CS 421: Computer Networks **FALL 2011**

MIDTERM
November 30, 2011
120 minutes

Name: _____

Student No: _____

Show all your work very clearly. Partial credits will only be given if you carefully state your answer with a reasonable justification.

Q1	
Q2	
Q3	
TOT	

1)

a) (5 pts) A study on the traffic at root DNS name servers shows that 95% of the queries they receive are for non-existent top level domains, e.g., typos such as bilkent.edu.ty **(**.ty domain does not exist**)** instead of bilkent.edu.tr. On the other hand, a study of a representative cross-section of local DNS servers shows that only 0.1% of the queries they receive are for non-existent top level domains. How can you explain this discrepancy?

b) (5 pts) Give an advantage of using hierarchical overlays in peer-to-peer applications, where some nodes act as super nodes or cluster heads (such as in Kazaa or Skype) compared with using a fully distributed architecture (such as in Gnutella).

c) (5 pts) Give two reasons why interactive applications, such as Internet telephony, prefer using UDP over TCP.

d) (5 pts) Suppose that you are asked to implement a reliable data transfer mechanism for a connection that has a very small packet loss/error rate. Which one of the Go-Back-N and Selective Repeat mechanisms will you choose? Justify your answer.

e) (5 pts) There are two loss events in TCP: timeout and fast retransmission. Which of these two loss events indicate a more severe state of congestion in the network? Justify your answer.

f) Answer the following questions regarding Project 1:

i) (2 pts) Did you use UDP or TCP sockets to talk to SMTP server? Why?

ii) (2 pts) Did you use UDP or TCP sockets to talk to POP3 server? Why?

iii) (1 pts) Which command of POP3 is used by your POP3 client to learn about the messages in your mailbox?

iv) (1 pts) Which command of POP3 did you used to download a message from the POP3 server and see it on your screen?

v) (2 pts) What kind of encoding does SMTP protocol use? Would you use the same encoding today as well, if you would be designing SMTP today? Why or why not?

2)

a) (8 pts) Assume that there are 3 links on a path connecting hosts A and B passing through routers R1 and R2 as shown in the following figure. Each link has a distance of 200 km and the transmission rate of each link is shown in the figure. We are transmitting a file composed of **six packets** from node A to node B using datagram packet switching. Each packet has a length of 1500 Bytes including all headers. Assume that the processing and queuing delays in each intermediate node are negligible and the propagation speed is $2x10^5$ km/s. Calculate the total delay incurred in transferring the file from node A to node B.

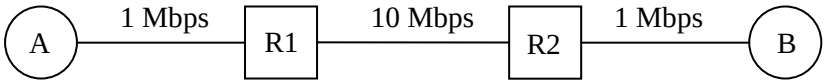

b) (12 pts) Consider a connection with a 10 msec end-to-end, i.e., 20 msec roundtrip, delay (including all delays incurred within the network, but excluding the packet transmission time of the sender). We want to transfer a file composed of 16 segments (with sequence numbers from 1 to 16), where each segment has a transmission time of 1 msec. Assume that ACK segments have negligibly small size and there is no processing delay at the receiver. Assume also that the processing delay at the sender after an ACK is received is negligible. We assume that the communication between the sender and receiver is fully duplex, i.e., sender can send data segments while receiving an ACK segment. **Selective Repeat** protocol is used with a window size of N = 10 segments. Assume that all data segments are received correctly while the **first transmissions** of the **ACK packets** with acknowledgment numbers **4 and 11** are errored, whereas all other **ACK packets are fully reliable**. The timeout for each data segment is set to 30 msec **starting from the end of the transmission of the segment**. How much time is required to complete the transfer of the whole file and receive the final ACK at the sender?

c) (12 pts) Answer the above Question b) when **Go-Back-N** protocol is used with a window size of N = 10 segments. Each window of the sender has a timeout of 30 msec starting from the time when the window is set by the sender.

3)

a) (8 pts) Suppose two applications running at Hosts A and B are communicating using a TCP connection. Assume by time t, Host B has already received all the bytes up to and including byte 500 from Host A and the application running on Host B has read all these bytes from the TCP (socket) buffer at Host B (i.e., all bytes [1,500] are received and delivered to the application). After time t, Host A is sending 3 TCP segments back-to-back to Host B. Segments 1, 2 and 3 each have 100, 120 and 80 bytes of data, respectively, excluding the headers. Assume, segment 3 arrives first to Host B, then segment 1, and then segment 2. After each segment is received, Host B sends an ACK to Host A immediately, i.e., sending an ACK is not delayed. Assume there are no packet/ACK losses and no timeouts happened. Assume the size of TCP receive-buffer at Host B for this connection is 2000 bytes. Assume, after time t, the application at Host B does not read data from the receive buffer until all ACKs corresponding to these 3 segments are sent. What are the acknowledgement and receive-window field values in the three ACK packets sent from Host B to Host A?

b) (6 pts) Consider the RTT estimation algorithm for setting the retransmission TimeOut used by TCP as we discussed in the class:

$$\text{EstimatedRTT}(k+1) \leftarrow (1 - \alpha)\ \text{EstimatedRTT}(k) + \alpha\ \text{SampleRTT}(k+1)$$

for $k = 1,2,\ldots$, where $\text{EstimatedRTT}(1) = \text{SampleRTT}(1)$. Assume that $\alpha = 1/8$.

The first segment of the TCP connection times out in its first transmission attempt, RTT of this first segment in its second transmission attempt is 12 msec, RTT of the second segment in its first transmission attempt is 8 msec, and RTT of the third segment in its first transmission attempt is 4 msec. What value of EstimatedRTT(2) will be generated by the above algorithm?

c) (9 pts) Assume that the congestion window of a TCP connection was 14 segments long when a timeout occurred. Assume that there are no segments or acknowledgments of this connection that were in transit when the timeout occurred. The connection has a 20 msec end-to-end, i.e., 40 msec roundtrip, delay (including all delays incurred within the network, but excluding the packet transmission time of the sender). The transmission time for a segment is 5 msec. The receive window is fixed at 100 segments for the entire duration of the connection. How long will it take for the connection to reach the Congestion Avoidance phase after the timeout, assuming that no further segments are lost until reaching the Congestion Avoidance phase?

d) Consider the following network. Hosts A, B and C are connected to each other via router R. There are two TCP connections: A-B and C-B as shown below, and the roundtrip delays for TCP connections A-B and C-B are RTT_{A-B} = 10 msec and RTT_{C-B} = 15 msec. Let T_{A-B} and T_{C-B} denote the throughputs achieved by connections A-B and C-B, respectively. Assume that TCP's AIMD algorithm reaches the steady-state for both connections.

i) (6 pts) The transmission rate of links A-R, B-R and C-R are all 10 Mbps. What are the values of T_{A-B} and T_{C-B}?

ii) (6 pts) Assume now that the transmission rate on link A-R is changed to 2 Mbps. What are the values of T_{A-B} and T_{C-B}?