CS 421: Computer Networks **FALL 2005**

MIDTERM I
November 9, 2005
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.

1)

a) (5 pts) Consider a time instant, *t*, at which a user has one *FTP* file download in progress at the same time as three browser windows are open, in the process of downloading three different web pages. The browser uses *HTTP*/1.1 with no pipelining. How many concurrent *TCP* connections will be open at the time instant *t*? (**You must briefly explain/justify your answer to receive full credit.)**

b) (5 pts) Suppose two IP datagrams are sent over the Internet, each carrying a different UDP segment. The first datagram has source IP address S1, destination IP address D, source port SP1, and destination port DP. The second datagram has source IP address S2, destination IP address D, source port SP2, and destination port DP. Assume that S1 is different from S2 and SP1 is different from SP2. Assuming that both datagrams reach their final destination host, will the two UDP datagrams be received by the same socket? Why or why not? Fully justify your answer.

c) (6 pts) Consider a Web server where the maximum number of clients running at the same host that can be concurrently connected to the server is one. Determine the smallest subset of the fields, source IP address, destination IP address, source port number and destination port number, necessary for demultiplexing so that to each segment can be delivered to the appropriate socket.

d) (6 pts) Give an example application that we discussed in this course which uses the hybrid application layer architecture. Point out application functionalities that use the client-server paradigm, and those that use the peer-to-peer paradigm.

2)

a) (10 pts) Consider a 1 Mbits/sec channel with a 20 msec one-way propagation delay, i.e., 40 msec roundtrip propagation delay. We want to transfer a file of size 11000 Bytes. Each packet has a total size of 1500 Bytes including the 125 Bytes header, i.e., each packet contains 1375 Bytes of data. When there is data to be transmitted, each packet contains the maximum number of bytes. Assume that ACK packets are of 125 Bytes long and there is a processing delay of 1 msec after a packet is fully received at the receiver until the transmission of the corresponding ACK is started. **Selective Repeat** protocol is used with a window size of N = 4 packets. Assume that every 6^{th} packet crossing the forward channel is lost while ACKs are not lost or corrupted. 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 packet is set to 3ms larger than the **total** round-trip delay (including transmission, propagation and processing delays).

b) (15 pts) Consider the use of Go-Back-7 protocol for communication from Node A to Node B. Assume that when the sender reaches at the end of the window, i.e., all packets in the window are sent but not ACK'd, the sender goes to the beginning of the window and retransmits all packets in the window, as if the timeout occurred. In the following figure, indicate the sequence number (SN) for packets sent from A to B, the ACK number (ACKN) for packets sent from B to A, the times and SN of the packets at B delivered to the application layer, and the window kept at A. **Note that packets received during the transmission of another packet 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 packet.**

