CS 168 Section 6: TCP

1. TCP Sequence Numbers

A TCP connection has been established between hosts A and B. B receives the following packet from A with the field values shown below:

Sequence Number 101	**ACK** 10001
Window Size 5000	**Payload Size** 1400

All numbers are in terms of *bytes*. B processes the packet and updates its sliding window. For parts a-c, assume the congestion window is not a factor (i.e. CWND = infinity).

a) B wants to respond to A. What is the range of sequence numbers B can use for its next packet?

Lowest:

Highest:

B sends a packet with the following fields to A:

Sequence Number 13001	**ACK**
Window Size 50	**Payload Size** 1400

b) Fill in the ***ACK*** field in the packet above.

c) What is the range of sequence numbers that A can send in response?

Lowest:

Highest:

d) Assume B's operating system always sets the receive window size to 65535B (the maximum possible). Can A can always send 65535B of data at once, assuming it's own window is big enough and packets can be infinitely large?

2. Congestion Control: Slow Start

Anna wants to send Elsa some data. She wants to sends as fast as possible, but doesn't want to congest the network. Anna decides to use *slow start* to ramp up the amount of traffic she sends to Elsa. Anna establishes a connection, and starts sending 100B packets. A couple reminders:

- *On each ACK,* ***CWND += MSS*** *during slow start.*
- *For this question,* ***MSS = 100B***

a) Fill out the following table. The sender **receives an ACK for a packet sent at time T at time T+1**. There should be one ACK per row in the table. Multiple ACKs may arrive at the same time (so some rows may have the same "time" value).

At Time:	When the sender receives ACK #:	She sets her CWND to:	And Sends Packets with the following sequence numbers:
0	---	100	1

b) Keeping in mind that each packet is 100B long, plot the amount of data the sender has sent and is not ACKed at **the end of each timestep** below.

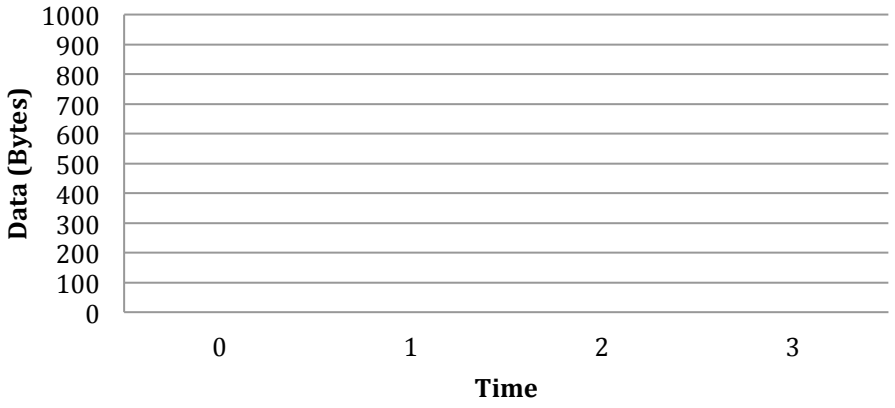

c) Which of the following best describes the growth in the graph above?

Logarithmic **Linear** **Quadratic** **Exponential**

3. Congestion Control: Additive Increase and Fast Retransmit

Hans is sending data to The Duke of Weselton. Because Hans has been sending data for some time already, he's currently in the *Additive Increase, Multiplicative Decrease (AIMD)* phase of TCP congestion control. Assume Hans uses ***simple fast retransmit***. Again, the MSS is 100B for all parts.

a) Fill in the *CWND* and *SSTHRESH* values for the following series of ACKs. When dividing integers, always round down (e.g. 55/10 = 5). Some reminders:

- *During congestion avoidance,* ***CWND += MSS / ⌊ CWND/MSS ⌋*** *on new ACK*
- *On triple duplicate ACK,* ***SSTHRESH = CWND/2****, then* ***CWND = SSTHRESH***

Time	ACK	CWND	SSTHRESH
0	801	1000	800
1	901		
2	1001		
3	1101		
4	1101 (1)		
5	1101 (2)		
6	1101 (3)		

(space for scratch work)

b) Which of the following does fast retransmit *primarily* respond to?

a. Congestion in the Network
b. Isolated Packet Loss
c. Overflowing the Receiver's Buffer

c) Which of the following best describes the growth of traffic sent per timestep during the AIMD phase?

Logarithmic **Linear** **Quadratic** **Exponential**

d) Say Hans decided to use ***advanced fast retransmit with fast recovery*** instead of simple fast retransmit. Fill out the *CWND* and *SSTHRESH* values in the table below. Assume Hans is in the AIMD phase at T=0. Reminders:

- *In Congestion Avoidance, Fast Retransmit/Fast Recovery works the same as above*
- *On triple duplicate ACK,* ***SSTHRESH = CWND/2,*** *then* ***CWND = SSTHRESH + 3 MSS***
- *In fast recovery,* ***CWND += MSS*** *on every duplicate ACK*
- *Exit fast recovery on new ACK, setting* ***CWND = SSTHRESH***

Time	ACK	CWND	SSTHRESH
0	801	1000	800
1	901		
2	1001		
3	1101		
4	1101 (1)		
5	1101 (2)		
6	1101 (3)		
7	1101		
8	1101		
9	1101		
10	1801		
11	1901		

e) Say a timeout occurs at the sender. Which of the following does this generally signify?

a. Congestion in the Network
b. Isolated Packet Loss
c. Overflowing the Receiver's Buffer

f) Does your answer in part e) differ from your answer in part b)? Explain why TCP responds differently in the two scenarios.

4. Congestion Control: Analyzing the TCP Congestion Window

Olaf sends some traffic to Kristoff. While this happens, you record the size of the congestion window (CWND).

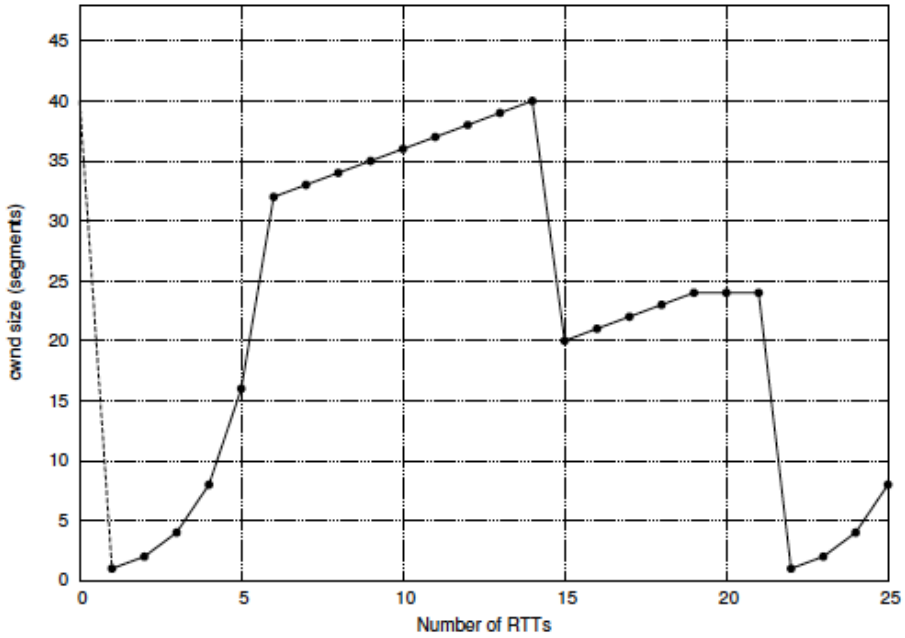

Assume that the connection has been operating for some time, meaning that the number of RTTs shown is with respect to when you started observing the connection behavior.

a) Identify the time periods when TCP slow start is operating.

b) Identify the time periods when TCP congestion avoidance is operating (AIMD).

c) After the 14th RTT, is the segment loss detected by a triple duplicate ACK or by a timeout?

d) What is the initial value of SSTHRESH, before the first congestion avoidance interval?

e) What is the value of SSTHRESH at the 19th RTT?

f) What is the value of SSTHRESH at the 24th RTT?

g) Assuming a packet loss is detected after the 25th RTT by the receipt of a triple duplicate ACK, what will be the values of the congestion-window size and of SSTHRESH?