Final for ECE374
05/06/13
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 there is exactly one packet switch between a sending host and a receiving host. The transmission rates between the sending host and the switch and between the switch and the receiving host are R_1 and R_2, respectively. Assuming that the switch uses store-and-forward packet switching, what is the total end-to-end delay to send a packet length L? (Ignore queuing, propagation delay, and processing delay.)

Answer: At time t_0 the sending host begins to transmit. At time $t_1 = L/R_1$, the sending host completes transmission and the entire packet is received at the switch (no propagation delay). Because the switch has the entire packet at time t_1, it can begin to transmit the packet to the receiving host at time t_1. At time $t_2 = t_1 + L/R_2$, the router completes transmission and the entire packet is received at the receiving host (again, no propagation delay). Thus, the end-to-end delay is $L/R_1 + L/R_2$.

b. Given the following forwarding table, complete the table below by specifying on which of the outgoing interfaces each destination address will be forwarded.

Destination Address Range	Link Interface
11001000 00010111 00010*** *********	0
11001000 00010111 00011000 *********	1
11001000 00010111 00011*** *********	2
otherwise	3

Destination Address	Link Interface
11001000 00010111 00010110 10100001	
11001000 00010111 00011000 10101010	
11001000 00010111 00011100 10101010	
11001000 00010111 10010110 10100001	

Answer:

Destination Address	Link Interface
11001000 00010111 00010110 10100001	0
11001000 00010111 00011000 10101010	1
11001000 00010111 00011100 10101010	2
11001000 00010111 10010110 10100001	3

c. 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. What is the approximate throughput of each of these four TCP sessions? Explain your answer briefly.

Answer: R/4 since TCP shares bandwidth fairly.

d. Briefly describe how web caching can reduce the delay in receiving a requested object. Will Web caching reduce the delay for all objects requested by a user or for only some of the objects? Why?
Answer: Web caching can bring the desired content "closer" to the user, possibly to the same LAN to which the user's host is connected. Web caching can reduce the delay for all objects, even objects that are not cached, since caching reduces the traffic on links.

e. How many keys are derived from the *Master Secret* that is exchanged during an SSL handshake? Describe the individual keys!
Answer: 4, K_c = encryption key for data sent from client to server, M_c = MAC key for data sent from client to server, K_s = encryption key for data sent from server to client, M_s = MAC key for data sent from server to client

f. If a node has a wireless connection to the Internet, does that node have to be mobile? Suppose that a user with a laptop walks around her house with her laptop, and always accesses the Internet through the same access point. Is the user mobile from a network standpoint? Explain!
Answer: No. User usually stays connected to a single access point while roaming around the house.

g. What is the difference between end-to-end delay and packet jitter? What are the causes of packet jitter?
Answer: End-to-end delay is the time it takes a packet to travel across the network from source to destination. Delay jitter is the fluctuation of end-to-end delay from packet to the next packet.

h. Suppose an application generates chunks of 40 bytes of data every 20 msec. and each chunk gets encapsulated in a TCP segment and then an IP datagram. What percentage of each datagram will be overhead, and what percentage will be application data? (Neither IP nor TCP use any of their option header fields).
Answer: 50% overhead. IP -> 40 bytes, TCP -> 40 bytes.

Problem 2: TCP (24 Points, 20 minutes)

For this problem you should familiarize yourself with Figure 1 first. Now assume that in the network shown in Figure 1 two parallel TCP transmissions are performed. *TCP1* is a transmission between Source A and Sink A that uses *TCP Tahoe*. *TCP2* is a transmission between Source B and Sink B that uses *TCP Reno.* Initial *ssthresh* for both TCP transmissions is set to 32. In this specific scenario no additional delay through forwarding is introduced. Thus, the RTT is only composed of the sums of the delay indicated on each link, times two.

