

This homework is due at the beginning of class on February 25, 2015 and is worth 1.5% of your grade.

Name: _____

CCIS Username: _____

Problem	Possible	Score
1	20	
2	20	
3	15	
4	10	
5	15	
6	15	
7	15	
8	15	
Total	125	

1a. TCP packets are being sent from a client to a server. The MSS is equal to 1460 bytes, and each TCP packet is sent with the maximum capacity. How many TCP packets can be sent before the sequence number field in the TCP header will wrap around? (10 pts)

1b. How much time (in seconds) will this take on a 1 Mbit/s link? (5 pts)

1c. How much time (in seconds) will this take on a 1 Gbit/s link? (5 pts)

2a. In the stop-and-wait protocol (i.e., the simple reliable transport protocol where the sliding window size is always one packet), what is the minimum number of bits required to encode the sequence number for the protocol to work correctly? (10 pts)

2b. If sequence number is not used, the resulting "broken" stop and go protocol may fail to transfer data correctly. Describe a specific scenario where not using a sequence number in the stop-and-go protocol would result in incorrect data transmission. (10 pts)

3a. If we generalize from the stop-and-wait protocol to a sliding window protocol, where k packets can be sent but unacknowledged, what is the minimum number of distinct sequence numbers that we would need for this protocol to work correctly? Why? (10 pts)

3b. Suppose we are using a sliding window protocol with a window size of 128 KB and a round-trip time of 100 milliseconds. What is the expected sending rate of this protocol? (5 pts)

4. Host A is transferring a file of size S to host B using TCP. A sends the file data in fixed size packets equal to the Maximum Segment Size (MSS), a predetermined value. B sends an acknowledgement immediately upon receiving a data segment. Let R be the round trip delay between A and B . The advertised receiver window size of host B is W . In this problem, we assume the TCP connection is already established and that the transmission time is negligible. TCP performs the slow start and congestion avoidance mechanisms, and there is no error or packet loss during transmission.

4a. Given $W = 3 * \text{MSS}$, $S = 10 * \text{MSS}$, how long does it take for the file to be sent and acknowledged? Show your work. (5 pts)

4b. Given $W = 5 * \text{MSS}$, $S = 15 * \text{MSS}$, how long does it take for the file to be sent and acknowledged? Show your work. (5 pts)

5a. Suppose that a TCP connection at a sender has a receiver's advertised window of r and a congestion window of c . What is the value of the sender's window? (5 pts)

5b. Suppose that we design a variant of TCP that uses MIAD (multiplicative increase, additive decrease) as the congestion control mechanism for updating the congestion window. What would happen if two of our TCP flows compete at a bottleneck router? (10 pts)

6a. Recall that TCP has two phases: *slow start* and *congestion avoidance*. What is the primary purpose of the slow start phase? (5 pts)

6b. Recall that TCP uses timeouts to determine if a packet has been lost (meaning if the ACK for a packet does not arrive within a specific period of time, the packet is considered to be lost). However, this requires the sender to wait for a long period of time to detect if the packet is indeed lost. It would seem that TCP could use negative-acknowledgements (NACKs), where a receiver would proactively notify the sender with a NACK message if the packet was not received. Why do you think the designers of TCP chose not to use NACKs? (10 pts)

7a. Recall that we discussed two TCP variants in class: TCP Tahoe and TCP Reno. However, in practice, most machines use TCP Reno. What is the primary reason why people have moved away from TCP Tahoe? (5 pts)

7b. Give an example of a case where the fast retransmit in TCP Reno would result in unnecessary packets being sent. (5 pts)

7c. What sequence of events would lead to TCP Reno *re-entering* the slow start phase after being in the congestion avoidance phase? Give a concrete example. (5 pts)

8a. When closing a TCP connection, why is the two-segment-lifetime timeout not necessary on the transition from LAST ACK to CLOSED? (5 pts)

8b. A sender on a TCP connection that receives a 0 advertised window periodically probes the receiver to discover when the window becomes nonzero. Why would the receiver need an extra timer if it were responsible for reporting that its advertised window had become nonzero (i.e., if the sender did not probe)? (5 pts)

8c. The sequence number field in the TCP header is 32 bits long, which is big enough to cover over 4 billion bytes of data. Even if this many bytes were never transferred over a single connection, why might the sequence number still wrap around from $2^{32} - 1$ to 0? (5 pts)