

Exam - Midterm

Due: 18 March 2022, 11:59pm (No late days)

Exam Code of Conduct

- This is a “take-home”-style exam. You can work on the exam on your own time until the assigned deadline.
- You can work on the exam by printing this document, working on your own paper, or typing your responses. Either way, your work must be legible to be graded.
- To turn in your work, submit a PDF or a scan/photo of your exam to the assignment labeled “Midterm” on Gradescope.
- This exam is open notes. You can use any notes you have taken during class, class resources, or other online resources. However, *the solutions you write down must be your own work*. Copying answers verbatim from online sources is not permitted, nor is it likely to help you—often, searching online for an answer to a conceptual question will give you a generic answer, but it is up to you to put it in context for the problem!
- This exam is an individual assignment: **collaboration with others is not permitted**, other than asking the course staff questions. Since this is a take-home exam, you are on your honor to adhere to this requirement. *Violations of this policy are considered a violation of the collaboration policy and will be handled accordingly.*
- You are encouraged to ask the course staff questions about the exam. You may do so during office hours (check the class calendar for details), or by posting on EdStem. **When posting on EdStem, please make all exam-related questions private**—we will decide on which questions can be made public.

By taking this exam, you agree to abide by this code of conduct.

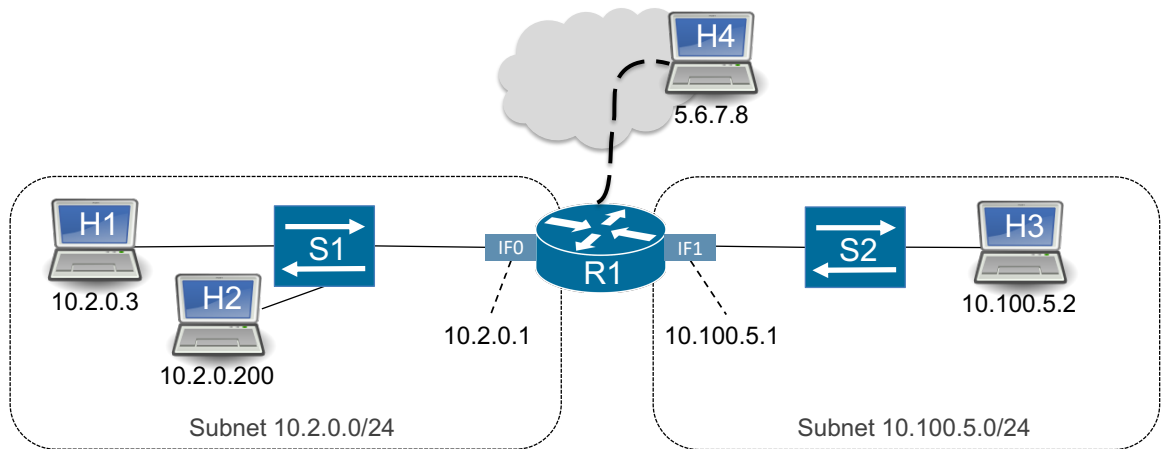
When you are done, submit your work on Gradescope to the assignment titled “Midterm.” Your completed exam is due: Friday, 18 March 2022 by 11:59pm EDT.

1. IP and Ethernet Addressing [20 pts]

- a. Sending frames on an Ethernet network is very convenient, as you don't have to configure anything. How do nodes get unique Ethernet addresses? [4 pts]

- b. Ethernet hosts use the ARP protocol to discover the Ethernet address of other nodes on the network. Consider the network pictured below, which has two private subnets connected to router R1 via separate switches S1 and S2.

Suppose H1 wants to send a packet to H4. H1 will send an ARP request of the form "Who has <IP>?", where <IP> is some IP address. **Which IP is being requested by the ARP request, and which nodes (out of H1–H4, and R1) will receive it?** (Assume all hosts start with an empty ARP cache.) [4 pts]



- c. While Ethernet is very convenient, why doesn't it scale to very large networks? (Write one or two sentences describing at least one challenge to making an Ethernet network scale to, eg., thousands of hosts.) [4 pts]
- d. In contrast to Ethernet addresses, hosts need a new IP address configuration (either manually, or using a protocol like DHCP) each time it joins a new network. Why is this necessary? (In other words, why does my laptop need a different IP on my home network vs. when I bring it to the CIT?) [4 pts]
- e. Why does IP routing, unlike a link-layer protocol like Ethernet, scale to the entire Internet? (Again, your answer need only be 1–2 sentences.) [4 pts]

2. FTP vs. NAT Network Address Translation (NAT) allows multiple clients in a private network to share a single public IP address (on the NAT gateway).