a. (4 Points) For the *TCP 1* transmission, draw the resulting congestion window, assuming that a packet loss (triple duplicate ACKs) is detected at time t=900ms in Figure 2.
   Answer: see Figure below
b. (4 Points) For the *TCP 2* transmission, draw the resulting congestion window, assuming that a packet loss (triple duplicate ACKs) is detected at time t=650ms in Figure 2.
   Answer: see Figure below
c. (4 Points) Describe the benefit of TCP Reno over TCP Tahoe.
   Answer: higher throughput since Reno doesn't reset ccwnd to 1 in case of packet loss and doesn't enter slow start phase.
d. (4 Points) In general, explain the purpose of the receiver-advertised window in TCP.
   Answer: It allows the receiver to signal the sender how much unacknowledged data can be in flight. Flow-control
e. (4 Points) Now TCP 1 is closed. What is the new throughput achieved by the remaining session? Briefly describe how this new throughput is reached. (Assume that the link between routers 2 and 2 is the bottleneck link with capacity R.)
   Answer: R since there is no more competing traffic on the bottleneck link. Source increases congestion window size until outgoing queue at router 1 fill up and a packet is dropped. Missing ACK will lead to reduction in congestion window. Sawtooth shape throughput around max. rate R.
f. (4 Points) Now suppose that Source B starts a second TCP connection (TCP 3) to Sink B. What would be the throughput for this new TCP session and TCP 2? How would that change if the sink for that TCP 3 would be Sink A?
   Answer: Both sessions chair the same bottleneck link and the RTT is identical: R/2. Due to the higher RTT to Sink A the throughout for the new session is lower.

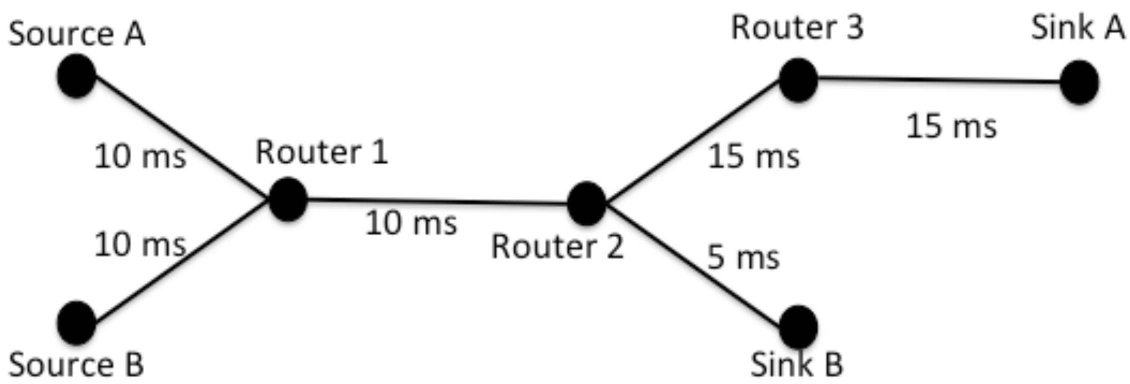

Figure 1 Network layout for problem 5.

