

Second Midterm for ECE374
04/08/15
Solution!!

Instructions:

- a. Put your name and student number on each sheet of paper!
- b. The exam is closed book.
- c. You have 90 minutes to complete the exam. Be a smart exam taker - if you get stuck on one problem go on to another problem. Also, don't waste your time giving irrelevant (or not requested) details.
- d. The total number of points for each question is given in parenthesis. There are 100 points total. An approximate amount of time that would be reasonable to spend on each question is also given; if you follow the suggested time guidelines, you should finish with 10 minutes to spare. The exam is 90 minutes long.
- e. Show all your work. Partial credit is possible for an answer, but only if you show the intermediate steps in obtaining the answer. If you make a mistake, it will also help the grader show you where you made a mistake.
- f. Good luck.

Problem	Max. Points	Points
1	32	
2	20	
3	24	
4	24	
Total	100	

Problem 1: (Quickies 32 Points (4 each), 25 minutes)

- a. Briefly explain at what point the TCP congestion control algorithm switches from the slow-start to the congestion avoidance phase? Is it possible that the algorithm is also switching back from the congestion avoidance to slow-start. Briefly explain your answer.

Answer: This transition happens when the slow-start threshold (sshtresh) is reached. Yes, in the case of packet loss or time out.

- b. Specify the IP network address, the host range, and the subnet mask for a subnet that can contain up to 30 clients. What is the maximum number of hosts per subnet if you have 24 subnet mask bits?

Answer: 192.168.0.0; Host range: 192.168.0.1 - 192.168.0.30; Mask bits: 27; 254;

- c. Suppose a web server runs in Host C on port 80. Suppose this Web server uses persistent connections, and is currently receiving requests from two different Hosts, A and B. Are all of the requests being sent through the same socket at Host C? If they are being passed through different sockets, do both of the sockets have port 80? Explain briefly.

Answer: For each persistent connection, the Web server creates a separate "connection socket". Each connection socket is identified with a four-tuple: (source IP address, source port number, destination IP address, destination port number). When host C receives an IP datagram, it examines these four fields in the datagram/segment to determine to which socket it should pass the payload of the TCP segment. Thus, the requests from A and B pass through different sockets. The identifier for both of these sockets has 80 for the destination port; however, the identifiers for these sockets have different values for source IP addresses. Unlike UDP, when the transport layer passes a TCP segment's payload to the application process, it does not specify the source IP address, as this is implicitly specified by the socket identifier.

- d. Do routers have IP addresses? If so, how many IP addresses and does each IP address have to belong to the same or a different subnet?

Answer: Yes; As many as the router has interfaces; Each has to belong to a different subnet;

- e. Briefly explain how IPv6 datagrams can be tunneled over an IPv4 backbone. You can use a simple network (as sketched below) that consist of 4 routers that are daisy chained to explain how the tunneling works.

IPv6 ---- IPv4 ---- IPv4 ---- IPv6

Answer: IPv6 ---- IPv6 encapsulated in IPv4 ---- IPv6 encapsulated in IPv4 ---- IPv6

- f. Why is an ARP query sent within a broadcast frame? Why is an ARP response sent within a frame with a specific destination MAC address?

Answer: Sender does not know MAC address of host that has specific IP address,

yet. Learns MAC of host from incoming ARP request, no broadcast necessary.

- g. Briefly describe how Ethernet's exponential backoff works. What is one reason why Ethernet's exponential backoff might be better than randomizing retransmission attempts over a fixed-length time interval?

Answer: Ethernet maintains an interval of time T over which it will randomize when it will attempt a retransmission. After each collision for the same packet, it doubles the length of T up to some fixed max. This is better than just a single, fixed value of T since when there are a lot of collisions the interval over which randomization is done will be large, allowing just one node to successfully be transmitting. When there are only a small number of colliding nodes, the retransmission will be randomized initially over a small T , allowing a node to transmit more quickly.

- h. Briefly describe the difference between ad-hoc and infrastructure mode in wireless networks. Give one example advantage of ad-hoc over infrastructure mode.

Answer: Ad-hoc: node-to-node communication; Infrastructure: Nodes do only communicate with base station and not with other nodes. Fault tolerance; if base station fails, no more communication is possible.

Problem 2: Congestion Control & Avoidance & Routing (20 Points, 20 minutes)

- a. (2 Points) Why do you think TCP does not include RTT measurements for retransmitted packets in the SampleRTT?

Answer: Let's look at what could go wrong if TCP measures SampleRTT for a retransmitted segment. Suppose the source sends packet P1, the timer for P1 expires, and the source then sends P2, a new copy of the same packet. Further suppose the source measures SampleRTT for P2 (the retransmitted packet). Finally suppose that shortly after transmitting P2 an acknowledgment for P1 arrives. The source will mistakenly take this acknowledgment as an acknowledgment for P2 and calculate an incorrect value of SampleRTT.

