Final Exam for ECE374
05/03/12
Solution!!

Instructions:

- Put your name and student number on each sheet of paper!
- The exam is closed book.
- 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.
- 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.
- 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.
- Good luck.

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

a. Suppose that Bob and Alice have access to a public key system that makes their public keys available to each other. Each knows its own private key.
   1. Suppose Bob has a document, m, that he wants to digitally sign. How does he do so? Efficiency is not a concern here, any digital signature technique is fine.
      Answer: Bob will encrypt the document, m, using his private key. This serves as Bob's digital signature on the document. If we want to be more efficient, Bob could compute a hash of m, and then encrypt hash(m).
   2. What does Alice do to verify Bob's digital signature?
      Answer: Alice applies Bob's public key to the signed version of m, i.e., computes $K_B^+(K_B^-(m))$. If that yields the message m, then Alice knows that Bob has signed the document.

b. Where and why does packet loss occur within a router?
   Answer: packet loss occurs in buffer in either the input or output line cards, because the memory in the buffer is finite, and the input rate to the buffer exceed the output rate of the buffer over some period of time.

c. Suppose that we want to change the IP address of gaia.cs.umass.edu from 128.119.40.186 to 128.119.40.187 and change this mapping in the DNS authoritative name server for gaia.cs.umass.edu. Once this mapping is changed in the authoritative name server, will all future references (generated anywhere in the Internet) to gaia.cs.umass.edu then be sent to 128.119.40.187? Explain briefly (in two or three sentences).
   Answer: Local DNS caches throughout the Internet will not time out the old mapping of gaia.cs.umass.edu to 128.119.40.186 until the valid interval originally associated with that mapping times out. Until that happen, local DNS caches will not query into the system for gaia.cs.umass.edu and hence would not learn the new mapping.

d. What is meant by the term "encapsulation"?
   Answer: This means that a protocol executing at a given layer in the protocol stack takes a message for the upper layer, wraps it in a new packets with new header (and possible trailer) fields.

e. Consider a server-side socket that is used to communicate from server to client. In the case of a TCP socket, can data being read from the server-side socket have been sent by more than one client? Explain briefly. In the case of a UDP socket, can data being read from the server-side socket have been sent by more than one client? Explain briefly.
   Answer: With TCP, the client-to-server socket is created by the return from the accept() call at the server, binding the client that was accepted to the server via the newly created socket. Thus only one client is associated with the TCP socket. With UDP, any client can send to the (same) UDP socket on the server.

f. What is meant by *demultiplexing* a protocol data unit up to an upper-level protocol?
   Answer: this refers to passing the decapsulated data unit up to the appropriate higher level protocol. This is done by looking at the upper-layer protocol field.

g. Does the Internet checksum always detect errors in thetransmitted segment?

Explain your answer in a sentence or two.
Answer: No. For example if two 16-bit word values are swapped, this would not be detected since the sum is unchanged.

h. Suppose a TCP SYN message is sent from a client with IP address 128.119.40.186 and client port number 5345 to a server with IP address 41.123.7.94 and server port number 80 (HTTP).
    1. Once the TCP connection has been established, what will be the client-side IP address, client-side port number, server-side IP address and server-side port number of the TCP segment carrying the HTTP GET message.
    Answer: exactly as specified in the problem statement.
    2. Will the TCP SYN message and the HTTP GET message be directed to the same socket at the server? Explain in one or two sentences.
    Answer: NO. The GET will be directed to the new socket that was created when the TCP SYN messages was accepted (i.e., the socket returned from the wait on the .accept() on the welcoming socket). Note that the TCP SYN and the GET will both be addressed to port 80 on the server, however.

Problem 2: TCP Potpourri (24 Points, 20 minutes)

a. What is the purpose of the receiver-advertised window in TCP?
Answer: this allows the receiver to tell the sender how much unacknowledged data can be in flight.
b. Identify 6 fields in the TCP header, and give a one-sentence description of each of their uses.
Answer: souce port, dest port, seq num, ack num, checksum, receiver window, options, SYN, FIN, ACK, urgent pointer, PSH.
c. Consider two TCP sessions that must share a link's bandwidth. One of the TCP connections has been running for quite some time and has built up a large TCP sending window. The second connection then starts up with an initially small window. Long term, what will be the relative throughput achieved by these two TCP sessions? Explain.
Answer: We saw that TCP will cause two senders to eventually fairly share the links bandwidth, so each will eventually have the same size window (assuming their RTT is the same).
d. Consider a TCP session and a UDP session that must share a link's bandwidth. Of course, both sessions would ideally like to send as fast as they can. Long term, what will be the relative throughput achieved by these two sessions. Explain.
Answer: Since UDP can send as fast as it wants, it can use all of the bandwidth (e.g., in the limit that it sends infinitely fast, as soon as buffer space becomes free in a router, that buffer will be filled by a UDP segment. TCP segments will always be lost, causing TCP to keeps its window at 1 segment, which when sent is always lost.
e. Consider the scenario in which the RTT between a TCP sender and receiver is 20ms and *ssthresh* is initially set to 4 segments. Complete the diagram below until after the first transmission in the congestion avoidance phase is finished. *Note: Assume that no losses occur in this case.*

