First Midterm for ECE374

03/04/13

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. We saw that TCP and UDP provide two very different service models. Suppose that an application wants all of the functionality provided by UDP but only some of the functionality provided by TCP (e.g., the application wants reliable message transfer and flow control, but not congestion control). How would an application get this different service in today's Internet?
Answer: The application would use UDP sockets and implement the desired additional functionality (e.g., reliability and flow control) in the application itself.

b. Explain why you can actually use the *telnet* protocol to perform a simple communication with a web server?
Answer: Telnet is a text-based communication protocol, likewise is HTTP. If a client composes a text message that is HTTP standard compliant the server sends back the expected reply, which is shown in text (including the HTML content at the client.

c. Consider two hosts, A and B, connected by a single link of rate *R* bps. Suppose that the two hosts are separated by *m* meters, and suppose the propagation speed along the link is *s* meters/sec. Host A is to send a packet of size *L* bits to Host B. Determine the transmission time of the packet, d_{trans}, in terms of *L* and *R.*
Answer: $d_{trans} = L/R$ seconds.

d. Suppose that Alice wants to send an email message to Bob. This will involve four entities: Alice's mail client (for email composition and sending), Alice's outgoing mail server, Bob's incoming mail server, and Bob's mail client (for email retrieval and viewing). Between which of these four entities does the SMTP protocol operate? What about the IMAP protocol?
Answer: SMTP runs between Alices mail client and her server, and also (separately) between her server and Bob's server. IMAP runs between Bob's server and his mail client to retrieve messages from Bob's server.

e. A client sends a TCP segment to the sever with Sequence Number 1400 and the payload included in the segment is 1399 bytes long. A) What is the ACK Number in the acknowledgement that is returned from the server? B) Assume this packet is lost but the following packet is received. What is the ACK Number in the acknowledgement that is returned from the server for this packet?
Answer: A)2800, B)1400.

f. Suppose TCP uses AIMD for its congestion control without slow start. Assuming *cwnd* increases by 2 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 4 MSS to 12 MSS (assuming no loss events)?

Answer: It takes 1 RTT to increase CongWin to 6 MSS; 2 RTTs to increase to 8 MSS; 3 RTTs to increase to 10 MSS; 4 RTTs to increase to 12 MSS.

g. Assume two neighboring nodes in Circular DHT. Node A has ID 1100 and node B's ID is 1111. Another peer (neither A nor B) in the ring searches for key 1110. Is either A or B sending a reply message to the requesting node informing that it holds the <key, value> pair?
Answer: B.

h. Cable and DSL are two residential Internet access technologies. Which one can be described as dedicated and which one as shared? Explain why!
Answer: Cable, shared, one coax cable shared between clients and head end. DLS, dedicated, each client is connected by a separate pair of wires to the DSLAM.

Problem 2: Delays, Throughput and Caches (24 Points, 20 minutes)
Consider the scenario shown in Figure 1 in which a server is connected to a router by a 100Mbps link with a 50ms propagation delay. Initially this router is also connected to two routers, each over a 50Mbps link with a 200ms propagation delay. A 1Gbps link connects a host and a cache (if present) to each of these routers and we assume that this link has 0 propagation delay. All packets in the network are 20,000 bits long.