Suppose a router performing NAT with public address 105.10.2.5 is operating on behalf of the nodes in your home, which are part of the internal network 192.168.4.0/24. [20 pts]

- a. When a client with address 192.168.4.2 sends a TCP packet to an outside server at 201.20.4.7 (as shown in the table below), **how does NAT modify the packet?**

To express your answer, fill in the table with the translated version of the packet header and explain your reasoning below. When writing the packet fields, be as specific as you can—if there are multiple options for certain fields, just choose an arbitrary value. [7 pts]

Original packet (before NAT)			After NAT		
	Source	Dest		Source	Dest
IP	192.168.4.2	201.20.4.7	IP		
TCP Port	5555	20	TCP Port		

- b. **What information does the router need to keep internally in order to properly translate and forward the server's response?**

Assume that the response packet has the same IPs and ports as the server receives, with the roles of source and destination reversed. You don't need to precisely describe what data structures the router is using—just specify what information it will need to store to forward the response. [3 pts]

- c. FTP (File Transfer Protocol) is a protocol that runs on top of TCP, and needs two parallel connections: one for commands, and one for data. First, a client opens an FTP control connection to a server on port 20. Then, (in one particular protocol mode) the client sends a `PORT` command to the server, which includes the client's IP address (as known by the client) and the port the client is listening on (say, port 7000). The server then opens a data connection to the client on that IP and port.

In our example, the client would open a connection to 201.20.4.7 on port 20 and would send a command `PORT 192.168.4.2,7000` to the server. Since our client is behind a NAT gateway, the data connection does not work. **Why not?** [4 pts]

- d. Suppose the FTP client is smart and is able to determine NAT's outside IP address. Thus, it instead sends the command ~~`PORT 201.20.4.7,7000`~~¹ `PORT 105.10.2.5,7000` to the server, but this still does not work. **Why not?** [3 pts]

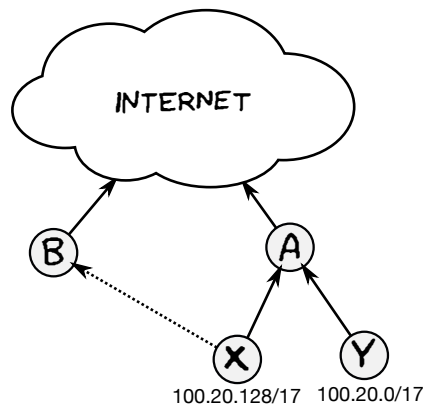
¹This was a typo in a previous version of the exam.

- e. Based on its description in this problem, why does FTP's `PORT` command violate good principles of layering in network protocols? [3 pts]

Bonus If you can change the router's NAT implementation (including understanding the FTP protocol and potentially modifying the the FTP messages), **how could you make the data connection work?** [6 pts]

3. BGP [20 pts]

Looking at the figure below, where nodes are ASes, arrows mean a customer-provider AS relationships. Note that the arrows do not constrain the direction of traffic, they only relate to BGP announcements! Assume that ASes follow the Gao-Rexford model we discussed in class, and that ASes A and B eventually learn all of the advertisements that the other one makes, through their respective providers.



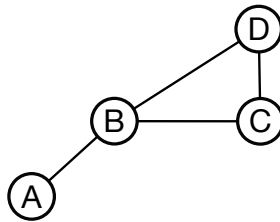
- What is the largest prefix that A can advertise to its providers, given that it has the two customers X and Y with the prefixes in the figure? [4 pts]
- If X decides to also become a customer of B (creating the dashed line), what new prefix will B advertise to its providers? [4 pts]
- What happens if B also decides to advertise Y's prefix, 100.20.0.0/17? [4 pts]

- d. **True or False:** If B and A decide to become peers, B will start advertising Y's prefix. **Explain your reasoning.** [4 pts]
- e. X receives BGP announcements about Y from A, which allows nodes in X to know how to reach nodes in Y. When X becomes a customer of B (ie, when the dashed line is created), **does B receive a route to reach Y via X? Why or why not?** [4 pts]

(This space intentionally left blank—you don't need to fill it with a long answer.)

4. Distance Vector [20 pts]

Consider the following topology in a distance vector protocol, such as RIP, that **does not** use poison reverse or split horizon.



- a. **Suppose the A-B link is disconnected.** Describe a sequence of events (e.g., A announces distance of X, B adds route for A with cost X+1, etc) that leads to a count-to-infinity scenario. [7 pts]
- b. In this topology, does using Split Horizon (when you don't advertise routes learned from some neighbor router X back to X) prevent count-to-infinity from happening? Why or why not? [7 pts]

c. Why can't count to infinity happen in a Link State protocol? [3 pts]

d. Why can't count to infinity happen in a Path Vector protocol? [3 pts]

(This space intentionally left blank—you don't need to fill it with a long answer.)

Physical layer You are working for a company that does satellite communications. You are tasked with designing a new physical layer protocol to replace an older one. The current (old) protocol uses a channel bandwidth of 2MHz, and encodes the signal using ASK (Amplitude-Shift Keying, ie. using just two levels of amplitude), which achieves a rate of 4Mbps. [15 pts]

Recall that Hartley expresses ideal channel capacity C (in bps) based on levels of modulation, such that $C = 2B \log_2(M)$. According to Shannon, the channel capacity based on the signal to noise ratio: $C = B \log_2(1 + S/N)$. For the ranges we care about in this problem, we can ignore the 1—so we can express Shannon's formula as $C \approx B \log_2(S/N)$.

- a. You are designing a new radio to be placed on the satellites. If you assume that the power of the signal will always be at least 512 times stronger than the power of the noise, what is the capacity of the link, in bps? [5 pts]

- b. How many modulation levels can you use to achieve this maximum rate? [5 pts]

- c. List two properties of the signal that you can vary (other than the amplitude) in order to achieve a higher number of modulation levels compared to ASK. [5 pts]