CS 421: Computer Networks **SPRING 2016**

MIDTERM I
March 24, 2016
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) An application can choose UDP as the transport layer protocol instead of TCP because UDP provides finer application layer control than TCP for determining what data is sent in a segment and when. In answering the following questions, recall the socket programming examples using TCP and UDP that we discussed in class.

i) (4 pts) Why does an application using UDP have more control (compared to TCP) of what data is sent in a segment?

ii) (4 pts) Why does an application using UDP have more control (compared to TCP) on when the segment is sent?

b) Consider a modification of DNS that runs over **TCP rather than UDP** and uses **iterative queries**. Suppose that **Host A** wants to obtain the IP address of **Host B**. Assume that the IP address of **Host B** is **not cached** in **Host A** or the local DNS server of **Host A**. Also assume that the local DNS server does not know the IP address of the TLD DNS server for the domain associated with Host B. **Assume that the round trip time between all entities (between Host A and each DNS server and between DNS servers) in the system is RTT.**

Assume that this version of DNS runs over **persistent** TCP connections (A TCP connection will be closed only after the DNS response message is received by the querying agent).

i) (4 pts) Write on the table given below the number of TCP connections that will be established between the entities in the corresponding row and the column of the table (Rows denote the TCP clients while columns denote the TCP servers**.** A value *n* in row *i* and column *j* will indicate that *n* TCP connections are established between the TCP client in row *i* and TCP server in row *j*).

TCP client/TCP server	Host A	Local DNS Server	TLD DNS Server	Root DNS Server	Authoritative DNS Server
Host A					
Local DNS Server					
TLD DNS Server					
Root DNS Server					
Authoritative DNS Server					

ii) (8 pts) Calculate the time from the initiation of the first TCP connection by Host A until Host A obtains the IP address of Host B.

c) (5 pts) Consider a web server running on port 80 with IP address 132.179.1.183 (Each host in the Internet has a unique IP address that is 32 bits long. Port numbers are 16 bits long). Assuming that the memory and bandwidth are unlimited, what is the theoretical maximum number of simultaneous TCP connections that this web server can have?

d) (5 pts) Consider a client that wants to establish simultaneous TCP connections with the server in part c). Assuming that the memory and bandwidth are unlimited, what is the theoretical maximum number of simultaneous TCP connections that this client can establish with the server.

2)

a) Assume that the bandwidth of a connection is 10 Mbps ($10x10^6$ bits/sec) and the round-trip propagation delay for the connection is 30 msec. Assume that each data segment is 2,500 Bytes long and the ACK segments have negligible lengths. **Separately answer the following questions for both Selective Repeat and Go-Back-N.**

   i. (8 pts) Assuming that no segments are lost, what should be the minimum window size (in packets) in order to achieve a bandwidth utilization of 100%, i.e., $U_{sender} = 1$?

   ii. (6 pts) What is the minimum number of bits necessary to represent the sequence numbers for proper operation using the window size that you calculated above?

b) (12 pts) Consider a 10 Mbps channel with a 20 msec roundtrip propagation delay. We want to transfer a file of size 18,000 Bytes. Each segment has a total size of 1,250 Bytes including the 125 Bytes header, i.e., each segment contains 1,125 Bytes of data. When there is data to be transmitted, each segment contains the maximum number of bytes. Assume that ACK segments have negligible lengths and there is a processing delay of 1 msec after a segment is fully received at the receiver until the transmission of the corresponding ACK is started. We use **Selective Repeat** protocol with a window size of N = 8 segments. Assume that **every 8th segment** crossing the forward channel is lost while ACKs are not lost or corrupted. Assume further that the processing delay at the sender after an ACK is received is negligible. How much time is required to complete the transfer of the whole file and receive the final ACK at the sender? Assume that the timeout for each segment is set to 30 msec starting from the end of the transmission of the segment.

c) (12 pts) Consider the use of Go-Back-6 protocol for communication from Node A to Node B. Assume that when the sender reaches the end of the window, i.e., all segments in the window are sent but not ACKed, the sender goes into timeout, i.e., it goes to the beginning of the window and retransmits all segments in the window as if timeout occurred. In the following figure, indicate the sequence number (SN) for segments sent from A to B, the ACK number (ACKN) for segments sent from B to A, the times and SN of the segments at B delivered to the application layer, and the window kept at A. **Note that segments received during the transmission of another segment will be immediately processed, but the corresponding action, e.g., update of SN/ACKN, will take effect with the start of transmission of the next segment.** Please use consecutive integers starting from 0 as sequence numbers.

