

This version of the exam was released in Dec 2024 as a practice exam for students. It only contains a selected subset of questions from the original SP24 final exam. The questions and solutions may contain bugs or errors.

PRINT your name: \_\_\_\_\_, \_\_\_\_\_  
(last) (first)

PRINT your student ID: \_\_\_\_\_

---

- You have 170 minutes.
- If you have a question during the exam, raise your hand and a proctor will come over to you. We will not answer your question on the spot, but we will make a clarification to everyone if we believe it's necessary.
- If you need to use the restroom, please bring your exam, answer sheet, phone, and ID to the front. Only one student is allowed to use the restroom at a given time.
- If you finish early, please bring your exam and answer sheet up to the front. **You may not leave early once there are 15 minutes remaining in the exam.**

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

- Unselected option (completely unfilled)
- Only one selected option (completely filled)

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

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

**Q1** *Honor Code*

**(0 points)**

Read the following honor code 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: \_\_\_\_\_

**Q2 True / False**

**(13 points)**

Q2.1 (1 points) TRUE or FALSE: In the Internet architecture, we think of transport as layer 3.

- TRUE  FALSE

Q2.2 (1 points) TRUE or FALSE: The TCP header is the outermost header on packets that use TCP.

- TRUE  FALSE

Q2.3 (1 points) TRUE or FALSE: Routers never fragment IPv6 packets.

- TRUE  FALSE

Q2.4 (1 points) TRUE or FALSE: AIAD (Additive Increase Additive Decrease) is “unfair” in that it could lead to different average bandwidth allocations to flows with the same RTT.

- TRUE  FALSE

Q2.5 (1 points) TRUE or FALSE: UDP never uses a checksum.

- TRUE  FALSE

Q2.6 (1 points) TRUE or FALSE: DNS typically runs over UDP.

- TRUE  FALSE

Q2.7 (1 points) TRUE or FALSE: DHCP typically runs over UDP.

- TRUE  FALSE

Q2.8 (1 points) TRUE or FALSE: ARP typically runs over UDP.

- TRUE  FALSE

Q2.9 (1 points) TRUE or FALSE: ARP requests are typically broadcast.

- TRUE  FALSE

Q2.10 (1 points) TRUE or FALSE: ARP responses are typically broadcast.

- TRUE  FALSE

Q2.11 (1 points) TRUE or FALSE: TCP packets with the ACK flag set cannot carry payloads.

- TRUE  FALSE

Q2.12 (1 points) TRUE or FALSE: In DNS, the MX record is the name of the authoritative server for that domain.

- TRUE  FALSE

Q2.13 (1 points) TRUE or FALSE: In a typical TCP connection, all packets aside from the initial SYN have the ACK flag set.

- TRUE  FALSE

**Q3 MCQ****(8 points)**

Q3.1 (2 points) Propagation delay

- Tracks the end-to-end transmission time of the entire packet.
- Dominates total latency for extremely large packets.
- Depends on transmission delay.
- Is proportional to the physical distance between two hosts.
- None of the above

Q3.2 (1 points) Which is the faster way of delivering a 1 TB file from Berkeley to Stanford (assuming the distance between these two locations for both methods is 45 miles): (Pick one.)

- Sending it over a 1Gbps link.
- Driving at an average speed of 25mph (due to heavy traffic on 101) carrying a 1 TB disk.

Q3.3 (1 points) Which of the following is provided by both UDP and TCP? (List all that apply.)

- Reliable transfer
- Mux and demux from/to application processes
- Congestion control
- A byte-stream abstraction

Q3.4 (1 points) When a host wants to obtain the IP address of www.ebay.com, where does it send its initial DNS request?

- The DNS root
- The authoritative server for ebay.com
- The .com TLD server
- The host's local DNS server

Q3.5 (1 points) Which of the following statements are in agreement with our discussion of the end-to-end principle.

- It is acceptable to implement a function in the network if we cannot trust hosts to implement the function correctly.
- A multi-hop wireless network should not implement per-link retransmissions.
- Switches should not be implementing Explicit Congestion Notification
- Having an SDN controller compute routes is a violation of the end-to-end principle

Q3.6 (1 points) To find network paths, which of the following does Traceroute leverage?

- IGP                       IGMP                       ICMP                       RCMP

Q3.7 (1 points) Which of the following is true of SDN?