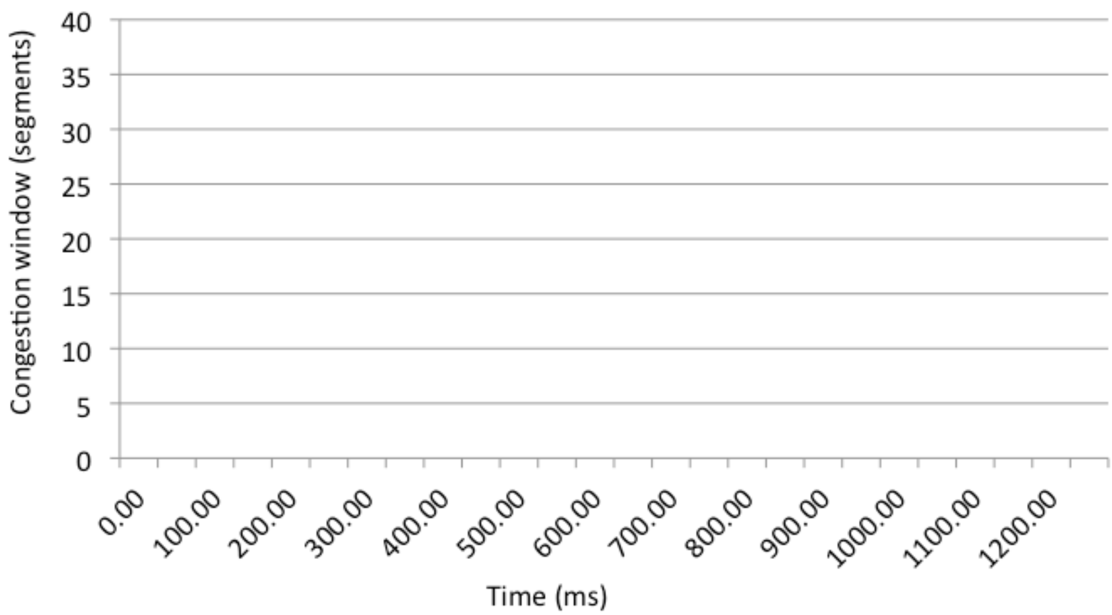

Figure 2 Solution

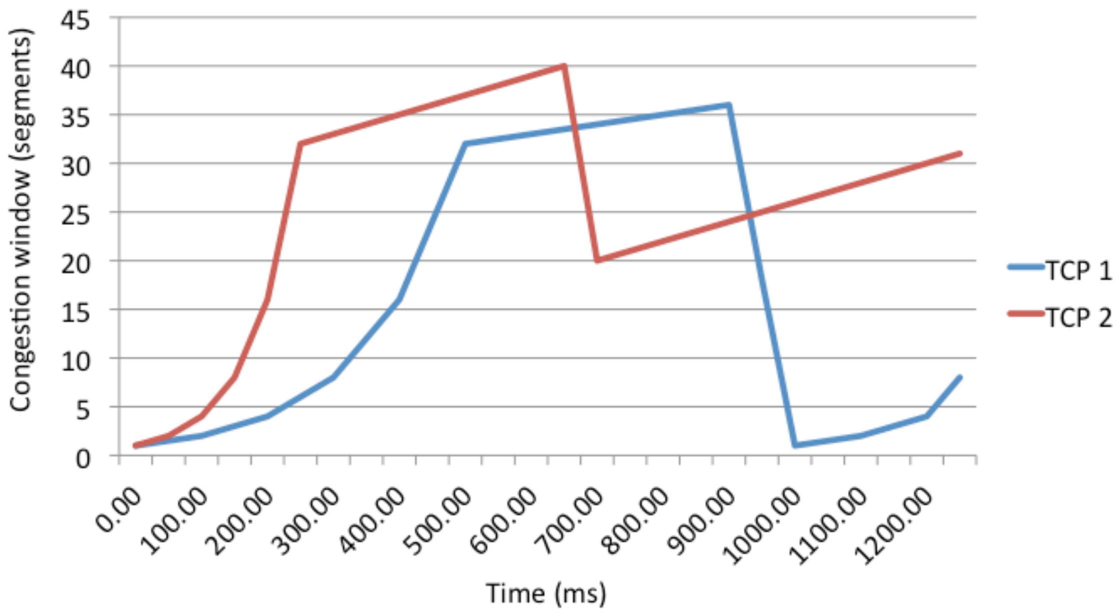
Congestion window (segments)
45
40
35
30
25
20
15
10
5
0
0.00
100.00
200.00
300.00
400.00
500.00
600.00
700.00
800.00
900.00
1000.00
1100.00
1200.00
Time (ms)
TCP 1
TCP 2

Problem 3: Multimedia Potpourri (20 Points, 15 minutes)

a. (4 Points) Suppose an analog audio signal is sampled 16,000 times per second, and each sample is quantized into one of 1024 levels. What would be the resulting bit rate of the PCM digital audio signal?
   Answer: Quantizing a sample into 1024 levels means 10 bits per sample. The resulting rate of the PCM digital audio signal is 160 Kbps.
b. (2 Points) Why is a packet that is received after its scheduled playout time considered lost?
   Answer: A packet that arrives after its scheduled play out time cannot be played out. Therefore, from the perspective of the application, the packet has been lost.
c. (8 Points) Consider the figure below. A sender begins sending packetized audio periodically at *t=1*. The first packet arrives at *t=8* at the receiver.
   1. What are the delays (from sender to receiver, ignoring any playout delays) of packets 2 through 8? Note that each vertical and horizontal line segment in the figure has length of 1,2, or 3 time units.
      Answer: The delay of packet 2 is 8 slots. The delay of packet 3 is 8 slots. The delay of packet 4 is 8 slots. The delay of packet 5 is 10 slots. The delay of packet 6 is 9 slots. The delay of packet 7 is 8 slots. The delay of packet 8 is 8 slots.
   2. If audio playout begins as soon as the first packet arrives at the receiver at *t=8*, which of the first eight packets sent will *not* arrive in time for playout?
      Answer: Packets 2, 3, 4, 5, 6, 7, and 8 will not be received in time for their playout if playout begins at t-8.
   3. If audio playout begins at *t=9*, which of the first eight packets sent will *not* arrive in time for playout?
      Answer: Packets 3, 5 and 6 will not be received in time for their playout if playout begins at t=9.
   4. What is the minimum playout delay at the receiver that results in all of the first eight packets arriving in time for their playout?
      Answer: No packets will arrive after their playout time if playout time begins at t=10.

