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.

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

Lowest:
10001

Highest:
15000 (B's buffer only reserved space for bytes 10001 to 15000, a total of 5000 bytes).

B sends a packet with the following fields to A:

Sequence Number 13001	**ACK** 1501
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:
1501

Highest:
1550

d) Assume B's operating system always sets the 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?

No – A sends MIN(CWND, receiver window), so it may send less data if the congestion window is smaller.

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
1	101	200	101, 201
2	201	300	301, 401
2	301	400	501, 601
3	401	500	701, 801
3	501	600	901, 1001
3	601	700	1101, 1201
3	701	800	1301, 1401

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.

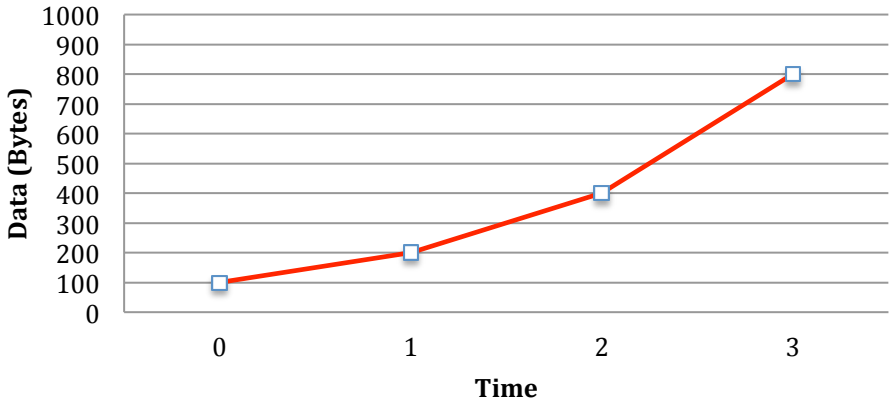

c) Which of the following best describes the growth in the graph above?

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	1010	800
2	1001	1020	800
3	1101	1030	800
4	1101 (1)	1030	800
5	1101 (2)	1030	800
6	1101 (3)	515	515

(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	1010	800
2	1001	1020	800
3	1101	1030	800
4	1101 (1)	1030	800
5	1101 (2)	1030	800
6	1101 (3)	815	515
7	1101	915	515
8	1101	1015	515
9	1101	1115	515
10	1801	515	515
11	1901	535	515

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.

TCP responds differently because when there's an isolated drop, there's no need to aggressively throttle the sending rate. On timeout, when multiple packets are dropped, the sending rate is dropped to 1MSS to alleviate network congestion.

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).

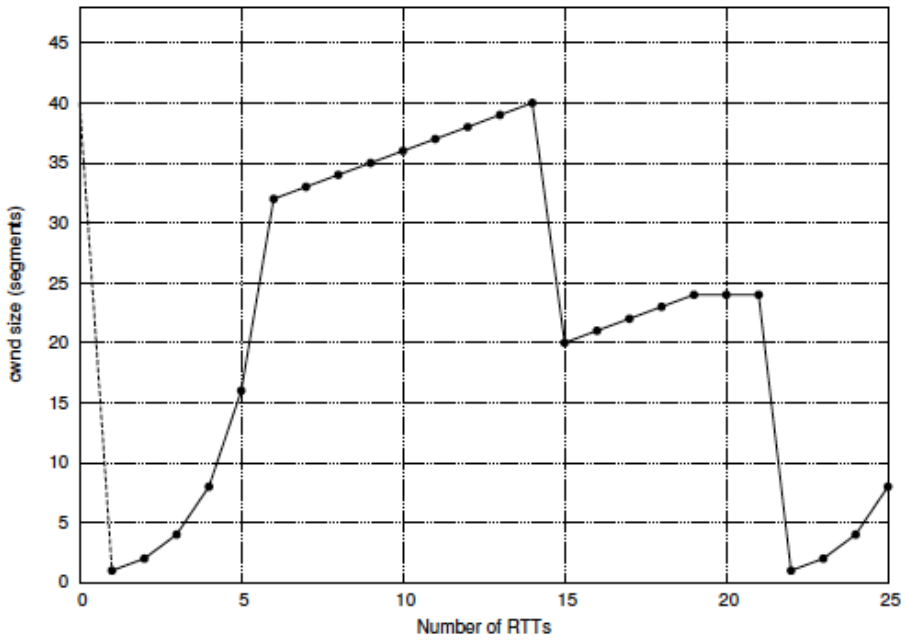

Assume that the connection has been operating for some time, meaning that the number of RTTs shown is with respect to when you started observing the connection behavior.

a) Identify the time periods when TCP slow start is operating.
1-6, 22-25

b) Identify the time periods when TCP congestion avoidance is operating (AIMD).
6-14, 15-21 (also accepting 15-19, 6-21)

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

d) What is the initial value of SSTHRESH, before the first congestion avoidance interval?
32

e) What is the value of SSTHRESH at the 19th RTT?
20

f) What is the value of SSTHRESH at the 24th RTT?
12

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?

Solution 1:

The three duplicate ACKS arrive immediately after CWND = 8, in which case

SSTHRESH = CWND / 2 = 8/2 = **4**
CWND = CWND / 2 = **4**

Alternatively, we also accepted **CWND** = CWND / 2 + 3 = **7**

Solution 2:

The three duplicate ACKS occur immediately before the end of the 26th RTT. This case is tricky because between t = 25 and t = 26, the CWND hits the previous SSTHRESH of 12 and enters congestion avoidance.

Of the 8 ACKs within this RTT, the first five are new ACKs, and the last three are duplicate ACKs.

After the first 4 new ACKs, CWND = 12. On the arrival of the fifth new ACK, however, CWND += 1/8, so CWND = 12.125. Then three duplicate ACKs are received, which do not change the value of CWND, and then the window is halved:

SSTHRESH = CWND / 2 = 12.125 / 2 = **6.0625**
CWND = CWND / 2 = 12.125 / 2 = **6.0625**

We also accept SSTHRESH = 6 and CWND = 6. Alternatively, if the +3 version of Fast Recovery is used, then CWND = 9.