Exam - Midterm

Due: 9:00pm, 23 Oct 2017

Closed Book. Maximum points: 100

NAME: (-1 pt if missing!)

1. Almost Local Networks [35 pts]

Consider the figure below. There are two subnets connected to each other via router R1. S1 and S2 are Ethernet switches. R1 also connects them to the Internet. Since the two networks use private IP addresses, R1 does Network Address Translation (NAT) between the internal networks and the Internet. m1 … m8 are MAC addresses; p0 … p2 are port identifiers at switch S1 (not MAC addresses).

Initially, all ARP tables and MAC learning caches are empty, but router R1's routing table is filled, as are the network configuration of all hosts (address, mask, and gateway). We assume that nodes know each other's IP addresses.

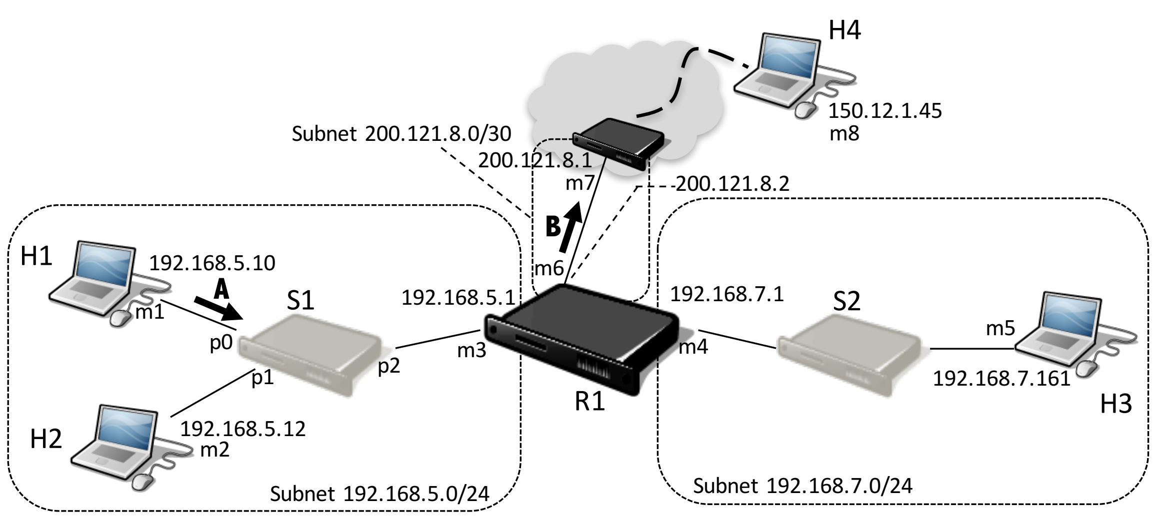

a. [7 pts] H1 sends an IP packet to H3. It has to send an ARP request before sending the packet, and then sends the IP packet. For this packet, while it is crossing link **A** in the figure, fill in the table with the source and destination addresses, both at L2 (Ethernet) and L3 (IP):

	Src Addr	Dest Addr
Ethernet		
IP		

b. [7 pts] Router R1 is a cheap model and can only hold 3 entries in its routing table. Given that it has to route packets to both local subnets, and to the entire Internet (including H4), what are the routing entries in the table? (There would also be a next hop entry, but you don't have to worry about it in this question…)

Destination Network	Out Port

c. [7 pts] After the IP packet successfully reaches H3, what are the contents of H1's ARP table, and of switch S1's MAC learning table? Use the tables below, but you don't necessarily have to use all entries.

H1 ARP Cache

IP	HW Addr

S1 MAC Learning Table

HW Addr	Port

d. [7 pts] Now H1 wants to send a packet to H4, which is a server out on the Internet. Note that R1 does NAT for the outgoing packet, as the 192.168/16 addresses are not routable. What are the source and destination L2 and L3 addresses for this packet when it is crossing link **B** in the figure?

	Src Addr	Dest Addr
Ethernet		
IP		

e. [7 pts] IP addresses enable small routing tables that don't sacrifice reachability. Can we do the same with MAC addresses? Why or why not?

2. Physical Layer [18 pts]

You have just been hired by Blackhole Networks, a company that wants to disrupt the networking with the new Blacktooth protocol (B*l*T). Their slogan is *"No packet ever arrives with errors."*, and they care about reliability. Your are tasked with improving the reliability and efficiency of the B*l*T protocol.

You start with the physical layer. The current protocol is using a channel of 2MHz, and encoding the signal using ASK (just two levels of amplitude). The achieved (binary) bandwidth is 4Mbps.

Recall that Hartley extended Nyquist's result to include levels of modulation, such that $C = 2B\log_2(M)$ bps. According to Shannon, the capacity of the channel is $B\log_2(1 + S/N)$. For the ranges we are worried about, let's ignore the 1, such that $C \approx B\log_2(S/N)$.

a. [6 pts] If you assume that the power of the signal will be always at least 4096 times stronger than the power of the noise, what's the capacity of the link, in bps?

b. [6 pts] How many total levels can you use to achieve the maximum rate?

c. [6 pts] List two additional properties of the signal you can vary (beyond the amplitude) to improve the rate.

3. Link Layer [20 pts]

At the link layer, B*l*T (see previous question) is currently a stop-and-go protocol, with timeouts, retransmissions, and no sequence numbers.

a. [6 pts] You start thinking about ways to improve the protocol. In the list below, given the **properties** on the left, and the protocol **features** on the right, draw lines connecting the features with the properties, i.e., which feature(s) is(are) required to provide the given properties? (There may be more than one feature per property, and features that are required by more than one property.)

Properties	Features
At least once semantics •	• Acknowledgments
At most once semantics •	• Timeouts + Retries
High utilization •	• Sequence Numbers
	• Sliding window

b. [7 pts] You want to add a sliding window protocol to B*l*T. To convince the rest of the team that this is a good idea, draw a graph in the axes below, of the throughput, in bytes per second, versus the window size (in packets or bytes). Using B for the bandwidth of the link, and l for the round-trip time, indicate any inflection points in the graph.

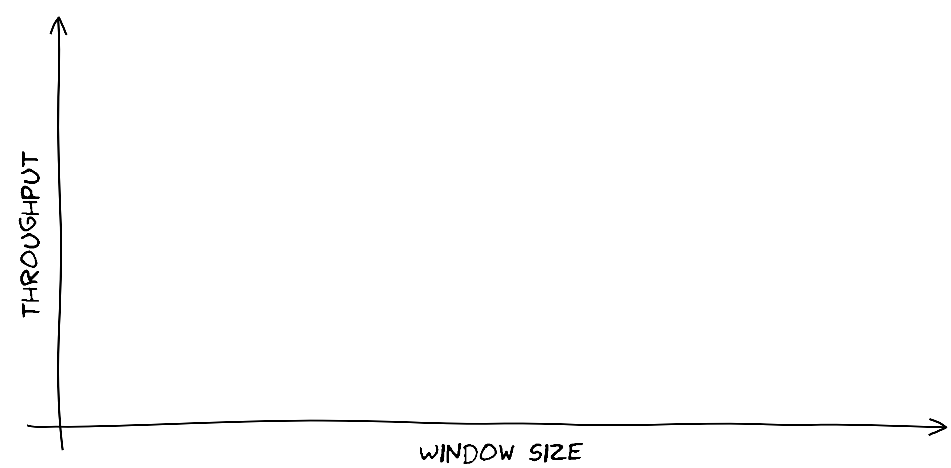

c. [7 pts] You convince them, and you all agree that the window size should be 4 packets. One of your colleagues goes and designs the header, allocating 2 bits for the sequence number. Draw an example, in the timing diagram below, that shows this will not be sufficient, and that the protocol will break. (Hint: you will need drops of packets or acknowledgments...). Briefly explain your reasoning.

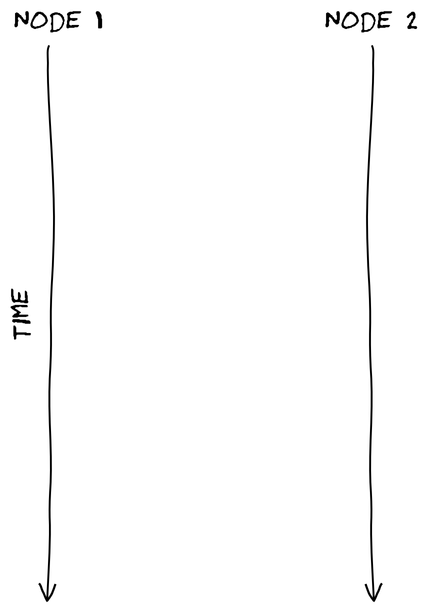

4. BGP [27 pts]

As Blackhole Networks grows, management decided to establish a stronger presence online, and the company bought its own prefix, 128.8.12/22, together with its own AS number, 68231! You are also buying transit from another AS, number 173.

a. [6 pts] It is AS 173's job to make sure the rest of the Internet can reach you through BGP announcements to its neighbors. What are the destination and the AS path in these announcements?

b. [7 pts] An AS always relays all announcements it receives from all neighbors to its customers. True or false, and why?

c. [7 pts] One day your pager rings, because there was a drop of almost 50% of the incoming traffic to *all* your addresses (Let's just assume that all your addresses receive a steady and even spread of traffic, perhaps from all the data you collect from users of your products worldwide). Upon inspection, you noticed that some other AS around the world, number 54668, was making erroneous announcements that included at least part of your prefix. What could these announcements be to achieve this effect (get some, but not all of your traffic)? How could you change your announcements to counter that, given your knowledge of how BGP speakers choose routes?

d. [7 pts] After too many of these incidents, all ASes agree to really work hard to maintain the list of which AS can originate each prefix: an AS now looks at an announcement, and if the first AS in the path does not own the prefix, it rejects the announcement. (Note that the announcements are not signed, as they would be in sBGP.) At ease, you revert to your old announcement (before your counter effect in c. above). This turns out to be a failure, and a couple of days later AS 54668 manages to steal traffic from you again with a different announcement. How can AS 54668 be doing this, if it can't be the originating AS in announcements with any portion of your address space?