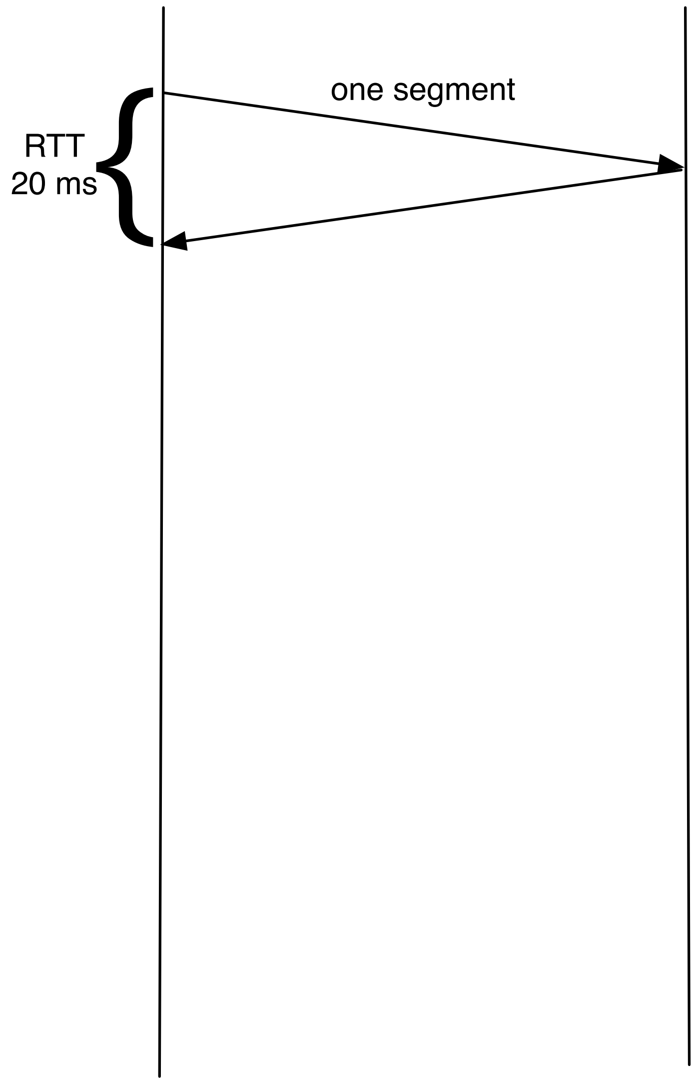

Answer:

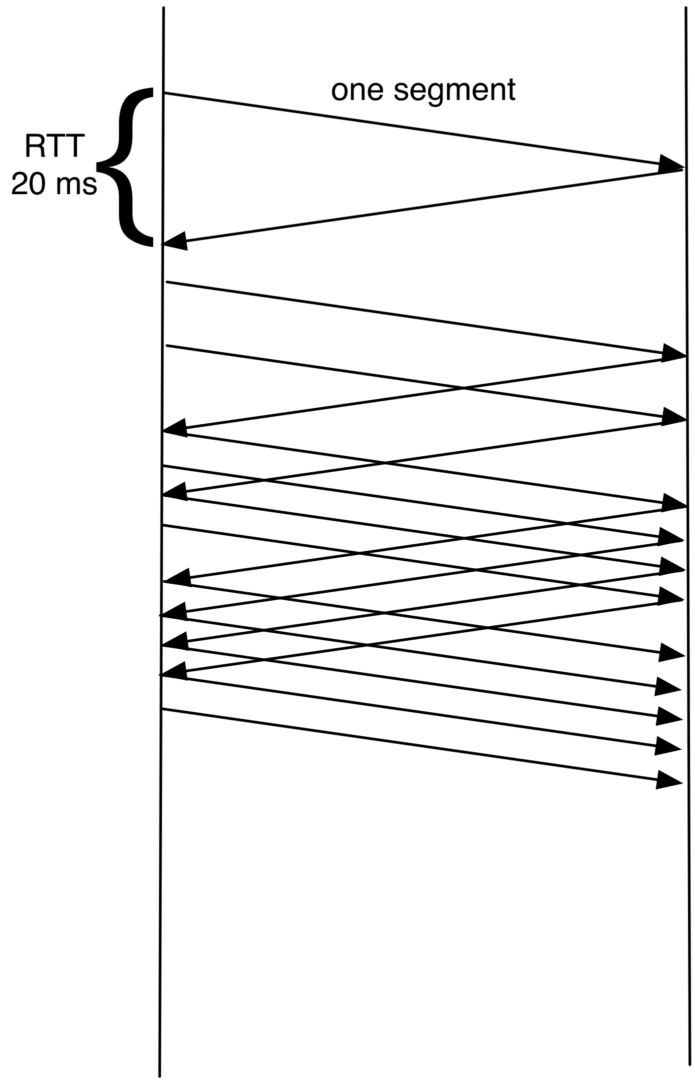
one segment
RTT
20 ms

Problem 3: Delays and Throughput (20 Points, 15 minutes)