- Openflow is an implementation of SDN
- SDN centralizes management of network state
- BGP programming routes into routers in an example of SDN
- SDN allows for arbitrary computations over network state that can be programmed into switches that expose the appropriate APIs

#### Q4 Network Virtualization

(7 points)

Assume we are a cloud provider that uses network virtualization to enable communication between customer VMs.

Q4.1 (1 points) What is the abstraction that network virtualization provides a customer?

- A network of VMs in a Clos topology
- A network of VMs all routable to one another but not reachable from other VMs
- A server hosting their VMs in a datacenter network
- A network of VMs all routable to one another and other hosts

Q4.2 (1 points) We have a customer Pepsi with the address space 10.0.0.0/16, and a customer Coke with the address space 10.1.0.0/16. A new customer, Dr. Pepper, wants a virtual network with address space 10.0.128.0/17. Can we give customer Dr. Pepper the address space it is requesting?

- Yes, because network virtualization configures a separate virtual network for each customer
- No, because this address space overlaps with another customer's
- No, because these addresses will not be routable from the network
- Yes, but it cannot be hosted in the same servers as the other customers

Q4.3 (1 points) Will hosts in Coke's network be able to initiate connections with other hosts on the Internet with their given IP addresses? Why?

- Yes
- No

Q4.4 (1 points) Assume that Pepsi provisions VM A with address 10.0.0.2 and Coke provisions VM B with address 10.1.0.2 on the same underlying server. What will the SDN controller program into the server that hosts VM A and VM B?

- Match action rule mapping the virtual network ID to the VM
- Match action rules mapping the virtual network ID to the IP address
- Route that matches the prefix of the customer address space to the VM

Q4.5 (1 points) Assume that Pepsi provisions VM C with address 10.0.128.2 and Dr. Pepper provisions VM D with address 10.0.128.2. Could VM C and VM D be placed on the same server?

- No, because the network will not be able to tell the VMs apart.
- No, because the hypervisor cannot host two VMs with the same IP address.
- Yes, because the packets will be delivered to the correct server either way.
- Yes, because the SDN controller can program a way to disambiguate them.

Q4.6 (1 points) As these VMs are provisioned and deprovisioned, does anything in the routers in the underlay network need to change?

- Yes
- No

Q4.7 (1 points) Consider a packet traversing a link between two routers in the datacenter network in a path between two Coke VMs. What headers must be present at this point to deliver the packet correctly?

- Ethernet header (source = source server MAC, destination = destination server MAC)
- IP Header (source = source VM's IP and destination = destination VM's IP)
- Virtual Network Header
- IP Header (source = source server's IP and destination = destination server's IP)
- Traffic Engineering Header
- HTTP Header

## Q5 Host Networking

(7 points)

In this question, we consider a host in a datacenter.

### Transport

Q5.1 (1 points) First, we consider whether to use an RDMA or TCP stack. Why might we choose RDMA?

- To reduce implementation complexity
- To save CPU cycles
- For faster design iteration
- For a message-passing interface
- For better performance

Q5.2 (1 points) Suppose we choose TCP, but we are considering making modifications. We want our congestion control solution to avoid filling buffers in the network. We thus consider different signals of congestion and how they relate to queueing delay.

a) Loss: Is loss a good choice given our goal?

- Yes
- No

b) ECN: Assume that the buffer at our bottleneck can hold one BDP of packets and that the ECN marking threshold is  $1/2$  BDP (ie, packets will be marked as having experienced congestion if they leave the queue when the queue is longer than the threshold). Is it possible for a single sender to cause a packet to get marked in the first BDP of packets sent?

- Yes
- No

c) Delay: Assume that all buffers in the network are 1MB and all link rates are 8Mbps. Assume all packets are 1000B. Consider a case where we see an increase of 100us in the measured delay. Could this extra delay be from queueing at multiple hops, just one, or both?

- One hop
- Multiple

d) Receive rate: Assume the receive rate for packets in our flow is lower than the rate they were sent? Should we expect an increase or decrease in RTT?

- Increase
- Decrease

Q5.3 (1 points) Now, consider a scenario where the measured RTT increases between two consecutive packets received by a socket in our host. (Assume that the flow is between our host and another host in the same datacenter.)

- a) Can we estimate what the range of possible RTT values will be before sending any packet? Why?

- b) Can we reasonably assume that this increase due to congestion?

Yes  No

Q5.4 (1 points) Inspired by delay-based congestion control in datacenters, a colleague, Alice, is proposing to do the same in the WAN between hosts in your network and anywhere else on the Internet. Why might you argue this is a bad idea?