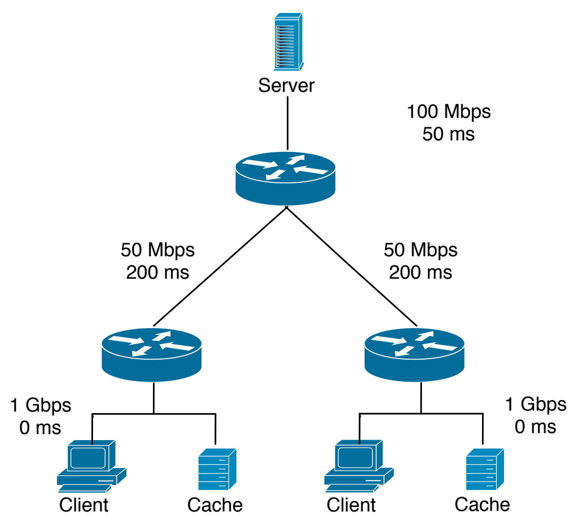

Figure 1

a. (4 Points) What is the end-to-end delay from when a packet is transmitted by the server to when it is received by the client? In this case, we assume there are no caches, there's no queuing delay at the routers, and the packet processing delays at routers and nodes are all 0.
Answer: If all packets are 20,000 bits long it takes 200 usec to send the packet over the 100Mbps link, 400 usec to send over the 50Mbps link, and 20 usec to send over the 1Gbps link. Sum of the three-link transmission is 620 usec. Thus, the total end-to-end delay is 250.62 msec.

b. (4 Points) Here we assume that client hosts send requests for files directly to the server (caches are not used or off in this case). What is the maximum rate at which the server can deliver data to a single client if we assume no other clients are making requests?
Answer: Server can send at the max of the bottleneck link: 50Mbps.

c. (4 Points) Again we assume only one active client but in this case the caches are on and behave like HTTP caches. A client's HTTP GET is always first directed to its local cache. 65% of the requests can be satisfied by the local cache. What is the average rate at which the client can receive data in this case?
Answer: We assume that requests are serially satisfied. 35% of the requests can be delivered at 50Mbps and 65% at 1Gbps. So the average rate is 667.5Mbps.

d. (4 Points) Now clients in both LANs are active and the both caches are on (similar to c). 65% of the requests can be satisfied by the local caches. What is the average rate at which each client can receive data?
Answer: 50Mbps link remains bottleneck link, which is not shared between LANs. Answer is same as in c. Note that we assume that the 100Mbps link is shared at fine granularity, so that each client gets up to 50Mbps over that link.

e. (8 Points) Now consider Figure 2 where the network has been extended by two additional LANs, and all LANs are connected via 25Mbps links and the throughput of the server link is only 80Mbps. In all four LANs 60% of the requests can be satisfied by the local caches. What is the average rate at which each client can receive data?
Answer: Now the server link is the bottleneck link an equally shared by all clients, each getting 20Mbps. 40% of the client requests are served at 20Mbps and 60% of the requests are delivered at 1Gbps, resulting in an average rate of 608Mbps. Only a reduction by 2Mbps on average despite a 20% reduction in server link bandwidth. This shows how efficient caching can be!

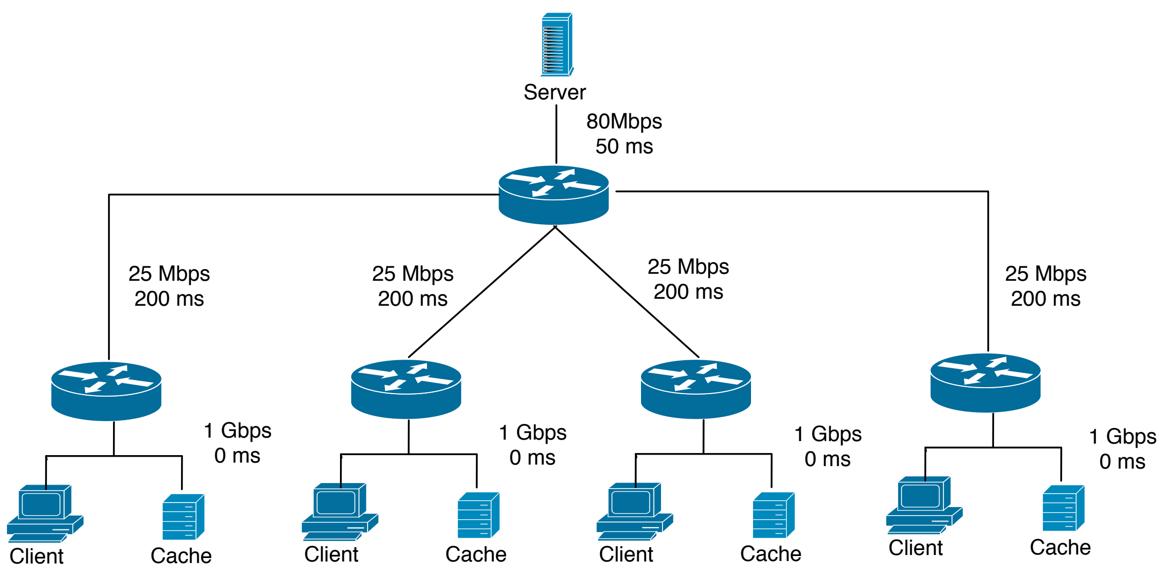

Figure 2

Problem 3: Putting it all together (20 Points, 15 minutes)
For this problem you should familiarize yourself with Figure 3 first. Initially, assume that the client wants to retrieve the www.cnn.com home page but has no information about the www.cnn.com web server IP address.

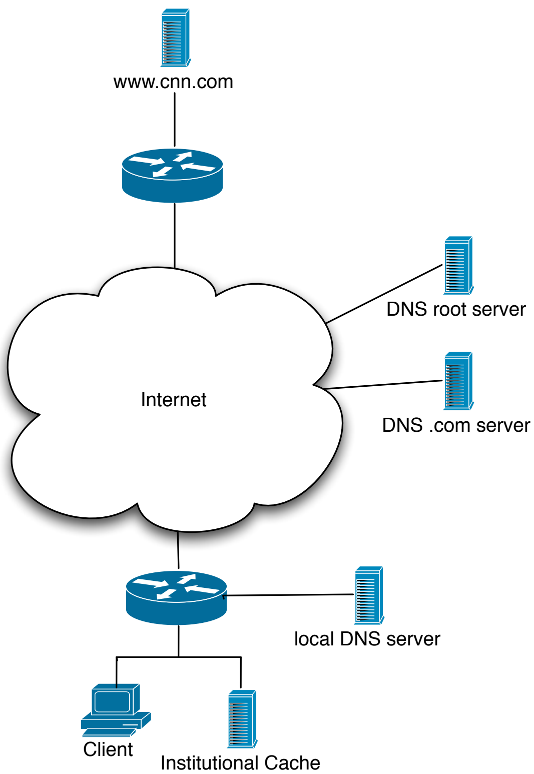

Figure 3

a. (4 Points) Describe the process of the client obtaining the IP address for the hostname www.cnn.com under the assumption that it is not cached at the local DNS server and that the local DNS server has not cached an entry for the .com DNS server. (Describe this for the recursive case!)
Answer: Client contacts local DNS, local DNS contacts root DNS, root DNS contacts TLD DNS, TLD DNS contacts authoritative DNS, IP for www.cnn.com is propagated back on the same rout.
b. (4 Points) We know that www.cnn.com is a very popular web site and the many client requests cannot be handled by a single server but rather by a cluster of web servers (each having a different IP address). Describe the process that DNS offers for load balancing.
Answer: Server responds with an entire set of IP addresses for canonical name but rotates ordering within each reply. Client typically sends request to first in list.
c. (4 Points) Give an example for the source and destination port numbers in a TCP segment sent from the client to the www.cnn.com web server. Now assume a second browser is opened on the client which also wants to retrieve the www.cnn.com start page. What are source and destination port for a TCP packet that belongs to this connection?

Answer: 1.) src port: XXXX (greater 5000), dst port 80, 2.) src port: YYYY (greater 5000 and different than XXXX), dst port 80.

d. (4 Points) Let us turn our attention to the web server for this sub-problem. Assume that the main web page is a single HTML file (e.g., index.html) that is of size 350,000 bytes and the MSS is 1400 bytes. Further assume that the sequence number for the very first TCP segment carrying that data stream is 0. What is the sequence number for the 2^{nd}, 3^{rd}, and 4^{th} segment? Also give an expression for the n^{th} TCP segment in that stream. Assuming there are no packet losses and timeouts, how many TCP segments are needed in total to transmit that HTML page?
*Answer: 1400, 2800, 4200, 5600, (n-1)*1400, 350,000/1400=250.*

e. (4 Points) Now assume there's an institutional cache in the client's subnet. The client's initial request is cached by this cache. Since news pages change frequently the client wants to make sure that it does not get served an outdated HTML page from the cache. Explain the HTTP mechanism that prevents this from happening. What would be in the body of the second server reply if the reply would be `304 Not Modified?`
Answer: Last-Modified from server, If-Modified-since from client. The body would be empty in the second case since the HTML page would be served from the cache.

Problem 4: Sliding Window Protocols (24 Points, 20 minutes)

a. (4 Points) Consider the sliding window protocol in Figure 4 to the right. Does this figure indicate that Go-Back-N is being used, Selective Repeat is being used, or there is not enough information to tell? Explain your answer briefly. *Answer: there is not enough information to tell, since both GBN and SR will individually ACK each of the first two messages as they are received correctly.*

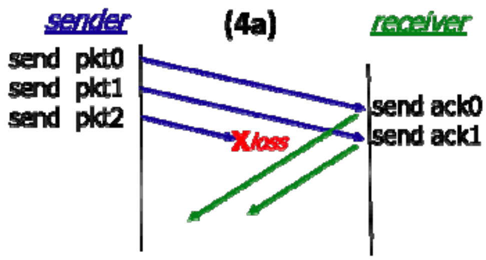

Figure 4

b. (4 Points) Consider the sliding window protocol in Figure 5. Does this figure indicate that Go-Back-N is being used, Selective Repeat is being used, or there is not enough information to tell? Explain your answer briefly. *Answer: This must be the SR protocol since pkt 3 is acked even though pkt 2 was lost. GBN uses cumulative ACKs and so would not generate an ACK 3 if pkt 2 was missing.*

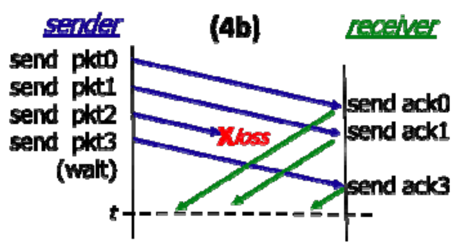

Figure 5

c. (4 Points) Consider Figure 5 again. Suppose the sender and receiver windows are of size N = 5 and suppose the sequence number space goes from 0 to 15. Show the position of the sender and receiver windows over this sequence number space at time t (the horizontal dashed line).
Answer:

Sender: [0 1 2 3 4] 5 6 7 8 9 ....

Receiver 0 1 [2 3 4 5 6] 7 8 9 ....

d. (6 Points) Suppose that it take 1 ms to send a packet, with a 10 ms one-way propagation delay between the sender and receiver. The sliding windows size is again N = 4. What is the channel/link utilization?
Answer: the utilization is 4/(1+20) or 0.19.

e. (6 Points) Assume a TCP sender transmits 4 TCP segments with respective sequence numbers 1200, 2400, 3600, 4800, 6000. The sender receives four acknowledgements with the following sequence numbers, 2400, 2400, 2400, 2400, 7200. Complete Figure 6 to show what TCP segments are exchanged between sender and receiver.

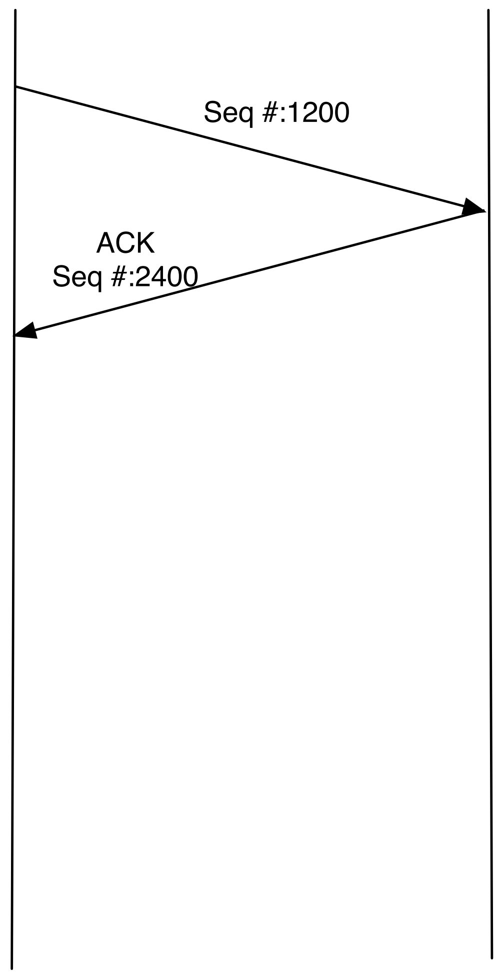

Figure 6

Solution for e:

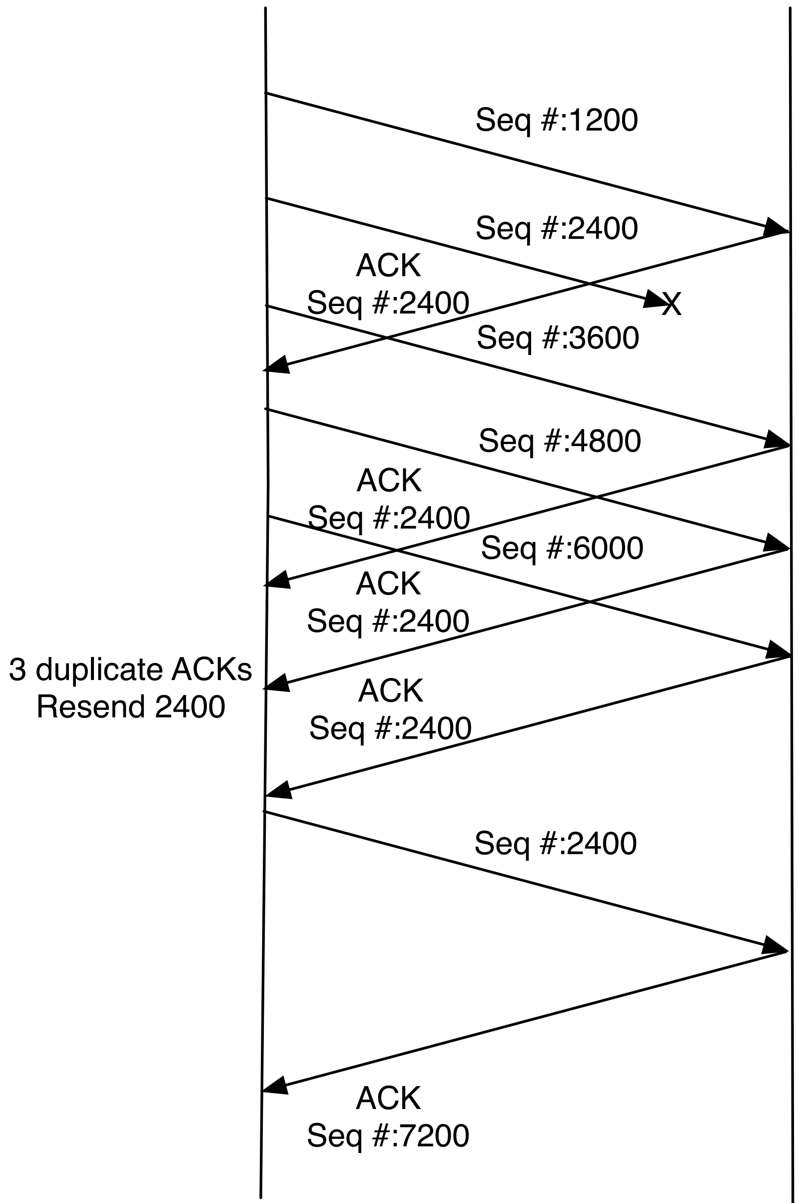
Seq #:1200
Seq #:2400
ACK
Seq #:2400
X
Seq #:3600
Seq #:4800
ACK
Seq #:2400
Seq #:6000
ACK
Seq #:2400
3 duplicate ACKs
Resend 2400
ACK
Seq #:2400
Seq #:2400
ACK
Seq #:7200