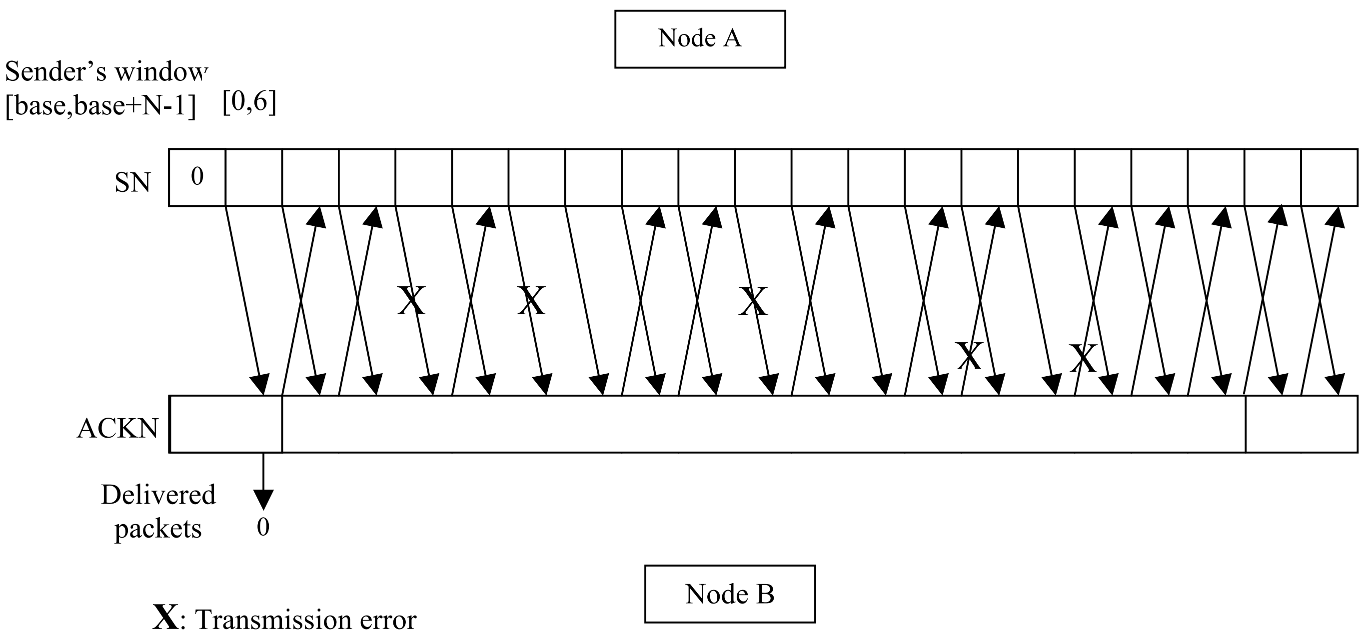

c) (8 pts) A satellite transmission link is to be designed for communication between two earth stations with a transmission rate of 64 Kbits/sec. The round-trip distance between each earth station and the satellite is 90,000 km and the propagation speed in the free space is 3×10^5 km/sec. The satellite just implements the physical layer, therefore it introduces a negligible delay before relaying a packet received from one earth station to the other earth station. Assume that all packets, i.e., data and ACK packets, are 1000 bits long. Once a data packet is received by an earth station, there is a processing delay of 5 msec before the corresponding ACK is generated. Suppose that we use the Selective Repeat protocol for communication between the two earth stations. What is the minimum window size N, in number of packets, that you would use so that full channel utilization can be achieved assuming that there are no lost or errored packets? How many bits are needed for representing the SN for the minimum window size you computed above?

3)

a) (6 pts) Suppose that TCP receives a segment for which the checksum fails. Consider the following two possible actions for the receiver:

   i. Discard the segment and send a duplicate ACK.
   ii. Just discard the segment and send nothing.

   Which one of these two possible actions is more appropriate? Or are these two actions equivalent? Why? **(You need to justify your answer for getting points. Use at most 3 sentences for justifying your answer.)**

b) (6 pts) Why does TCP use the Congestion Avoidance phase in addition to the Slow Start instead of continuing with the Slow Start even when the congestion window exceeds the threshold? Are there any performance problems that may arise if the sender just uses the Slow Start and never uses the Congestion Avoidance? **(Use at most 3 sentences.)**

c) (6 pts) Suppose upon receiving a segment containing *N* Bytes of data, the TCP receiver divides the received data into *M* parts, where $M \leq N$, and sends *M* separate acknowledgements each covering one of the *M* distinct pieces of the received data sequence instead of sending a single acknowledgement. Is there any advantage for the TCP receiver obtained as a result of this behavior? **(Use at most 3 sentences.)**

d) (6 pts) The adaptive timeout mechanism of TCP **does not** update EstimatedRTT in case an ACK is received for a segment that was retransmitted. Why? **(Use at most 3 sentences.)**

4)

a) (6 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 1 Mbits/sec (1×10^6 bits/sec). Assume that the receive buffer is initially empty and it has a fixed size of 150,000 Bytes. What is the value of the Receive Window advertised by the receiver at t = 2 sec?

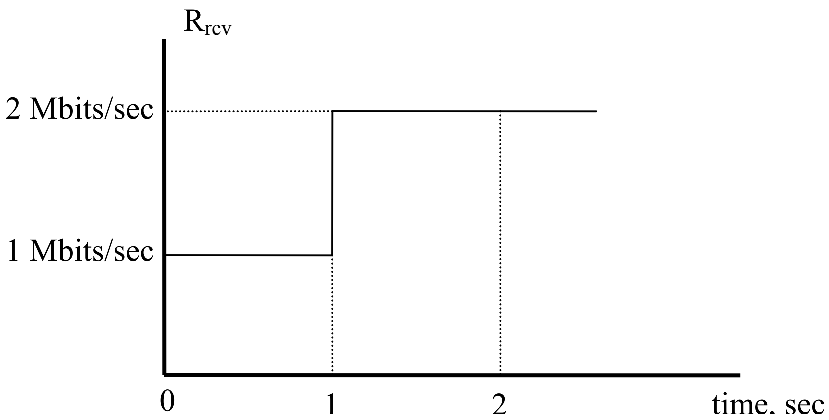

b) (8 pts) Assume a congestion feedback model for a system composed of two flows sharing a bottleneck link with bandwidth R bits/sec where both connections have the same RTT. Each flow gets binary synchronous feedback in discrete time steps of one RTT. If the aggregate consumption of the two flows is above the bottleneck bandwidth, both senders receive a congestion notification signal (**CN**), otherwise they receive no CN. The flows use a simple congestion control scheme: When no CN is received in a time step, each sender increases its window by one packet. On the other hand, for each congested time step, i.e., when both senders receive CN, each sender decreases its window by one packet. So, we can call this algorithm as *Additive Increase-Additive Decrease* (**AIAD**). **Prove or disprove** that AIAD achieves a fair allocation of bandwidth between the flows, i.e., each flow getting R/2, by using graphical arguments similar to the one we made in class in showing that TCP's AIMD algorithm is fair.

c) (7 pts) Consider the figure below showing the Congestion Window for a TCP connection as a function of time. The x-axis (time axis) denotes the Transmission Round (**TR**), where each tick corresponds to one round-trip-time (RTT) (assume that data packets and ACKs have negligible transmission times). The y-axis is the Congestion Window in segments at the beginning of each TR. Use the table provided on the next page in order to identify the time intervals during which Congestion Control algorithm is in a certain phase, i.e., Slow Start (**SS**) or Congestion Avoidance (**CA**), the value of the Slow Start threshold (**ssthresh**) during that time interval and the reason for the phase change which ends that time interval for TR = 1,…,36.

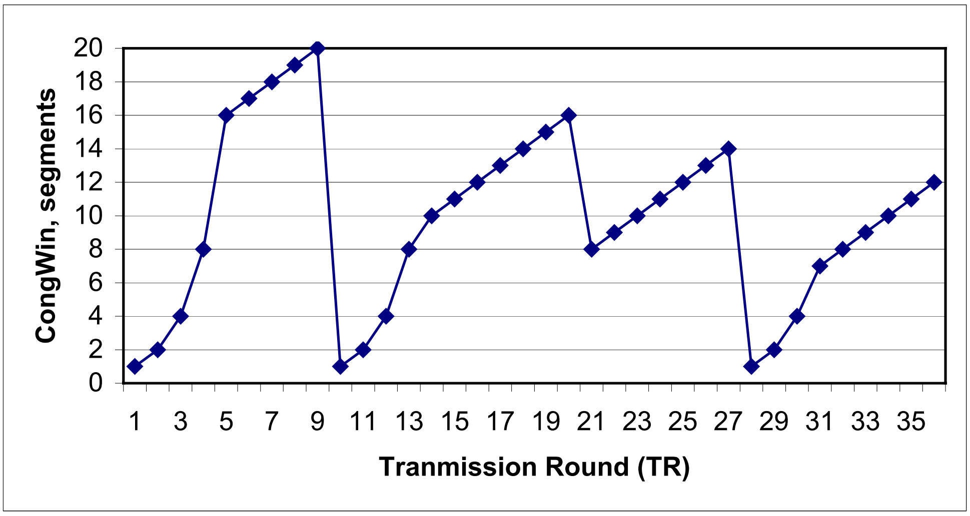

Time Interval [starting TR,ending TR]	Phase (SS or CA)	ssthresh (in segments)	Event causing phase change
[1 ,]			
[,]			
[,]			
[,]			
[,]			
[,]			
[, 36]			