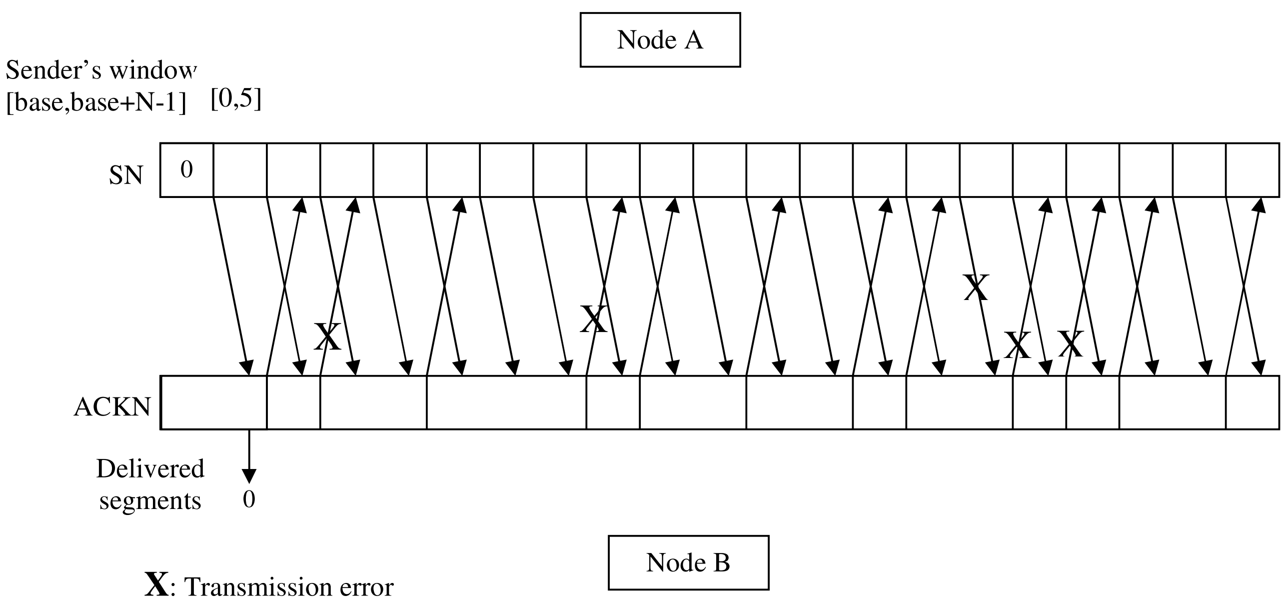

3)

a) (8 pts) Suppose a TCP connection experiences round-trip times (RTT) of 20 msec for 30% of its segments, 40 msec for 30% of its segments, 100 msec for 20% of its segments and no ACK is received for 20% of its segments. Suppose no segments are actually lost. Assume that the estimated RTT (according to the exponential weighted moving average) is equal to the true (ensemble) average of RTT, i.e., EstimatedRTT = average value of RTT = E[SampleRTT]. Assume that the timeout is set to **2 (two)** times the estimated RTT (as in the original version of TCP), i.e., TimeOut = 2 x EstimatedRTT. What fraction of the segments will be assumed lost by the TCP sender?

b) (8 pts) Assume that the congestion window of a TCP flow was 12 segments long when a **timeout** occurred. Assume that there are no segments or acknowledgments of this flow that were in transit when the timeout occurred. The round trip delay for the flow is fixed and is equal to 40 msec. 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 flow to reach the Congestion Avoidance phase after the timeout, assuming that no further segments are lost until reaching the Congestion Avoidance phase?

c) (8 pts) Consider the following network. Hosts A, B and C are connected to each other via router R. The bandwidths of the links A-R and R-B are R1 and R2, respectively, while the bandwidth on the link C-R is infinitely large. There are two TCP connections: A-B and C-B, and the roundtrip delays for both connections are equal. Let x and y 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 and that R1<R2/2. What are the values of throughputs x and y (as functions of R1 and R2)?

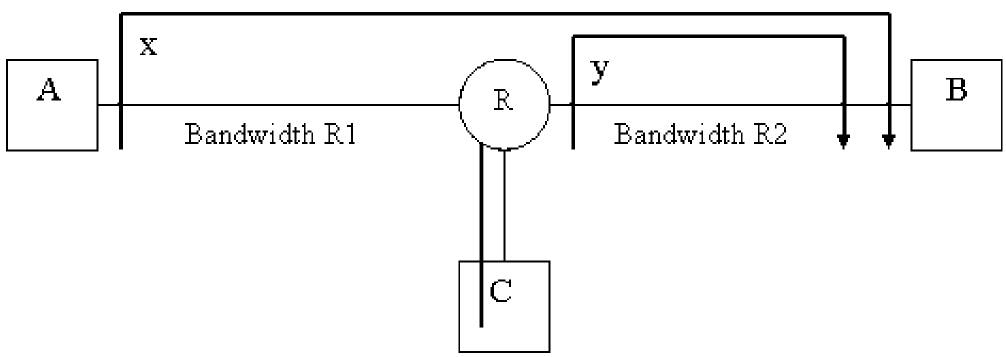

d) (8 pts) Suppose that R_{rcv}, the rate at which bits are arriving to a TCP receive buffer, is given in the following figure as a function of time. The application process at the receiver is removing bits from the receive buffer at the constant rate of 200 Kbits/sec (2 x 10^5 bits/sec). Assume that the receive buffer is initially empty at t = 0, and it has a fixed size of 50,000 Bytes. Considering the TCP flow control algorithm, what is the value of the Receive Window advertised by the receiver at t = 2 sec?