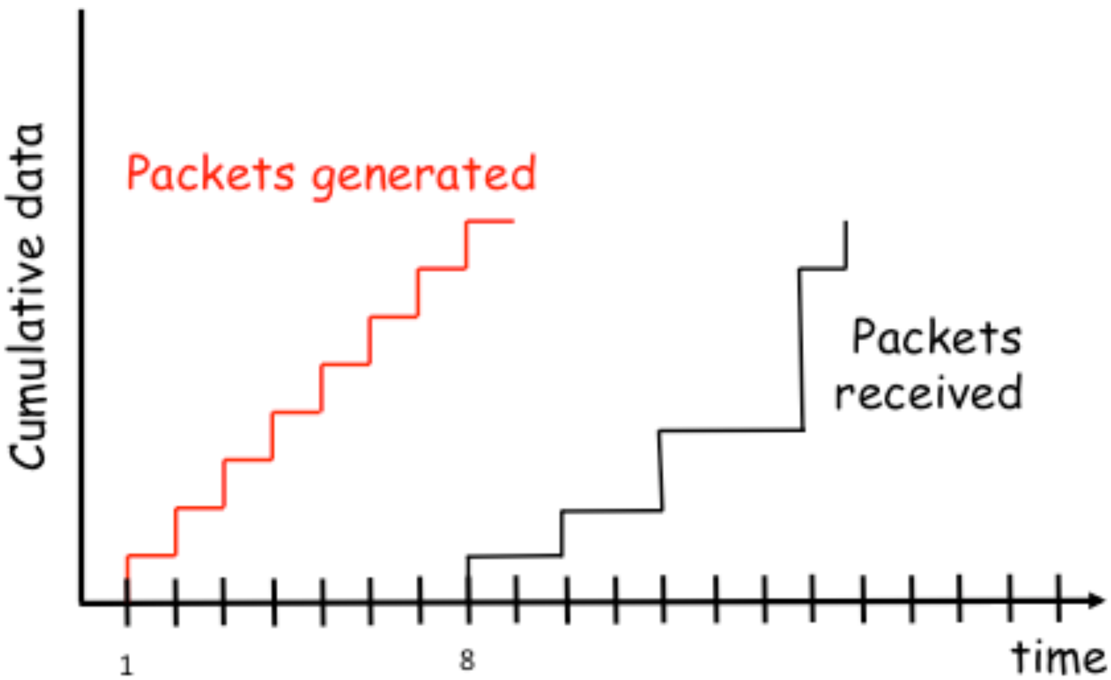

d. (3 Points) Consider a video streaming systems for which there are *N* video versions (at *N* different rates and qualities) and *N* audio versions (at *N* different rates and versions). Suppose we want to allow the player to choose at any time any of the N video versions and any of the *N* audio versions.

  a. If we create files so that the audio is mixed in with the video, so the server sends only one media stream at a given time, how many files will the server need to store (each been accessed via a different URL)?
  Answer: $N*N = N^2$

  b. If the server instead sends the audio and video stream separately and has the client synchronize the streams, how many files will the server need to store?
  Answer: $N+N = 2N$

e. (3 Points) Suppose we send into the Internet two IP datagrams, each carrying a different UDP segment. The first datagram has source IP address A1, destination IP address B, source port P1, and destination port T. The second datagram has source IP address A2, destination IP address B, source port P2, and destination port T. Suppose A1 and A2 are different and that P1 is different from P2. Assuming that both datagrams reach their final destination, will the two UDP datagrams be received by the same socket? Why or why not?
Answer: As discussed in Chapter 2, UDP sockets are identified by the two-tuple consisting of destination IP address and destination port number. So the two packets will indeed pass through the same socket.

Problem 4: LAN Potpourri (24 Points, 20 minutes)
In this problem we are looking in the "day in the life of a web request" scenario. Please refer to Figure 3 for this problem.

a) (8 Points) For the initial step we assume that the connecting laptop needs to get its own IP address, the address of the first hop router, and the address of the DNSserver. Describe in detail, how this is achieved by using the DHCP protocol.
b) (8 Points) Before the client can send out an HTTP request to www.google.com it has to obtain the IP address for that hostname via DNS. Describe in detail, how the DNS request is performed. Keep in mind the client does not yet no the MAC address of the first hop router interface!
c) (8 Points) After the client has retrieved the IP address, it can finally send out the HTTP request. First describe how the TCP connection is setup and then describe how the HTTP request is being performed.