- Because RTTs in the WAN may vary widely for reasons other than congestion
- You cannot get precise timestamps from an arbitrary NIC somewhere on the Internet
- Delay only matters in a datacenter environment
- Because in the WAN, loss-based congestion control is effective in avoiding queuing delay
- Because loss-based congestion control provides better fairness than delay-based congestion control

### Rate enforcement

Q5.5 (1 points) Assume we want to ensure our host never uses more than 5Gbps on the link between the host and the datacenter network, regardless of the number of competing flows. Can congestion control enforce this for us?

Yes  No

Q5.6 (1 points) Assume we want to ensure that our flow always gets at least 1/4 th of a link for any link in our network. Can we enforce this? How?



Q5.7 (1 points) Suppose we decide to use a timing wheel to pace traffic leaving our host. Assume there are two classes of traffic, A and B. Assume the following about these classes:

- A's last packet was sent at  $t = 10\mu\text{s}$
- A's rate limit is 1 Gbps, which is equivalent to  $12\mu\text{s}$  between packets
- B's last packet was sent as  $t = 6\mu\text{s}$
- B's rate limit is 0.8 Gbps, which is equivalent to  $15\mu\text{s}$  between packets

Assume the following about the timing wheel:

- The time horizon is 5s
- The time granularity is  $10\mu\text{s}$
- A given traffic class can have at most one packet per queue
- All queues are currently empty
- The now pointer last moved at  $t = 15\mu\text{s}$

Lastly, assume all packets are 1500B

a) At time  $t=16\mu\text{s}$ , a packet from A is ready to be scheduled. What will its timestamp be?

b) At time  $t=17\mu\text{s}$ , a packet from B is ready to be scheduled. What will its timestamp be?

c) Which traffic class sends its packet next?

A

B


d) What is the maximum rate supported by this timing wheel? You can answer with an expression and do not have to calculate your answer.

e) What is the minimum rate enforceable by this timing wheel? You can answer with an expression and do not have to calculate your answer.

**Q6** *Datacenters*

**(4 points)**

Suppose we had the following datacenter topology, where each link has a capacity of 20 Gbps.



Q6.1 (1 points) Provide the following information about this network:

a) Bisection bandwidth

b) Full bisection bandwidth

c) Oversubscription ratio

Q6.2 (1 points) Suppose ECMP is used to handle multi-path routing and TCP makes optimal, fair use of the bandwidth.

- a) Assume hosts H1, H2, and H3 are sending flows to H8, and H5 is sending a flow to H6. Assuming the hashing function places these flows in the best possible way, what is the approximate rate each host sends at?

H1:

H2:

H3:

H5:

- b) Do the above answers change if the hashing function places these flows in the worst possible way?

- c) What could we do to mitigate congestion when flows are placed non-optimally without causing other disruptions?

- Change the source IP address to force a repath
- Temporarily stop the flow to force a repath
- Use congestion control
- Change the flow label to force a repath

Q6.3 (1 points) Suppose we want to modify this topology so that all hosts can communicate at 20 Gbps. We propose adding links from R5 to R4 and R6 to R4. Does this achieve our goal? Why?

Yes


No

Q6.4 (1 points) Ignoring oversubscription, what major problem do you see with this network?

### Q7 *Shakey Signals*

(8 points)

You are part of the development team for a new wireless earthquake alert system on Berkeley's campus. Nodes for this system are distributed throughout the campus and are designed to detect unusual seismic activity. The nodes transmit seismic data to receivers using a wireless technology. You have been tasked with evaluating and fixing bugs in the medium access control protocol.



Sensor nodes are located at LBNL (a), International House (b), Hearst Gym (c), and Downtown Berkeley BART (d). Receivers are located in Cory Hall (a), Clark Kerr (b), and Valley Life Sciences (c). Transmission ranges of the different nodes are shown by circles around the nodes. For this problem you should assume that the transmission range is the same in all directions as indicated by the circles. Additionally, the transmission range of receivers is larger than the entire campus, so you can assume that transmissions from a receiver will reach all the nodes. Nodes do not have a shared clock and hence we cannot assume time synchronization.

#### Part 1

Your boss would like you to implement a protocol with the following characteristics.

- All nodes share a channel which they transmit on.
- All receivers have their own unique channel which they can transmit on, which is different from the channel that sensor nodes transmit on.
- Each node listens to see if the channel is clear and if it is clear, the node broadcasts data packets

for 10 seconds.

- If the channel is not clear, the node waits one minute, and then tries again.

Q7.1 (1 points) They claim that, using this protocol, Node (a) located at LBNL can always transmit to Receiver (a) without worrying about collisions. Is this true?

- Yes  No

Q7.2 (1 points) Identify the receivers if any which are, in this protocol, impacted by the hidden terminal problem.

- A  B  C

Q7.3 (1 points) Identify the receivers if any which are, in this protocol, impacted by the exposed terminal problem.

- A  B  C

## Part 2

You notice that the performance is not as good as expected, so you decide to design a new protocol which does not use carrier sense. In your new protocol:

- All nodes share a channel which they transmit on.
- All receivers have their own unique channel which they can transmit on, which is different from the channel that sensor nodes transmit on.
- **New:** Each node sends a request to send (RTS), and if the channel is clear the receiver sends back a (CTS), then upon receiving a CTS the node broadcasts data packets for 10 seconds.
- **New:** If a node does not receive a CTS, it backs off and retries later.

Q7.4 (1 points) You claim to your boss that your new RTS/CTS should prevent collisions of data packets at all three receivers. Is this true?

- Yes  No

Q7.5 (1 points) Identify the receivers if any which are, in this protocol, impacted by the hidden terminal problem.

- A  B  C  None

Q7.6 (1 points) Identify the receivers if any which are, in this protocol, impacted by the exposed terminal problem. If none, please select "None".

- A  B  C  None

Q7.7 (1 points) If the node does not receive a CTS, it backs off (since there must have been a collision or other failure) and then retries. The backoff is randomly drawn from a range of 0 to the contention window (CW) where:

- Initial Condition:  $CW_0 = 2$
- Maximum Case:  $CW_{max} = 64$
- Upon successful RTS/CTS:  $CW = CW_0$
- Upon failed RTS/CTS:  $CW = \min(2CW, CW_{max})$

Is this fair?

Yes

No

Q7.8 (1 points) Since you took CS168 and know the data is important, you are concerned about reliability. You decide that receivers should send back an ACK if the packet is received. You change your backoff code as follows:

- Initial Condition:  $CW_0 = 2$
- Maximum Case:  $CW_{max} = 64$
- Upon successful RTS/CTS:  $CW = CW - 1$
- Upon successful RTS/CTS, but failed ACK, no change
- Upon failed RTS/CTS:  $CW = \min(2CW, CW_{max})$

Is this fair?

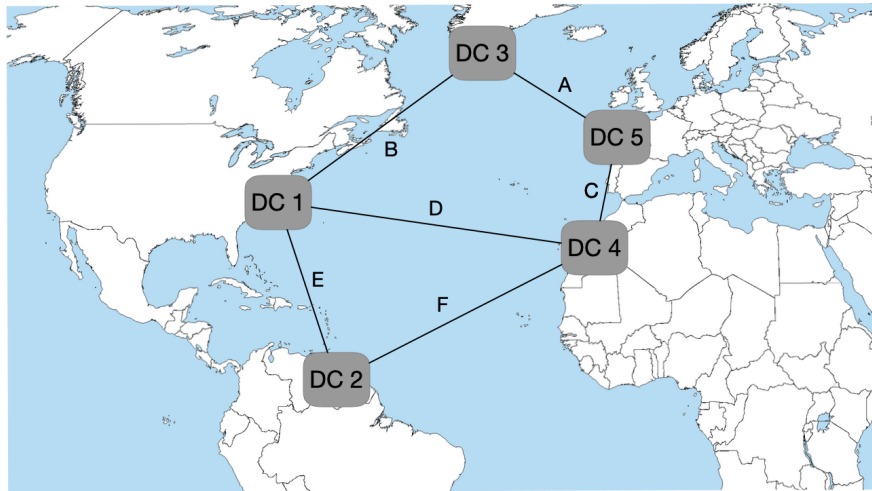
Yes

No

**Q8 Traffic Engineering**

**(6 points)**

Consider the following Wide Area Network (WAN) spanning multiple continents connecting datacenters operated by a single operator. You have been tasked with designing a policy to route packets between these data centers. The figure shows the border routers connecting these datacenters. You can assume the length of the links in the diagram are proportional to their physical distance.



