CS 536 Fall 2009 - Homework 3
Due 11/05/2009 to Blackboard by class time

Problem 1 (30 pts)

TCP. Consider the following plot of TCP window size as a function of time.

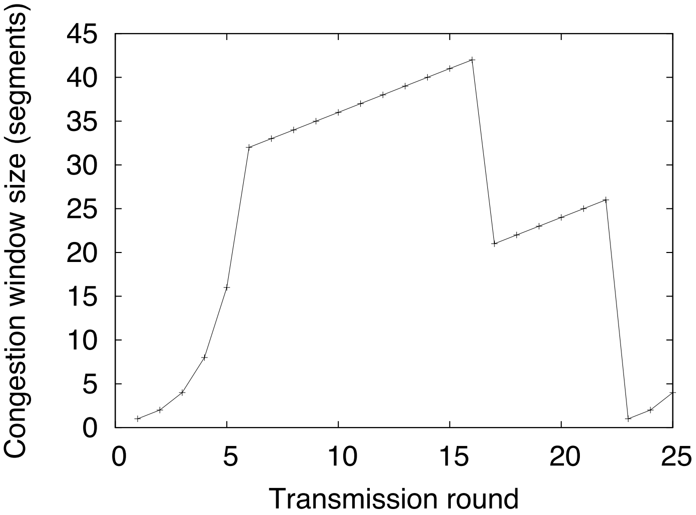

Assuming TCP Reno is the protocol experiencing the behavior shown above, answer the following questions. In all cases, you should provide a short discussion justifying your answer.

1. Identify the intervals of time when TCP slow start is operating.
2. Identify the intervals when TCP congestion avoidance is operating.
3. After the 16th transmission round, is segment loss detected by a triple duplicate ACK or by a timeout ?
4. After the 22nd transmission round, is segment loss detected by a triple duplicate ACK or by a timeout ?
5. What is the initial value of the `Threshold` at the first transmission round ?
6. What is the value of the `Threshold` at the 18th transmission round ?
7. What is the value of the `Threshold` at the 24th transmission round ?

8. During what transmission round is the 70th segment sent ?
9. Assuming a packet loss is detected after the 26th round by the receipt of a triple duplicate ACK, what will be the values of the congestion window size and of `Threshold`.
10. What is the average throughput in 20 rounds assuming an average of 10 ms per round and all segments are 1500 bytes?

Problem 2 (10 pts)

Consider the rdt3.0 protocol. Draw a diagram showing that if the network connection between the sender and receiver can reorder messages (that is, that two messages propagating in the medium between the sender and receiver can be reordered), then the alternating-bit protocol will not work correctly (make sure you clearly identify the sense in which it will not work correctly). Your diagram should have the sender on the left and the receiver on the right, with the time axis running down the page, showing data (D) and acknowledgement (A) message exchange. Make sure you indicate the sequence number associated with any data or acknowledgement segment.

Problem 3 (10pts)

No packet loss, no timer expiry. Host A wants to send an enormous file to Host B over a TCP connection. Over this connection there is never any packet loss and timers never expire. Denote the transmission rate of the link connecting host A to the Internet by R bps. Suppose that the process in Host A is capable of sending data into its TCP socket at a rate of S bps, where $S = 10 \times R$. Further suppose that the TCP receive buffer is large enough to hold the entire file, and the send buffer can hold only one percent of the file. What would prevent the process in Host A from continuously passing data to its TCP socket at rate S bps ? TCP Flow control ? TCP Congestion control ? Or something else ? Elaborate.

Problem 4 (10 pts)

In this problem we explore the delay introduced by the TCP slow-start phase. Consider a client and a server directly connected by one link of rate R. Suppose the client wants to retrieve an object whose size is exactly equal to $15S$, where S is the maximum segment size (MSS). Denote the round-trip time between client and server at RTT (assumed to be constant). Ignoring protocol headers, determine the time to retrieve the object (including TCP connection establishment) when

a. $4S/R > S/R + RTT > 2S/R$

b. $S/R + RTT > 4S/R$

c. $S/R > RTT$

Problem 5 (20pts)

RED. Let us suppose a router uses RED as a way to avoid congestion buildup in the network. The queue in the router is shared by two connections A and B. The average queue length is 10 packets while the min and max thresholds are 5 and 15 packets respectively, so the average queue length is halfway between the two thresholds. Assume the RED implementation uses a $maxP = 0.02$.

1. What is the drop probability for a new incoming packet for connection A if the number of packets queued from the time the queue length crossed the min threshold is 10.

2. Suppose A's congestion window consists of 8 packets. (Typically TCP congestion window is specified in bytes, but lets suppose all packets are of same size, so it does not matter.) Let B's congestion window also be of size 8 packets. Now a train of A's packets (8 of them) and B's packets (8 of them) arrive in sequence. When the first A packet arrived, the count of number of packets queued after is 5. What is the probability that A will not lose a packet ?. What is the probability that B will not lose a packet ?. To make these calculations easy, assume that none of the previous packets got dropped for each of A's as well as B's packets.

Problem 6 (20 pts)

Consider transferring an enormous file of L bytes from Host A to Host B. Assume an MSS of 1460 bytes.

a. What is the maximum value of L such that TCP sequence numbers are not exhausted? Recall that the TCP sequence number field has 4 bytes.

b. For the L you obtain in (a), find how long it takes to transmit the file. Assume that a total of 66 bytes of transport, network, and data-link header are added to each segment before the resulting packet is sent out over a 10 Mbps link. Ignore flow control and congestion control so A can pump out the segments back to back and continuously.

c. Suppose now the link were a 50 Gbps link. Retaining other parameters from (b), how long will it take to transmit the file now?

d. If the file were instead were longer than L, some sequence numbers would be re-used. Assume now that the Host B acknowledges segments from Host A, and that there exist two paths between B and A. If a handful of acknowledgements take an alternate path, which experiences severe queueing delay, they might be confused with new acknowledgements if they arrive after the data with the same sequence number is sent. If the propagation delay is 300ms on each path (in both directions), how much queueing delay would be necessary to cause this confusion to occur? What implications might this have for the design of high-bandwidth network transport protocols?