Midterm Exam II

CMPSCI 453: Computer Networks

Fall 2011

Prof. Jim Kurose

Instructions:

- **Please use two exam blue books –** answer questions 1, 2 in one book, and the remaining three questions in the second blue book.
- Put your name and student number on the exam books NOW!
- The exam is closed book.
- **You have 80 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.
- 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 8 minutes to spare. The exam is 80 minutes long.
- Show all your work. Partial credit is possible for an answer, but only if you show the intermediate steps in obtaining the answer.
- Good luck.

<u>Question 1: ``Quickies'' (22 points, 18 minutes)</u>

Answer each of the following questions *briefly*, i.e., *<u>in at most a few sentences.</u>* (Really – your answer should be very **brief** here)

a) (3 points) Consider a router with *N* input lines, each with input link rate *R* and an internal switching fabric that is *2N* times faster than *R*. Where in this router can packet queue form? Explain your answer. *Answer: Queueing will only occur at the output ports. Since the switch is more than N times fast than the input rate, all arriving packets in a slot can be move from input port to output port in that slot.*

b) (3 points) Suppose BGP router A sends a BGP path vector to BGP peer router B. BGP peer B is connected to BGP peer C. Must B advertise that path to C? *Answer: No. It is up to B's internal policy about what routes to advertise to others. Recall that we discussed how an ISP will generally only carry traffic to/from its customers, and not carry transit traffic (i.e., traffic that is both sources and destined in non-customer networks).*

c) (4 points) TRUE OR FALSE: *(i)* In IP forwarding, it's possible for two packets to take *different* paths from a common router to the same destination, based on their source IP address. *Answer: False. (ii)* In MPLS forwarding, it's possible for two packets to take *different* paths from a common router to the same destination, based on their source IP address. *Answer: True.*

d) (4 points) 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 is 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 being 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.*

e) (4 points) Describe the role of the home agent and the foreign agent in mobile IP. What is sent in a datagram from the home agent to the foreign agent, and what is sent from the foreign agent to the home agent? *Answer: The foreign agent will let the home agent know when the visiting node joins the foreign network, and assigns a care-of-address to the visiting node. It will also relay packets from the foreign agent to the visiting node. The home agent knows (see above) the network where the visiting node is visiting and will forward incoming datagrams for that node to the foreign agent at the visited network. The foreign agent send registration info to the home agent; the home agent sends encapsulated datagrams for the mobile node to the foreign agent.*

f) (4 points) Suppose that Bob and Alice have access to a public key system that makes their public keys available to each other. K_B^- and K_B^+ are Bob's private and public keys, respectively. Each knows its own private key. Describe how Bob and Alice can use these keys so that Bob can know that when he is communication with Alice that *(i)* he is indeed communicating with Alice, and *(ii)* Alice is live (i.e., he is not receiving a playback of earlier info sent from Alice? *Answer: Bob creates a nonce and sends it to Alice and asks her to encrypt it with her private key. Bob then applies Alice's public key to the nonce and if he recovers the nonce value, accepts Alice as live (since the nonce was signed) and* *as Alice (since only Alice has access to her private key, and only Alice's private has the property that* $K_A^+ K_A^- (m) = m$.

Question 2: Congestion Control and TCP (15 points, 10 minutes)

Consider four Internet hosts, each with a TCP session. These four TCP sessions share a common bottleneck link - all packet loss on the end-to-end paths for these four sessions occurs at just this one link. The bottleneck link has a transmission rate of *R*. The round trip times, *RTT*, for all fours hosts to their destinations are approximately the same. No other sessions are currently using this link. The four sessions have been running for a long time.

a) (4 points) What is the approximate throughput of each of these four TCP sessions? Explain your answer briefly. *Answer: R/4 since TCP shares bandwidth fairly.*
b) (4 points) What is the approximate size of the TCP window at each of these hosts? Explain briefly how you arrived at this answer. *Answer. Recall that roughly throughput = W/RTT or W = throughput * RTT = R*RTT/4.*
c) (4 points) Suppose that one of the sessions terminates. What is the new throughput achieved by each of the three remaining sessions? Briefly describe how this new throughput is reached (i.e., what do the TCPs in the remaining three hosts do that results in this new throughput being achieved). *Answer: R/3 since TCP shares bandwidth fairly*
d) (3 points) Now suppose that one of the three hosts starts a second session that also crosses this bottleneck link. What is the throughput achieved (in aggregate) by the one host with two sessions, and by each of the two hosts with one session each? *Answer: each session will again get R/4, so the one host with two sessions will get R/2 in aggregate and the other teo hosts will each get R/4.*

Question 3: Link State and Distance Vector Routing (20 points, 20 minutes)