Assume the following about each link:

- Link A: cost 10, 5 Tbps
- Link B: cost 10, 5 Tbps
- Link C: cost 10, 5 Tbps
- Link D: cost 25, 10 Tbps
- Link E: cost 10, 5 Tbps
- Link F: cost 10, 5 Tbps

As a first go, you implement least cost routing based on the above link costs.

Q8.1 (1 points) With no traffic engineering, assuming least cost routing, what path would a 2 Tbps constant workload take from DC1 to DC4?

Q8.2 (1 points) This works well for the workload in the previous part, but then the traffic increases to the following listed workloads. With least cost routing, is there congestion? Assume that the 2 Tbps workload from the previous part is no longer present.

- Workload 1: 6 Tbps from DC1 to DC4
- Workload 2: 3 Tbps from DC1 to DC5

Yes

No



Q8.3 (1 points) What route(s) should the workloads from the previous part take to use the least-cost-with-capacity route and avoid congestion? Assume that workloads cannot be split across multiple paths.

a) Workload 1

b) Workload 2

Q8.4 (1 points) Assume that we enforce these paths by programming forwarding rules based on a TE header into our WAN routers. For a packet from workload 1, when is the latest this header should be added?

- In the source host
- In a datacenter switch in DC 1
- In the WAN router connecting DC 1 to the other datacenters
- On the way to the next datacenter in the path

Q8.5 (1 points) Assume that a new workload begins between DC 1 and DC 4, Workload 3. Workload 3 is a streaming application which cares more about latency than every packet getting through consistently. Though, it cannot tolerate less than 0.5 Tbps. With our current TE algorithm, would this application's workload be placed on the best link for its needs?

- Yes
- No

Q8.6 (1 points) In order to get even better balancing of load across our links, we consider a design which could split workloads across paths. The implementation we are considering would assign packets within a workload to paths in a round-robin way as the packets arrive.

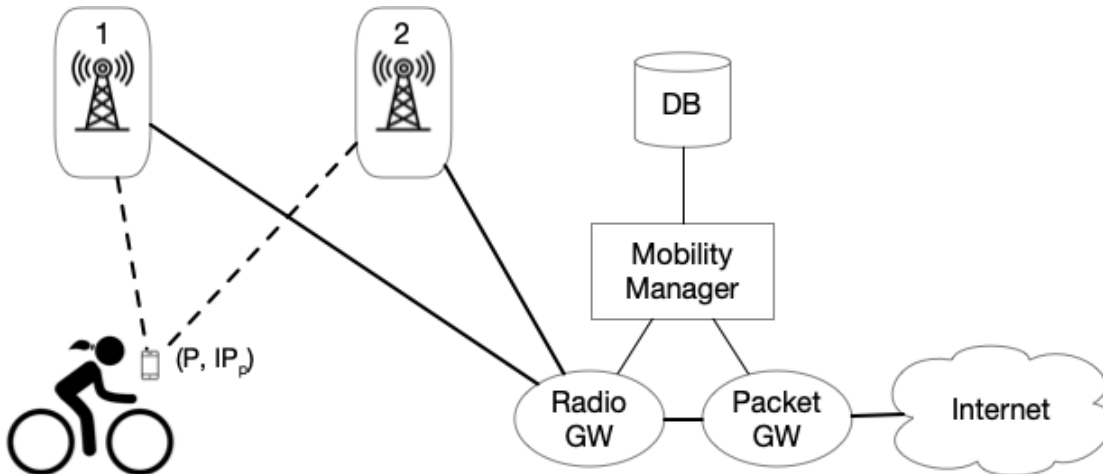
a) What problem could this cause?

b) How else could we split the traffic to avoid this problem?

**Q9 High Speed Handover**

**(6 points)**

Janet is biking to campus when she realizes she is late to the Friday morning CS168 staff meeting! Luckily, she can join the Zoom call using her cell phone. Janet's phone (noted as P with IP address  $IP_p$ ) is turned off but she quickly turns it on and gets ready to join the Zoom call. Janet is a customer of AT&T and you can assume that AT&T has excellent coverage in the area she is currently speeding through. The above diagram shows Janet biking within range of multiple towers along with the cellular core.



Q9.1 (1 points) What is the one top level central design goal for cellular?

- Reliability
- Distributed management
- Mobility
- Accountability

Q9.2 (1 points) Assuming signal strength decreases with distance, which tower should Janet's phone connect to?

- AT&T Tower 1
- AT&T Tower 2

Q9.3 (1 points) What entity in the above diagram determines whether to accept the attachment request?