- b. (2 Points) Does UDP require a mechanism to estimate the RTT between sender and receiver? Briefly explain your answer?

Answer: No, UDP does not perform retransmissions and, thus, does not require an RTT-based timeout timer.

Consider sending a large file from a host to another over a TCP connection that has NO loss

- c. (5 Points) Suppose TCP uses AIMD for its congestion control without slow start. Assuming cwnd increases by 1 MSS every time a batch of ACKs is received and assuming approximately constant round-trip times, how long does it take for cwnd to increase from 5 MSS to 10 MSS?

Answer: It takes 1 RTT to increase CongWin to 6 MSS; 2 RTTs to increase to 7

MSS; 3 RTTs to increase to 8 MSS; 4 RTTs to increase to 9 MSS; 5 RTTs to increase to 10 MSS.

- d. (5 Points) What is the average throughput (in terms of MSS and RTT) for this connection up through time 6 RTT?

Answer: In the first RTT 5 MSS was sent; in the second RTT 6 MSS was sent; in the third RTT 7 MSS was sent; in the fourth RTT 8 MSS was sent; in the fifth RTT, 9 MSS was sent; and in the sixth RTT, 10 MSS was sent. Thus, up to time 6 RTT, $5+6+7+8+9+10 = 45$ MSS were sent (and acknowledged). Thus, we can say that the average throughput up to time 6 RTT was $(45 \text{ MSS})/(6 \text{ RTT}) = 7.5 \text{ MSS/RTT}$.




Figure 1

Consider Figure 1 for the following problem.

- e. (2 Points) In this case, the two networks A and B are connected through network C to the rest of the Internet. What network address or addresses for A and B are advertised by the gateway router (that connects C to the Internet)? Specify network address and subnet mask?

Answer: 128.119.0/16

- f. (2 Points) Assume that the router that connects C to both A and B has interface 1 connected to A and interface 2 connected to B. For these two networks, specify the forwarding table entries in the router of network C.

Answer: 128.119.0/17 | Interface 1, 128.119.128/17, Interface 2.

- g. (2 Points) Now assume that a host somewhere in the Internet (the host is not located in networks A, B, and C) sends a packet to *i*) a host in network A and *ii*) to a host in network B. Explain how the router (from problem b.) in network C determines which interface to send each of these packets.

Answer: longest prefix matching.

Problem 3: DHCP & NAT & Routing (24 Points, 20 minutes)

This problem is concerned with DHCP and NAT. Let's first focus on DHCP.

- a. (4 Points) What four pieces of information does a DHCP client obtain from the DHCP server after a successful communication between the two?

Answer: IP address, IP address of first hop router, name and IP address of DNS server, network mask.

- b. (4 Points) What are the destination IP and MAC address in the discover request sent by the DHCP client? What is the source IP address and MAC address of the offer message sent by the DHCP server?

Answer: IP = 255.255.255.255 MAC: FF:FF:FF:FF:FF:FF, IP = IP address of DHCP server, MAC = MAC address of DHCP server.

Now let's look at how NAT work.





Figure 2

- c. (2 Points) Specify IP addresses for all interfaces of the local network. Also assign MAC addresses to these interfaces.

Answer:



- d. (6 Points) Now assume that the top-most client in the local network sends an HTTP request to <http://www.umass.edu> (IP addr: 128.119.103.148) and the local port for the TCP connection for this HTTP request is 7000. Fill out the NAT table below with the information that would be entered in that table after the HTTP request has been forwarded from the NAT router into the WAN.

In addition, specify destination and source IP addresses and ports for

segment 1, 2, 3, and 4 as indicated in Figure 2.

NAT translation table	
WAN side address/port	LAN side address/port
128.119.103.148, 5000	10.2.0.2, 7000
128.119.103.148, 5001	10.2.0.3, 7000
128.119.103.148, 5002	10.2.0.4, 7000

Answer:

- e. (2 Points) Also specify MAC source and destination address for 1 and 2.

Answer: *BB:BB:BB:BB:BB:BB, AA:AA:AA:AA:AA:AA*

When the packet leaves the NAT router it needs to be routed through the Internet to its destination (the UMass web server). For this problem, let's assume the network shown in Figure 3.




Figure 3

- f. (6 Points) Assume node u represents the NAT router and node z is the router the UMass web server is connected to. By applying Dijkstra's shortest path algorithm fill out the table below and determine the shortest path from u to z.

Step	N'	D(v), p(v)	D(w), p(w)	D(x), p(x)	D(y), p(y)	D(z), p(z)
0	u					
1						
2						
3						
4						
5						

Answer:

Step	N'	$D(v), p(v)$	$D(w), p(w)$	$D(x), p(x)$	$D(y), p(y)$	$D(z), p(z)$
0	u	2, u	∞	1, u	∞	∞
1	ux	2, u	4, x		2, x	∞
2	uxy	2, u	3, y			4, y
3	uxyv		3, y			4, y
4	uxyvw					4, y
5	uxyvwz					

Problem 4: Ethernet (24 Points, 25 minutes)

The first part of this problem focuses on 802.3.

