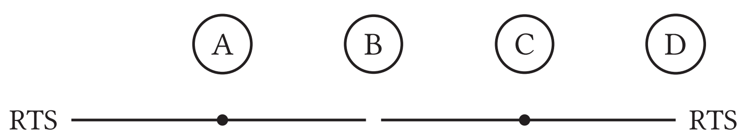

At this point, what will happen next at B?

○ B picks the RTS from the host who wants to transmit the least amount of data, and transmits a CTS packet for that host's RTS.

○ B receives both packets at the same time, causing a collision.

○ B randomly picks one of the two hosts, and transmits a CTS packet for that host's RTS.

○ B sends a CTS packet for both hosts' RTS packets, allowing both A and C to transmit to B.

For the next three subparts, consider this scenario:

Desired communications:

- C wants to talk to A.
- B wants to talk to D.

C transmits an RTS packet, and A, B, and D hear it. Then, A transmits a CTS packet, and B and C hear it. No other packets are transmitted.

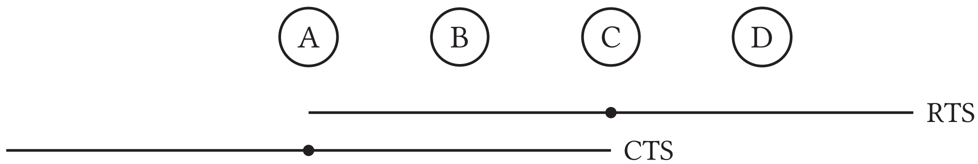

Q8.3 (1 point) At this point, what will happen next at B?

- ◯ B immediately transmits its data packets, without sending an RTS packet first.
- ◯ B immediately transmits an RTS packet.
- ◯ B waits for A and C to finish, and then attempts to transmit an RTS packet.
- ◯ B never sends an RTS packet, and gives up on transmitting any data.

Q8.4 (2 points) At this point, can A and C both transmit data packets to each other simultaneously (with no further RTS/CTS exchanges)?

- ◯ Yes, because A transmits a CTS, so everyone in A's range must be quiet. This guarantees that A's data transmission will be received by C.
- ◯ Yes, because in MACA, after a single RTS/CTS exchange, the two hosts can take turns transmitting data.
- ◯ No, but if C also transmitted a CTS at this point, then A and C can both transmit data simultaneously.
- ◯ No, because if A and C both transmit data packets to each other at this point, this could cause collisions.

Q8.5 (1 point) For this subpart only, suppose the CTS transmitted by A is dropped, and C does not receive the CTS.

At this point, what will happen next, and what happens to C's contention window (CW)?

- ◯ C increases its CW, and tries again by transmitting another RTS later.
- ◯ C decreases its CW, and tries again by transmitting another RTS later.
- ◯ C increases its CW, and starts transmitting its data packets.
- ◯ C decreases its CW, and starts transmitting its data packets.

Q9 *Cellular* (5 points)

Recall the steps of using a cellular network, as seen in lecture. For each subpart, select the step in which the event occurs. Not all choices may be used, and some choices may be used more than once. Each subpart is independent.

If the event never occurs in the process of using a cellular network, select None of the above.

Q9.1 (1 point) The user receives an IMSI (subscriber ID).

- ○ Step 0: Registration
- ○ Step 1: Discovery
- ○ Step 2: Attachment
- ○ Step 3: Data Exchange
- ○ Step 4: Handover
- ○ None of the above

Q9.2 (1 point) The routers in the cellular core use a routing algorithm (e.g. distance-vector) to find the shortest path to the user.

- ○ Step 0: Registration
- ○ Step 1: Discovery
- ○ Step 2: Attachment
- ○ Step 3: Data Exchange
- ○ Step 4: Handover
- ○ None of the above

Q9.3 (1 point) Suppose an attacker steals a victim user's IMSI. The attacker does not know any other information about that victim user.

The attacker tries to connect to the cellular network by presenting the victim's IMSI to the mobility manager. The mobility manager detects that the connection is invalid at this step.