- Radio Gateway
- Mobility Manager
- AT&T Tower 1
- Packet Gateway
- AT&T Tower 2

Q9.4 (1 points) From the events listed in the following list, select the events that will occur as Janet calls into and participates in the meeting and organize them in the order of their occurrence in the provided Table. Note that only 10 of the events listed will occur, and not all events that occur are listed.


Remember that P denotes Janet's phone and  $IP_p$  denotes Janet's phone's IP address.

- A. The mobility manager looks up P's IMSI in AT&T's cellular database
- B. P hears a beacon advertised by an AT&T Tower 1
- C. P transmits discovery beacons
- D. AT&T's mobility manager authenticates P using a shared key associated with P's IMSI
- E. AT&T's mobility manager authenticates P using a shared key associated with  $IP_p$
- F. P scans for nearby AT&T towers
- G. P discovers the channel on which it must communicate via AT&T Tower 1's beacon messages
- H. AT&T's radio gateway authenticates P using a shared key associated with P's IMSI
- I. P transmits its first IP data packet
- J. P completes a 3-way handshake to a zoom server (assume that Zoom is a TCP-based app)
- K. The mobility manager configures two tunnels, one from AT&T Tower 1 to radio gateway and another from radio gateway to packet gateway
- L. P sends an attachment request
- M. AT&T's mobility manager assigns an IP address to P
- N. The mobility manager configures a single tunnel, one from AT&T Tower 1 to packet gateway
- O. Janet discusses the cellular question on the final exam with the other GSIs on Zoom

Order	Event
1	
2	
3	
4	
5	
6	
7	
8	
9	
10	

Q9.5 (1 points) The handover is a collaborative process between many entities! Match the following entities to their correct letter (a, b, c, d, e) according to the following handover diagram.

- i) User device



- (a)
- (b)
- (c)
- (d)
- (e)

ii) Mobility manager

- (a)
- (b)
- (c)
- (d)
- (e)

iii) Radio gateway

- (a)
- (b)
- (c)
- (d)
- (e)

iv) AT&T Tower 1

- (a)
- (b)
- (c)
- (d)
- (e)

v) AT&T Tower 2

- (a)
- (b)
- (c)
- (d)
- (e)


Q9.6 (1 points) The commute to campus is never ending and Janet is now biking on the Bay Trail where AT&T has no coverage. Fortunately, Verizon does and AT&T has a roaming agreement with Verizon. Which of the following events will occur:

- AT&T sends a handover request to Verizon's Tower
- Verizon's mobility manager will authenticate P by looking up P's shared key in Verizon's database
- Verizon's mobility manager will set up tunnels for P from Verizon's Tower to Verizon's packet gateway
- AT&T updates the location for P in its local cellular database
- P's TCP connection to the zoom server breaks


**Q10 Routing**

**(12 points)**

Q10.1 (4 points) **Link-State:** For the network shown below, consider the case where all nodes have flooded their link-state information to everyone, shortest paths have been calculated, and the appropriate forwarding entries installed. But then the link between E and F goes down. The packets announcing this link-status change reach all nodes **except** B, and again routes are recalculated. When D sends a packet to F, what path does it take? On the answer sheet, list all the nodes traversed, including D and F. You can break ties between shortest paths arbitrarily.



Q10.2 (4 points) **Distance-Vector:** Distance-vector routing is plagued by the “count-to-infinity” problem. Poisoned reverse is one approach to mitigating this problem. For the following graphs, consider the case where distance-vector routing has converged to a stable set of routes, and then the destination (marked D) fails, taking down all of the links attached to it. Which of these graphs still have the count-to-infinity problem even when poisoned reverse is used? List all that apply.



Q10.3 (4 points) **A Distance-Vector Routing Mystery:** Consider a router A which is connected to another router B. Assume we are using a distance-vector routing protocol which uses poisoned reverse (for partial updates) and split horizon (for full updates), where all distances are in hop-counts (and all links count as a single hop). We focus on their routing entries for reaching a specific destination X. After the distance-vector routing algorithm has reached equilibrium (i.e., no information is changing), in A’s most recent full update to B (about X), A declared a distance of 4 hops. In B’s most recent full update to A (about X), B declared a distance  $d$ .

i. What is the smallest possible integer value for  $d$ ?

ii. What is the largest possible integer value for  $d$ ?

*Nothing on this page will affect your grade in any way.*

## Doodle

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

