CMPS 4750/6750 Computer Networks – Spring 20

Homework 4 (30+10 points)

Due on 03/31/20 before class

Questions marked with * are required for graduate students only.**

1. **rdt 3.0 and reordering (4 points)**

   Consider the rdt 3.0 protocol. Draw a trace (similar to Figure 3.16 in the main textbook) showing that if the network connection between the sender and the receiver can reorder messages (that is, two messages propagating in the medium between the sender and receiver can be reordered), then the rdt 3.0 protocol will not work correctly. Your trace should have the sender on the left hand and the receiver on the right, with the time axis running down the page, showing data (D) and acknowledge (A) message exchange. Make sure you indicate the sequence numbers assocated with any data or acknowledge segment.

2. **Go-Back-N (4 points)**

   Consider the GBN protocol with a sender window size of N. Suppose that at time t, the next in-order packet that the receiver is expecting has a sequence number of k. Assume that the size of the sequence number space is much larger than N, and that the medium does not reorder messages. What are all possible values of the ACK field in all possible messages currently propagating back to the sender at time t? Justify your answer.

3. **Sequence number space and window size (6 points)**

   Consider the GBN and SR protocols. Suppose the sequence number space is of size M. What is the largest allowable sender window that will avoid the occurance of problems such as that in Figure 3.27 in the main textbook (also discussed on slide #53) for each of these protocols?

4. **TCP sequence number (6 points)**

   Hosts A and B are communicating over a TCP connection, and Host B has already received from A all bytes up through byte 126. Suppose Host A then sends two segments to Host B back-to-back. The first and second segments contain **50** and **100** bytes of data, respectively. In the first segment, the sequence number is 127, the source port number is 302, and the destination port number is 80. Host B sends an acknowledgement whenever it receives a segment from Host A.

   (a) In the second segment sent from Host A to B, what are the sequence number, source port, and destination port?

   (b) If the second segment arrives before the first segment, in the acknowledgement of the first arriving segment, what is the acknowledgement number?

   (c) Suppose the two segments sent by A arrive in order at B. The first acknowledgement is lost and the second acknowledgement arrives after the first timeout interval. Draw a timing diagram (similar to Figure 3.34 in the main textbook), showing these segments and all other segments and acknowledgements sent (Assume there is no additional packet loss.) For each segment in your figure, provide the sequence number and the number of bytes of data; for each acknowledgemnet that you add, provide the acknowledgement number.

5. **AIAD and fairness (10 points)**

   We discussed the fairness of TCP in class and showed that under certain assumptions, the AIMD algorithm converges to an equal bandwidth share (see slide #96 and also refer to Section 3.7.1 in the main textbook). Suppose that instead of a multiplicative decrease, TCP decreases the window size by a constant amount when there is a loss. Assume that the two connections have the same MSS and RTT and ignore slow start as we did in class. Would the resulting AIAD (additive-increase, additive-decrease) algorithm converge to an equal share algorithm in each of the following two cases? Justify your answer using a diagram similar to the one given on slide #96 and Figure 3.55 in the main textbook.

   (a) Whenever there is a loss, each connection decreases its window by the same amount;

   (b) Whenever there is a loss, connection 1 decreases its window by twice the amount of connection 2.

6. *** **TCP congestion control and fairness (10 points)**

   Consider a simplied TCP's AIMD algorithm where the congestion window size is measured in number of segments, not in bytes. In additive increase, the congestion window size increases by one segment in each RTT. In multiplicative decrease, the congestion window size decreases by half (if the result is not an integer, round down to the nearest integer). Suppose that two TCP connections C_1 and C_2, share a single congested link of speed **50 segments** per second. The two connections have the same MSS, but a *different* RTT. Connection C_1's RTT is 50 msec and connection C_2's RTT is 100 msec. Assume that when the total data rates (from the two connections) in the link exceeds the link's speed, all TCP connections experience data segment loss. Ignore the slow-start phase and assume that both connections are operating in the congestion avoidance phase (AIMD) at all times.

   (a) If both C_1 and C_2 at time t_0 have a congestion window of 10 segments, show the evolution of the window sizes in the next 500 msec.

   (b) In the long run, will these two connections get the same share of the bandwidth of the congested link? Explain.