- ○ Step 0: Registration
- ○ Step 1: Discovery
- ○ Step 2: Attachment
- ○ Step 3: Data Exchange
- ○ Step 4: Handover
- ○ None of the above

Q9.4 (1 point) Consider a user who is already connected to the cellular network. The user is moving in and out of the range of various towers.

The user's device listens for nearby towers. The user's device (not the cellular network) picks the next tower to use during this step.

- ○ Step 0: Registration
- ○ Step 1: Discovery
- ○ Step 2: Attachment
- ○ Step 3: Data Exchange
- ○ Step 4: Handover
- ○ None of the above

Q9.5 (1 point) The mobility manager tells the cellular core routers to delete an existing tunneled path, and install a new tunneled path.

- ○ Step 0: Registration
- ○ Step 1: Discovery
- ○ Step 2: Attachment
- ○ Step 3: Data Exchange
- ○ Step 4: Handover
- ○ None of the above

This page intentionally left (mostly) blank.

The exam continues on the next page.

Q10 *Midterm Clobber: STP* **(12 points)**

This question will only be used to potentially replace your score on the midterm STP question. It will not affect your final exam score.

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.

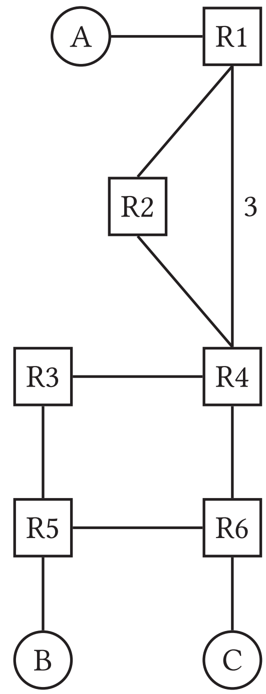

Q10.1 (2 points) After running STP, how many links get disabled?

- ○ 1
- ○ 2
- ○ 3
- ○ 4
- ○ 5
- ○ 6

Q10.2 (2 points) Select all routers that disable at least one of their links.

- ☐ R1
- ☐ R2
- ☐ R3
- ☐ R4
- ☐ R5
- ☐ R6

Q10.3 (1 point) During STP, the **first** message that R4 announces to its neighbors is:

"The root is _____, and I can reach the root with cost _____."

- ○ R1; 0
- ○ R1; 3
- ○ R4; 0
- ○ R1; 2
- ○ R3; 1
- ○ A; 3

Q10.4 (1 point) After STP converges, R4 thinks that the root is _____, and that R4 can reach the root with cost _____.

- ○ R1; 0
- ○ R1; 3
- ○ R4; 0
- ○ R1; 2
- ○ R3; 1
- ○ A; 3

Suppose STP has converged. Regardless of your answers to the previous subparts, assume that the R3-to-R4 and R1-to-R4 links are disabled:

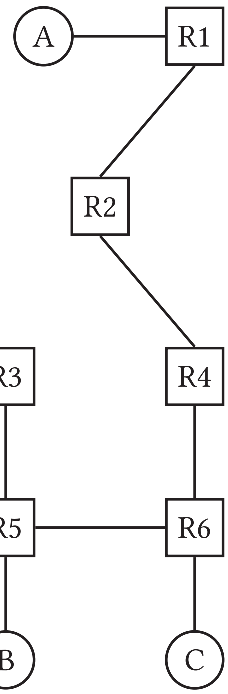

Switches R1 to R6 are all learning switches.

The forwarding tables start out empty, except for R5's forwarding table, which starts with one hard-coded entry:

R5's Forwarding Table	
Destination	Next Hop
A	R6

In each of the next 3 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.

Q10.5 (2 points) C sends a packet to A.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6

Q10.6 (2 points) A sends a packet to C.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6

Q10.7 (2 points) B sends a packet to A.

☐ R1 ☐ R2 ☐ R3 ☐ R4 ☐ R5 ☐ R6

Nothing on this page will affect your grade.

Comment Box

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

CS 168's new course mascot is an octopus! Suggest any designs, names, backstories, etc. below:

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.