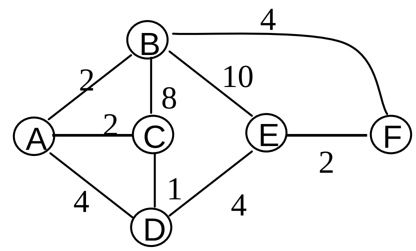

Consider the network shown to the right.

a) (10 points) Show the operation of Dijkstra's (Link State) algorithm for computing the least cost path from **E** to all destinations. *Answer:*

N	D(A),p(A)	D(B),p(B)	D(C),p(C)	D(D),p(E)	D(F),p(F)
E	infty	10,E	infty	4,E	2.E
EF	infty	6,F	infty	4,E	
EFD	8,D	6,F	5,D		
EFDC	7,C	6,F			
EFDCB	7,C				

b) (2 points) From these results, show the least cost path from E to A, and briefly describe (in a sentence) how you got that answer from your work in part a).

c) (6 points) What are distance vectors in nodes E, D, and C? In one or two sentences, explain how the least cost path from E to A was determined by E based on these three distance vectors. *Note: you do not have to run the distance vector algorithm; you should be able to compute distance vectors by inspection. Answer For nodes E, D. and C respectively:*
[D_E(A)=7, D_E(B)=6, D_E(C)=5, D_E(D)=4, D_E(F)=7]
[D_D(A)=3, D_D(B)=6, D_D(B)=5, D_D($)=4, D_D(F)=6]
[D_C(A)=2, D_C(4)=6, D_C(D)=1, D_C(E)=5, D_C(F)=7]
E's distance to A is the sum of its cost to D (4) plus D's cost to A which is 3. D's cost to A is D/sits cost to C which is 1, plus C's cost to A which is 2. Therefore the total cost from E to A is 2+1+4= 7.

d) (2 points) Let us focus again on node E and distance vector routing. Suppose all distance vectors have been computed in all nodes and now suppose that the link from E to B goes down. Approximately how many distance vector messages will be sent by node E as a result of this link going down? Explain your answer. *Answer: no messages would be sent out by E, since the EB link is not on E's shortest path to any destination.*

Question 4: Link and Network Layer operating (25 points, 20 minutes)

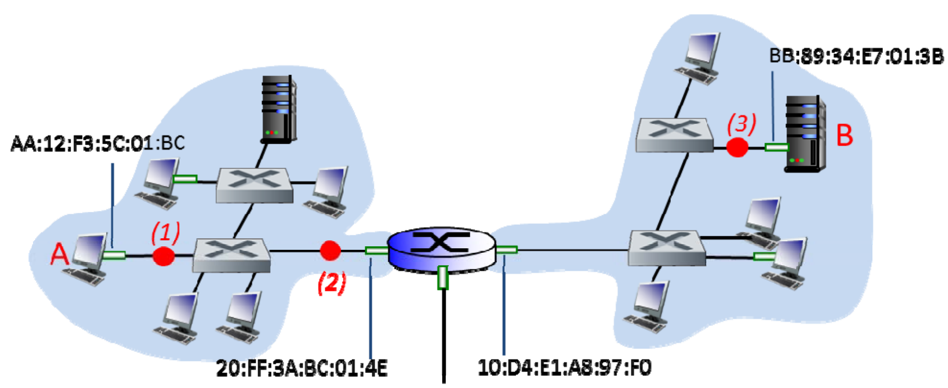

Consider the network shown above.

a) (4 points) Assign IP address ranges to the subnets containing hosts A and B, and assign IP addresses in these ranges to hosts A and B. (You don't have to assign IP addresses to any hosts except A and B, but you do need to specify the address range being used by each subnet). Your subnet addressings should use the smallest amount of address space possible. *Answer: Because there are less than eight but more than 4 nodes in each subnet, we'll need three address bits for each subnet. So let's assign the left subnet XX.YY.ZZ.xxxx0***/29, where the XX.YY.ZZ are 8 bit numbers. Each x is a bit and the three *'s correspond to the three address bits for this network. For the right subnet, well use XX.YY.ZZ.xxxx1***/29. A will have an IP address of XX.YY.ZZ.xxxx0000 and B will have an IP address of XX.YY.ZZ.xxxx1000.*

b) (4 points) What IP address range can the router advertise to the outside for all of the hosts reachable in these two subnets? Again, you should choose your answer in a) above so that the minimum-size address space is advertised here. *Answer: XX.YY.ZZ.xxxx/28*

c) (3 points) Does the router interface with link-layer address `20:FF:3A:BC:01:4E` have an IP address? If so, what is the role of the IP address of the router's IP interface in forwarding datagrams through the router. *Answer: Yes. That's the address that a host in the left network will use to determine the MAC address to send frames to, containing datagrams that need to be routed through the router. The router address however, won't appear in the IP datagram.*

