

TCP Congestion Control - Questions

Q1) Figure 1 shows the trace of Congestion window for a particular TCP implementation. Would it be a TCP Reno or TCP Tahoe? Why?




Figure 1. TCP Congestion Window trace

Q2) In Figure 1, what would be the value of TCP Congestion window at point B?

Q3) Figure 2 shows congestion window traces for both TCP Tahoe and TCP Reno where up to transmission round of 8, both follow the same blue curve, but after that TCP Tahoe follows the blue curve and TCP Reno follow the black curve.




Figure 2: Congestion window trace for TCP Tahoe and TCP Reno

Answer the following Questions:

- a) What has happened at transmission round No 8?
- b) Identify the regions where Slow Start (SS) and Congestion Avoidance (CA) are in operation?
- c) How many total segments has been transferred by TCP Tahoe and TCP Reno at the end of round 15?

Q4) TCP waits until it has received three duplicate ACKs before performing a fast retransmit. Why do you think the TCP designers chose not to perform a fast retransmit after the first duplicate ACK for a segment is received?

Q5) Consider sending a large file from a host to another over a TCP connection that has no loss. Suppose TCP uses AIMD for its congestion control with the SSthreshold at 6 MSS. Assuming approximately constant round-trip times, how long does it take for cwnd to increase to 12 MSS (assuming no loss events, neglect TCP connection establishment time)?