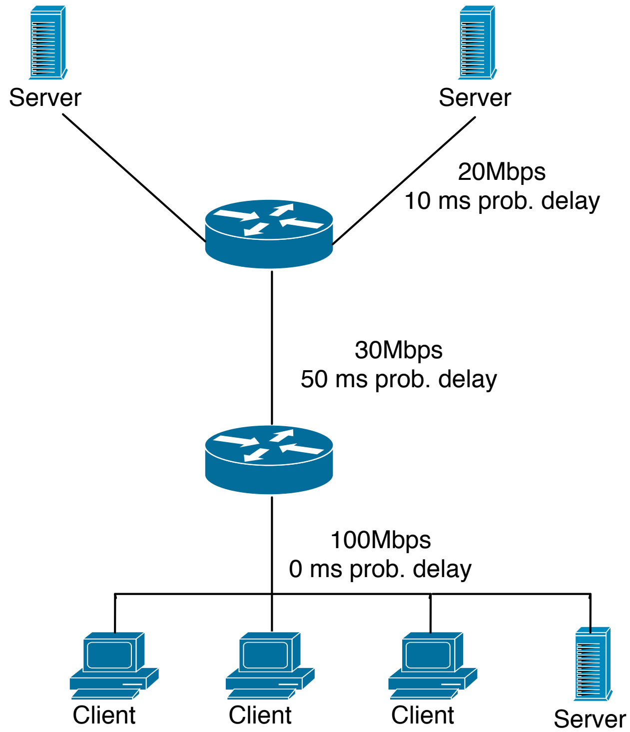

Figure 1

Consider the scenario in Figure 1, in which (from the bottom up) three hosts and a local logging server (that stores information that is sent to it) are connected to a router and to each other by a 100 Mbps link, with an near-zero ms propagation delay. That router in turn is connected to another router over a 30 Mbps link with a 50 ms propagation delay, and that latter router is connected to two remote logging servers, each over a 20 Mbps link with a 10 ms propagation delay.

a. Suppose a host sends a logging message directly to one of the remote logging servers. The logging message is 10K bits long. What is the end-to-end delay from when the logging message is first transmitted by the host to when it is received at the remote server? Assume that the request goes directly to the server, that there are no queueing delays, and that node (router) packet-processing delays are also zero.
Answer: given the 10K bit packet, it takes .0005 secs to send this packet over a 20 Mbps link. 0.000333 secs to send over a 30 Mbps link, and .0001 secs over the 100 Mbps link. The total transmission time end-to-end is this .0009333 secs. The total propagation delay is 60 ms. Therefore the total end-end delay is .0609333 secs.

b. Assume that each of the three hosts generate logging messages at the same rate; each host is equally like to send a logging message to either of the two remote servers. No traffic is directed to the local logging server. What is the maximum rate at which the clients can send logging messages to the remote

servers?
Answer: the link between routers is the bottleneck link, allowing 30 Mbps to be delivered to the two servers combined, or 15 Mbps to be delivered to each server. Since each message is 10K bits, this is 1.5K logging messages per second.

c. Now assume that the local logging server is ON and only one host is active (generating) logging messages and that host is only sending messages to one of the remote logging servers. Suppose that 50% of the logging messages are directed locally and the other 50% directed to this remote server. What is the maximum rate at which this host can generate and send logging messages (both local and remote combined, given there is a 50/50 ratio of local/remote transmissions) in this scenario?
Answer: The maximum rate at which the host can generate remote logging messages is 20 Mbps or 2K logging messages per second. Local messages can be generated that the same rate, so the overall rate is 40 Mbps or 4K logging messages per second.

d.

Problem 4: Wireless and Multimedia (24 Points, 20 minutes)

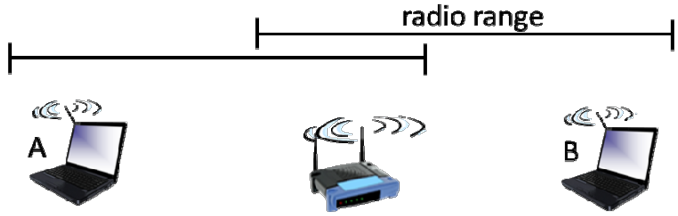

Consider the example above in which two wireless hosts, A and B want to exchange 802.11frames with each other. A frame sent from A will be received by the base station (assuming a transmission from B does not interfere with the transmission at the base station); after receiving the frame from A, the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

a. 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. 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. 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. 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. Now assume a video is streamed from the base station to host A. To ensure reliability on a wireless link 802.11 performs retransmissions if an ACK is not received on time or not received at all.

First, explain the definition of *jitter*. Second, how will the 802.11-based retransmission influence the jitter and thus the quality of the received video, assuming the number of retransmissions are randomly distributed between 0 and 7.
Answer: Jitter is the deviation from the expected periodicity of the data. In the case of video a frame is expected every .x seconds.
Retransmission can increase jitter significantly and lead to buffer underflow.

f. 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.