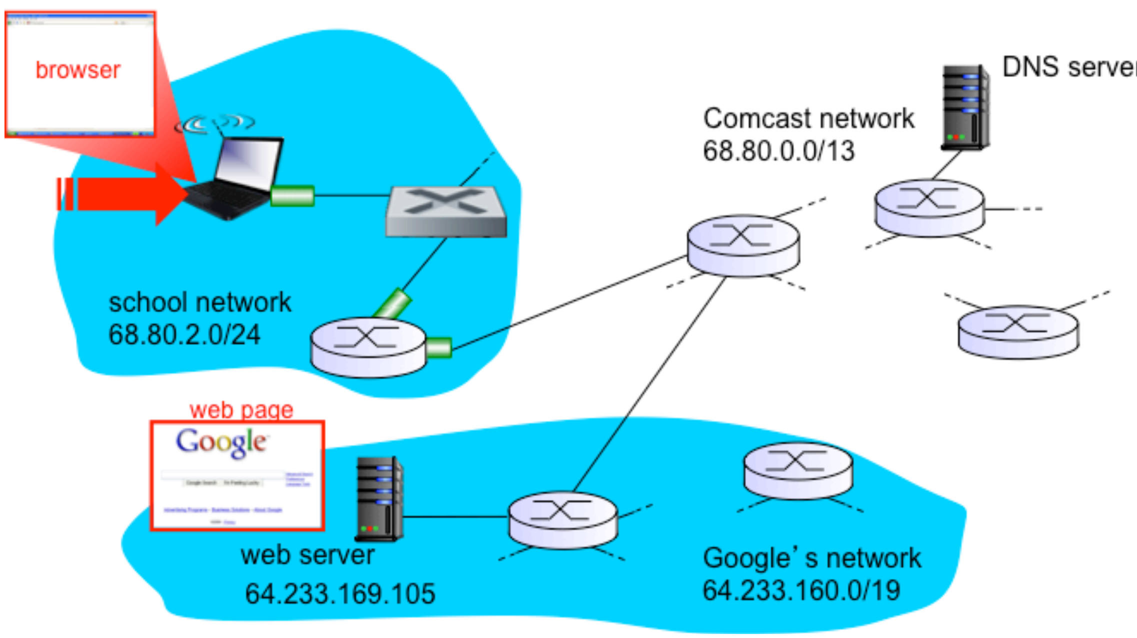

Figure 3

Solution:
(The following description is short, but contains all major key steps and key protocols involved.)

Your computer first uses DHCP to obtain an IP address. You computer first creates a special IP datagram destined to 255.255.255.255 in the DHCP server discovery step, *and puts it in a Ethernet frame and broadcast it in the Ethernet. Then following the steps in the DHCP protocol, you computer is able to get an IP address with a given lease time.*

A DHCP server on the Ethernet also gives your computer a list of IP addresses of first-hop routers, the subnet mask of the subnet where your computer resides, and the addresses of local DNS servers (if they exist).

Since your computer's ARP cache is initially empty, your computer will use ARP protocol to get the MAC addresses of the first-hop router and the local DNS server.

Your computer first will get the IP address of the Web page you would like to download. If the local DNS server does not have the IP address, then your computer will use DNS protocol to find the IP address of the Web page.

Once your computer has the IP address of the Web page, then it will send out the HTTP request via the first-hop router if the Web page does not reside in a local Web server. The HTTP request message will be segmented and encapsulated into TCP packets, and then further encapsulated into IP packets, and finally encapsulated into Ethernet frames. Your computer sends the Ethernet frames destined to the first-hop router. Once the router receives the frames, it passes them up into IP layer, checks its routing table, and then sends the packets to the right interface out of all of its interfaces.

Then your IP packets will be routed through the Internet until they reach the Web server.

The server hosting the Web page will send back the Web page to your computer via HTTP response messages. Those messages will be encapsulated into TCP packets and then further into IP packets. Those IP packets follow IP routes and finally reach your first-hop router, and then the router will forward those IP packets to your computer by encapsulating them into Ethernet frames.