d) (5 points) Consider an IP datagram being sent from A to B using Ethernet as the link layer protocol in all links in the figure above. What are the *(i)* Ethernet source and destination addresses and *(ii)* IP source and destination addresses of the IP datagram encapsulated within the Ethernet frame at points (1), (2), and (3) in the above example for a datagram going from A to B.
(1): ETH source: aa:12;F3:5C:01:BC, ETH dest: 20:FF:3A:BC:01:4E
IP source: XX.YY.ZZ.xxxx0000; IP dest: XX.YY.ZZ.xxxx1000 see (a) above
(2) same as (1)
(3) ETH source: 10:D4:E1:A:97:FO, ETH dest: BB:89:34:E7:01:3B*
IP source: XX.YY.ZZ.xxxx0000; IP dest: XX.YY.ZZ.xxxx1000 same as (1) above

e) (6 points) Suppose all switches in the above example are learning switches. Consider the datagram being sent from A to B; neither A nor B have sent any frames or datagrams in the network before.

- How many of the 11 hosts in the network receive the frame containing the datagram sent by A? Explain your answer briefly. *Answer: all 11, since no switch knows where B is located (since B hasn't sent anything), all switches will broadcast the frame containing the IP datagram from A. Note that different frames are broadcast on the left and right subnets (e.g., the frames have different source and destination MAC addresses, see above), but both contain the datagram from A.*
- Suppose the server in the upper part of the *left* network sends a datagram to A shortly *after* the A-to-B datagram is sent. How many of the 11 hosts in the network receive the frame containing the datagram sent by this server? Explain your answer briefly. *Answer: Only A will receive that, since all of the switches know the outgoing port leading to A, as a result of learning where A is, as a result of the initial A-to-B transmission.*

f) (3 points) Suppose A sends out an ARP request, and this ARP request is in the very first frame sent in the network above (i.e., even before the original A-to-B datagram). How many of the 11 hosts in the network receive the frame containing this ARP request? Explain your answer briefly. *Answer: ARP broadcasts are restricted to a subnet and generally do not pass through the router, so all 5 other hosts in the left network will receive the ARP broadcast (as will the leftmost interface on the router).*

<u>Question 5: Wireless (18 points, 20 minutes)</u>

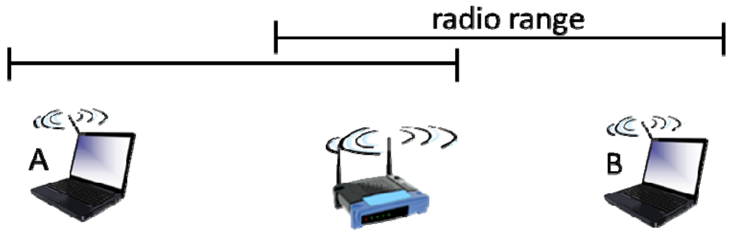

Consider the example above in which two wireless hosts, A and B want to exchange 802.11 frames with each other. A frame sent from A will be received by the base station (assuming a transmission from B does not interfere with the transmission at the base station); *after* receiving the frame from A, the base station will send the frame to B.

a) (4 points) Suppose that A (similarly B) will send an RTS messages to the base station when it wants to send; the base station responds with a CTS that is received by A and B. What is the purpose of the RTS and CTS messages? *Answer: The purpose of the RTS message is to request a reservation of the channel (i.e., so others will not send); the CTS is a grant to prevent all who hear the CTS from sending.*

b) (4 points) Does this use of RTS and CTS message ensure that data frames from A and B never collide? Briefly discuss your answer. *Answer: not always. If a CTS is not received (e.g., bit errors) then a node will not refrain from transmitting, so and can collide.*

c) (4 points) Suppose that the frames that A is sending to B contain TCP segments (i.e., that there is an A-to-B TCP session) and suppose that the wireless links are all noisy, so that packet loss occurs due to bit errors. Give one reason why TCP would *not* be a good protocol to be using in this scenario. Briefly discuss your answer *Answer: TCP interprets a lost segment as an indication of congestion and reduces its send window. In this case, if a segment is lost it is because of bit errors, not congestion, and so there is no reason for TCP to reduce is send window size.*

d) (4 points) Suppose that the wireless channel between A and the base station (similarly B and the base station) can support a transmission rate of 54 Mbps. What is the approximate A-to-B TCP throughput achieved in the scenario above? You can assume that TCP segments are large, so that frame/datagram/segment field overhead is relatively small. *Answer: Note that every TCP segment must cross two hops, and when something is being sent on one hop, one can't send on the other hop. Similarly, a TCP ACK must be sent across two hops. Thus the maximum throughput we would expect is around 54/4 or around 13 Mbps of TCP throughput.*

e) (2 points) Describe two purposes (uses) of the beacon frame in 802.11. *Answer: a beacon frame advertises an SSID, and also serves to let waking stations know if there are queued data frames for them.*