- a. (5 Points) In the scenario shown in Figure 4, A is sending a frame to A' and after receiving that frame, A' sends a frame back to A. Explain the steps that are carried out by A, A', and the switch to populate its table such that packets from A can be switched to A' and vice versa.




Figure 4

Answer:

MAC Addr	Interface	TTL
A	1	60
A'	4	60

Step 1: A sends frame to switch on port 1, switch know that it can reach A on that port.

Step 2: Switch floods packet out all ports except 1 because it does not no to which port A' is connected to.

Step3: A' replies and switch learns that A' is connected to port 4.

- b. (3 Point) Does each port on the switch have to have its own, unique MAC address? Briefly explain your answer.
Answer: No. Since each node is directly connected to the switch MAC addresses are not required at the switch ports (a packet sent from a host will always end up at the switch port that connects to the host). Actually, the switch is transparent for the hosts. E.g., in the case of A sending a frame to A', the frame only carries A and A' MAC addresses and no information about the switch or its ports.
- c. (2 Points) In Figure 5, A is sending an ARP request to determine the MAC address it should use for a frame destined to B. What MAC address will be returned to A in the ARP reply. Briefly explain your answer.

Answer: 20:FF:3A:BC:01:4E. Since B is not in the same switching domain as A frames have to be sent to the left interface of R1.

- d. (2 Points) How many interfaces of all the ones shown in Figure 5 receive the original ARP request sent by A and what is the destination MAC address on the frame containing the ARP request?

Answer: 6, FF:FF:FF:FF:FF:FF





Figure 5

Answer:

- e. (4 Points) Let's assume that router R1 shown in Figure 5 connects two IP subnets. In the scenario shown in this figure a host that is connected to either switch 1 or 4 cannot change from its current subnet to the other subnet without changing its physical connection to either switch 2 or 3. Can you think of a solution that would allow this host to stay connected to its original switch. (Hint: Think VLAN). If so, draw a sketch or your proposed solution.

Answer:



Let's now focus on 802.11

- f. (2 Points) Briefly explain the two methods through which a wireless 802.11 node learns about access points it can communicate with.

Answer:

i) *Passive scanning, beacon frame from AP, association from node, reply from AP.*

ii) *Active scanning, probe request from node, probe responses from APs, association request from node, reply from AP.*

- g. (2 Points) Explain how the ACK transmitted by the base station after the data transmission from the host signals that the transmission from the host has finished also to a potential hidden terminal.

Answer: *A host that cannot observe the transmission of another host but is in transmission range of the base station will be able to detect the ACK and no that the transmission of the other host has finished.*

- h. (4 Points) Figure 6 shows a typical setup for a infrastructure mode wireless network where the BSS includes several hosts and a base station. The later is connected to the Internet via a route. For this scenario show all MAC address in correct order in the 802.11 frame that is transmitted from the host H1 to the base station and the frame transmitted from the base station to the router. If the addresses in the two frames are different, briefly explain why.

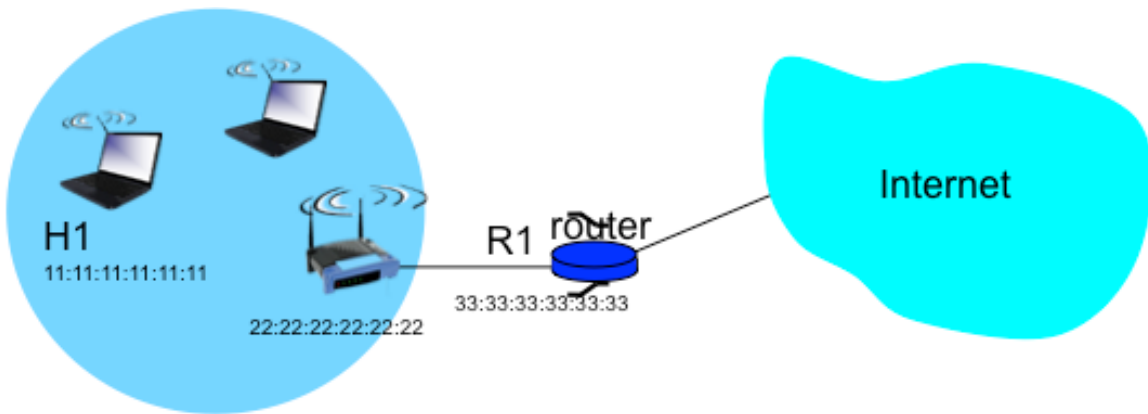


Figure 6

Answer:

