First Midterm for ECE374

02/25/15

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

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

a. Briefly describe why the actual bandwidth on your cable access network connection on, e.g., a Saturday evening, could be significantly lower than what has been advertise by the provider (Comcast, Charter, etc.)? Would you see the same with a DSL-based provider?
Answer: Cable is shared and the chances are high that many customers use the Internet from the homes in the evening (especially on weekends). The sharing will lead to reduced BW for the individual customer. DSL is not shared and that's why this behavior would not been seen in this case.

b. Server A runs a web server on port 80 and a SMTP server on port 25. Client 1 opens local port 5100 to establish a TCP session to the web server on Server A. What are the numbers of the destination and source ports in the TCP segments send to the server? Client 2 opens local port 6100 to establish a TCP session to the SMTP server on Server A. What are the numbers of the destination and source ports in the TCP segments send to the server?
Answer: Client 1 D=80, S=5100, Client 2 D=25, S=6100

c. The Internet is a network of networks. How many Tier levels to classify these individual networks exist? At what level would your residential access network reside? How likely do you think it is that traffic that is transmitted from a host in the US to a host in Europe will go through a Tier-1 network? How are the points called at which Tier-1 network providers and large content distributes peer?
Answer: 3, 3, 100%, Internet Exchange points.

d. A caravan of 10 cars "propagates" at 100 km/hr on a 100 km long toll road. At the toll booth it takes 12 seconds to service each individual car. How long does it take for the last car to arrive at the end of the toll road? Cars now "propagate" at 1000 km/hr. Will cars arrive at the end of the toll road before all cars are serviced at first booth? If so, at what time will that happen?
A:120 seconds + 60 minutes = 62 minutes, yes after 7 minutes.

e. Draw a sketch that shows all layers of the Internet protocol stack in correct hierarchy!
A: Application, Transport, Network, Link, Physical.

f. Suppose that we want to change the IP address of www.ecs.umass.edu from 128.119.91.47 to 128.119.91.50 and change this mapping in the DNS authoritative name server for www.ecs.umass.edu. Once this mapping is changed in the authoritative name server, will all future references (generated anywhere in the Internet) to www.ecs.umass.edu then be sent to 128.119.91.50? Explain briefly (in two or three sentences).
Answer: Local DNS caches throughout the Internet will not time out the old mapping of www.ecs.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 www.ecs.umass.edu and hence would not learn the new mapping.

g. 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.

h. Assume that in the case of a pipelined, reliable transport protocol, for which the sender can have up to *N=4* unacknowleged packets, packets with sequence numbers 0-3 are transmitted. Only the packet with sequence number 2 gets lost during the transmission. What are the sequence numbers of the packets that are retransmitted in the case of Go-Back-N? What are they in the case of Selective Repeat?
Answer: Go-Back-N: 2, 3; Selective Repeat: 2.

Problem 2: Delays, Throughput (24 Points, 20 minutes)

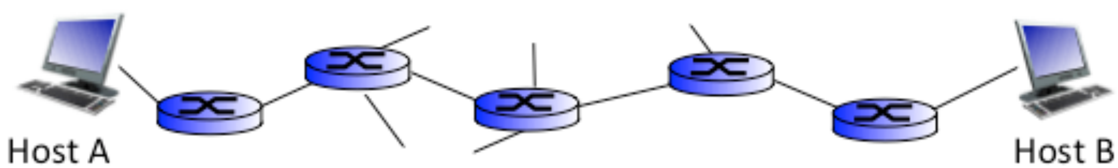

Figure 1

Consider the network shown in Figure 1.

a. (4 Points) Based on this figure explain how the traceroute application determines the round-trip-time (RTT) between your host, the individual routers on the path to the receiving host, and the receiving host?
A: sends three packets that will reach router i on path towards destination, router i will return packets to sender, sender times interval between transmission and reply.

b. (2 Points) The following line represents a typical single output line from traceroute:
`ar01.needham.ma.boston.comcast.net(68.87.147.241) 14.682ms 13.927ms  19.944ms`
Why do the values for the three RTTs shown at the end of that line change?
Answer: RTTs highly depend on the queue length of the routers located between the sender and the receiver. Based on other traffic there can be more or less waiting packets in a routers outgoing queue at a certain instance in time. Thus, the RTT between individual pings can vary.

c. (4 Points) Now assume a static RTT of 136 milliseconds between the two hosts. How long will it take in this case of an HTTP GET request to receive the first bit of the HTTP response from Host B? How long will it take to transmit

a 50,000 bit HTTP object from Host B to Host A if the bottleneck bandwidth on the path between the two hosts is 5 Mbit/s?
A: 2 RTT = 272 milliseconds, 10ms

d. (4 Points) Name the four sources of packet delay? Which of the four is the most significant one?
A: $d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$, d_{trans}

Now, consider the scenario shown in Figure 2 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.

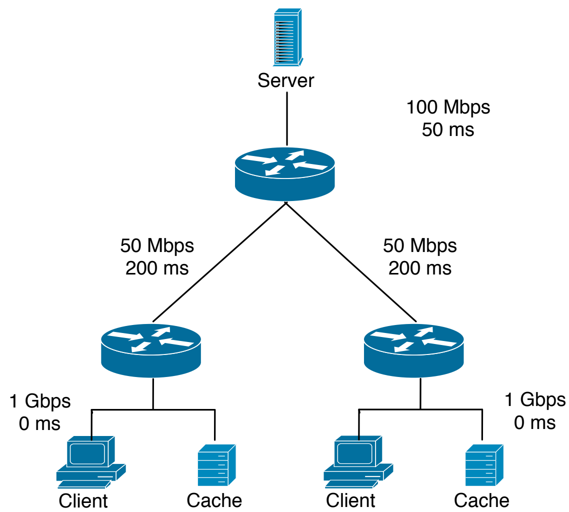

Figure 2

e. (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.

f. (2 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.

Problem 3: Putting it all together (24 Points, 20 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. All links have 10Mbps capacity and we assume there is no queuing delay at the routers, and the packet processing delays at routers and nodes are all 0.

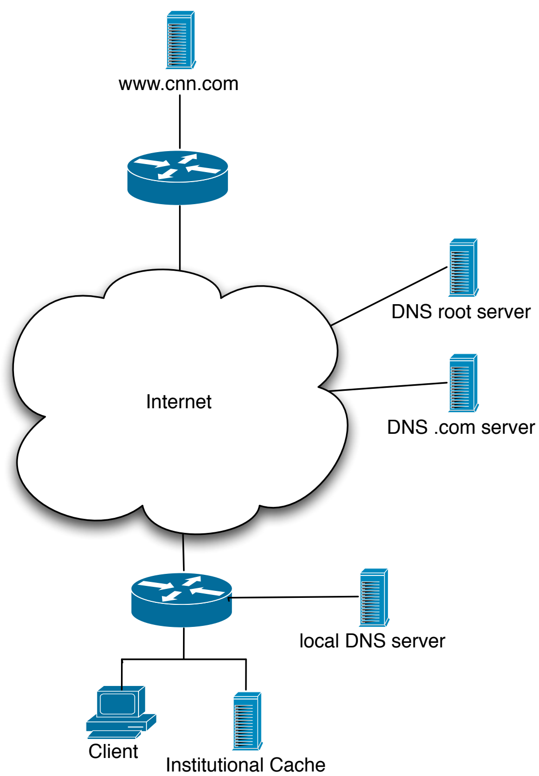

Figure 3

a. (6 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 BUT that the .com TLD DNS server has cached an entry for the .com DNS server. The local DNS server also does NOT have cached the IP address of the .com TLD DNS. (Describe this for the iterated case!)
Answer: Client contacts local DNS, local DNS contacts root DNS, root DNS replies to local DNS, local DNS contacts authoritative DNS, authoritative DNS replies to local DNS, local DNS send IP for www.cnn.com back to client.

b. (2 Points) After that first client, a second client (connect to the same network as the first client) also wants to obtain the IP address for www.cnn.com. Describe the process of the client obtaining the IP address in this case.
Answer: Client contacts local DNS, local DNS has entry cached and returns it to client.

c. (2 Points) Assume that the round trip time between local DNS server and DNS root server is 3RTT, between local DNS server and DNS TLD server is 2RTT, and between the clients and the local DNS server is RTT. How long does it take for the first client to obtain the IP address for www.cnn.com? How long for the second?
Answer: 6RTT, RTT.

d. (4 Points) The CNN main page is 500Kbits in size and refers to 7 additional web object (e.g., images) that all also have a size of 1250Kbits. Assume that the round trip time between client and server is 3*RTT*. Assuming non-persistent HTTP with no parallel TCP connections, how long does it take to download all web objects to the client?
*Answer: (6RTT + 1250Kbits/10Mbps)*8 = 48RTT + 10Mbits/10Mbps = 48RTT + 1s*

e. (4 Points) How long will it take to download all objects in the case of persistent HTTP?
Answer: (6RTT + 1250Kbits/10Mbits) + 7(3RTT + 1250Kbits/10Mbits) = 27RTT + 1s*

f. (2 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.

g. (4 Points) Now let us assume that independent of the URL, 30% of the web requests that originate from the local network can be satisfied by the local cache and 70% of the requests are satisfied at the origin (in this specific case we assume they are all satisfied by www.cnn.com). The bandwidth between client and cache is 100Mbps in this case. What is the total average delay with which web objects are served to clients in the local network?
Answer: We assume that requests are serially satisfied. 30% of the requests can be delivered at 100Mbps and 70% at 10Mbps. So the average rate is 37Mbps.

Problem 4: Transport Protocols (24 Points, 25 minutes)

a. (14 Points) Consider a scenario in which Host A wants to simultaneously send packets to Hosts B and C via a broadcast channel - a packet sent by A is carried by the channel to both B and C. Suppose that the broadcast channel connecting A, B, and C can independently lose and corrupt packets (for example, a packet sent from A might be correctly received by B, but not by C). Design a stop-and-wait like error control protocol for reliably transferring packets from A to B and C, such that A will not get new data from the upper layer until it knows that both B and C have correctly received the current packet. Give FSM descriptions of A and C. (Hint: The FSM for B should be essentially the same as for C.) Also, give a description of the packet format(s) used.

Answer:

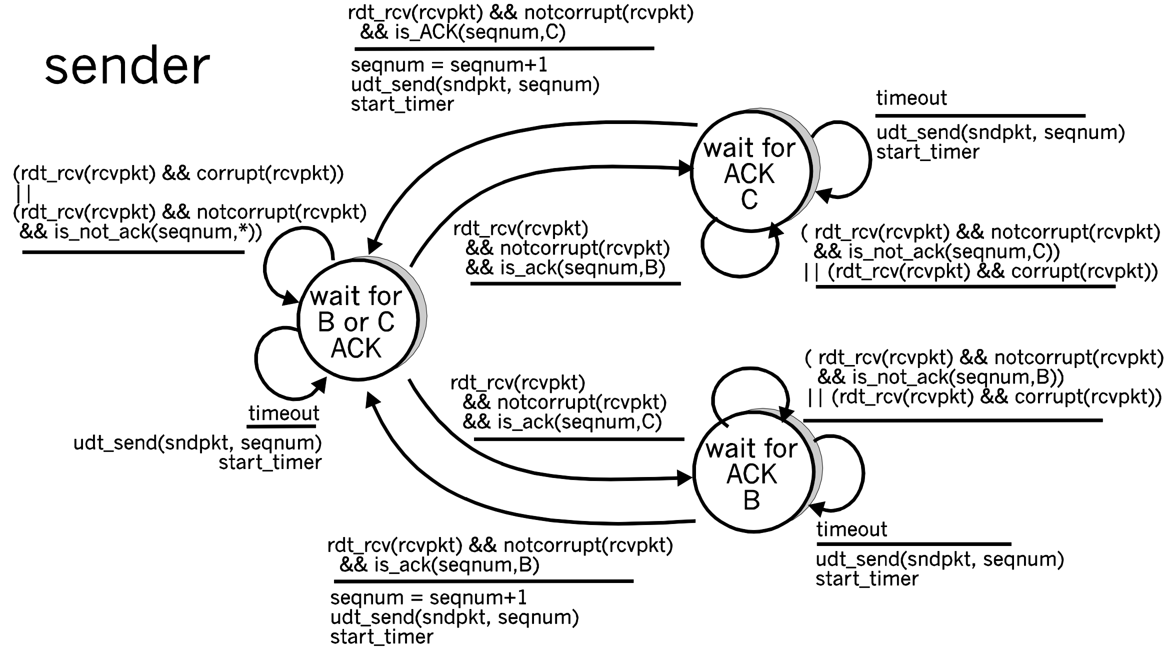

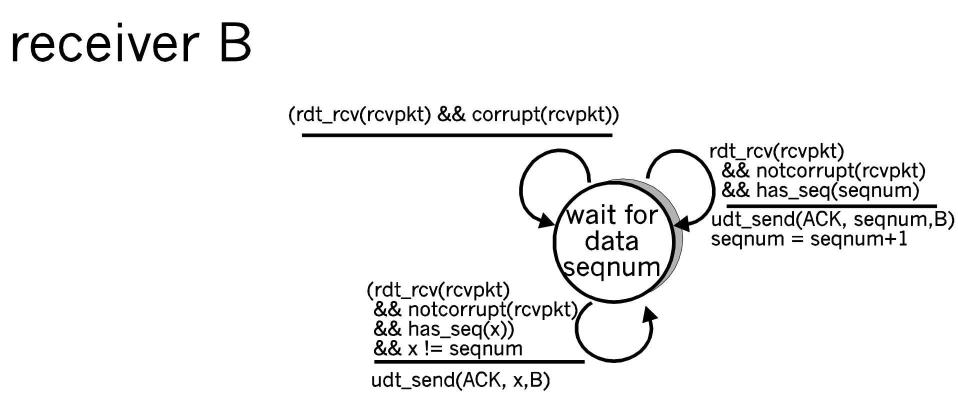

b. (4 Points) rdt3.0 in action! For this problem assume the reliable, error-correcting stop and wait protocol rdt3.0. Figure 4 shows the initial communication between sender and receiver. In this example, the packet with sequence number 1 is lost on the way to the receiver. Complete Figure 4 to illustrate how rdt3.0 recovers from that packet loss by illustrating the (error-free) transmission of 2 additional packets.

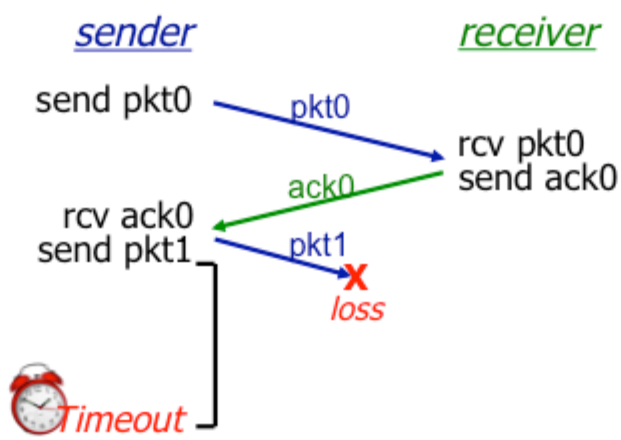

Figure 4

Answer:

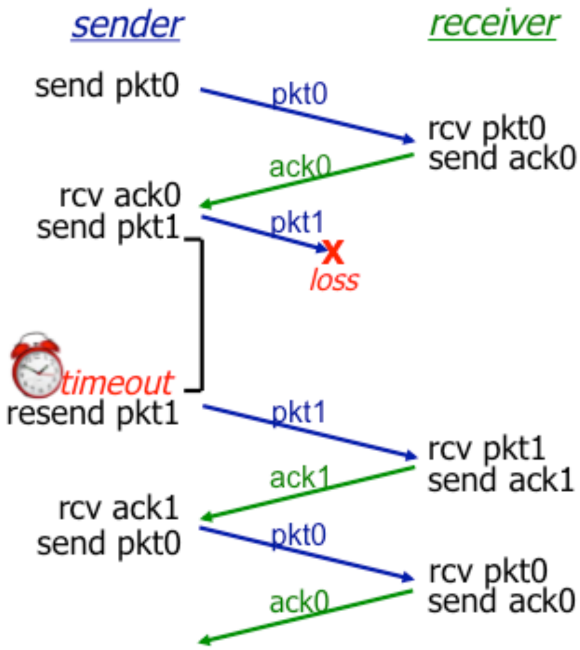

c. (4 Points) In the rdt 3.0 protocol the ACKs flowing from the receiver to the sender do NOT have sequence numbers (although they do have an ACK field that contains the sequence number of the packet they are acknowledging). Why is it that the ACKs do not require sequence numbers? (Hint: Think about the case of duplicated ACKs.)
Answer:
To best answer this question, consider why we needed sequence numbers in the first place. We saw that the sender needs sequence numbers so that the receiver can tell if a data packet is a duplicate of an already received data packet. In the case of ACKs, the sender does not need this info (i.e., a sequence number on an ACK) to tell detect a duplicate ACK. A duplicate ACK is obvious to the rdt3.0 receiver, since when it has received the original ACK it transitioned to the next state. The duplicate ACK is not the ACK that the sender needs and hence is ignored by the rdt3.0 sender.

d. (2 Points) Figure 5 shows the trace operation of a pipelined protocol. Does it show the trace of Go-Back-N or Selective Repeat?

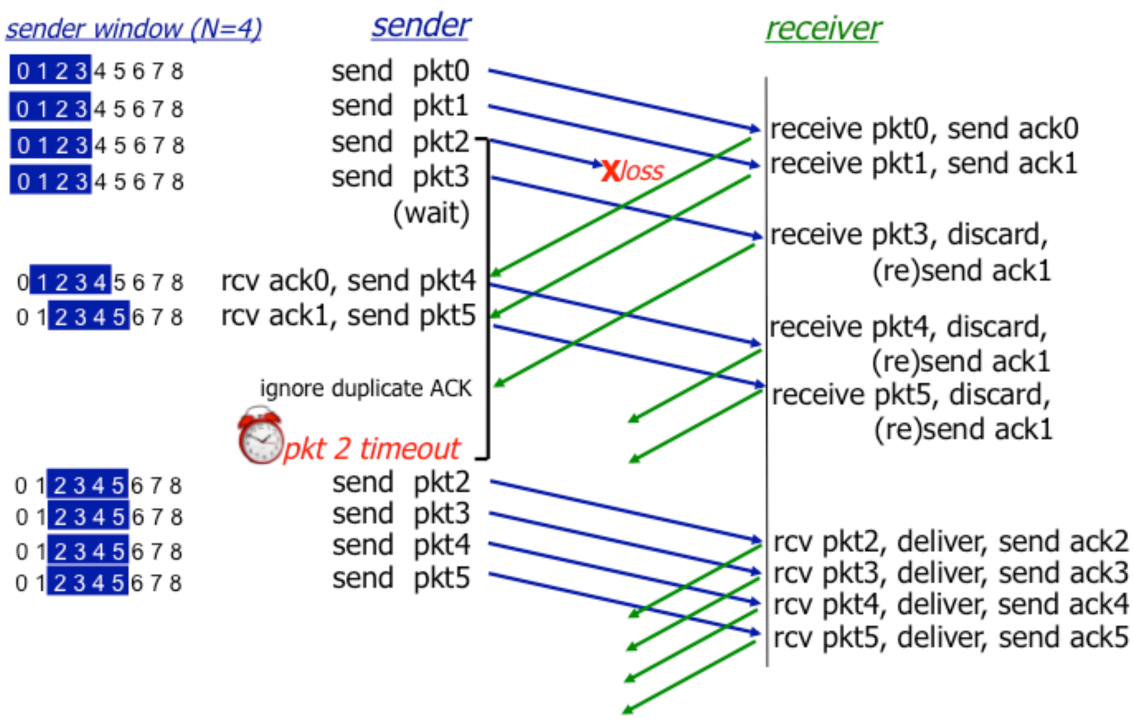

Figure 5

Answer: Go-Back-N