

CS4450 Problem Set #6

May 1, 2018

1

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 without slow start. Assuming $cwnd$ increases by 1 MSS every time an ACK is received and assuming approximately constant round-trip times, how long does it take for $cwnd$ to increase from 6 MSS to 12 MSS (assuming no loss events)?
- Suppose $MSS = 1000$ bytes. What is the average throughput (in terms of MSS and RTT) for this connection up through time = 6 RTT after connection setup, where throughput is defined as the total amount of data sent over the specified time?

2

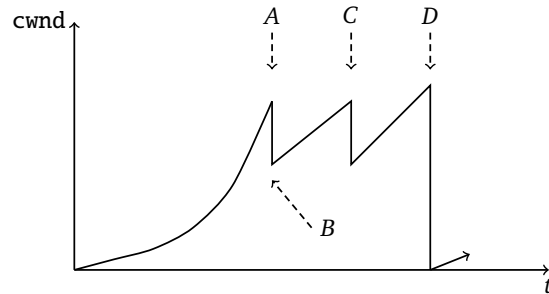
- (a) A sender and a receiver are communicating using a TCP-based protocol in which data is sent immediately after connection. The bandwidth of the link is 100Mbps, the round-trip time is 10ms and the maximum segment size is 1500 bytes (ignore header overhead). Draw a packet exchange diagram showing how long it takes to transfer 15000 bytes.
- (b) Consider that only a single TCP connection uses one 10Mbps link which does not buffer any data. Suppose that this link is the only congested link between the sending and receiving hosts. Assume that the TCP sender has a huge file to send to the receiver, and the receiver's receive buffer is much larger than the congestion window. We also make the following assumptions: each TCP segment size is 1,500 bytes; the RTT for this connection is 150 msec; and this TCP connection is always in congestion avoidance phase, that is, ignore slow start.
- What is the maximum window size (in segments) that this TCP connection can achieve?
 - What is the average window size (in segments) and average throughput (in bps) of this TCP connection?
 - How long would it take for this TCP connection to reach its maximum window again after recovering from a packet loss?

3

- (a) Consider two TCP flows, one over link X, and the other over link Y. X is a 100Mbps Ethernet link with a 10ms RTT, and Y is a 1Gbps satellite link with a 1000ms RTT. Assuming both links have the same loss probability p , and the flows always use the same packet size and never saturate the links, what's the link utilization ratio of X to Y?
- (b) Consider two TCP flows, one over wireless link X, and the other over wireless link Y. Both links have 25ms RTT. If the interference causes a 1% packet loss on X, and a 4% packet loss on Y, respectively, what's the link throughput ratio of X to Y? (Assume there's no forward error correction, and there is no link layer retransmission.)
- (c) From the above answers, we know that TCP does not perform very well in the presence of high RTT or lossy links. Give one solution to improve TCP throughput on links with high RTT, and one solution to improve TCP throughput on lossy links.
- (d) Suppose a TCP flow and a UDP flow are sharing the same link, and the UDP sends at a rate equal to the link capacity. Assuming the flows stay for long enough, how would they split the bandwidth eventually? If they would equally share the bandwidth, rationalize it. If a particular one would monopolize the bandwidth, explain which one and why.

4

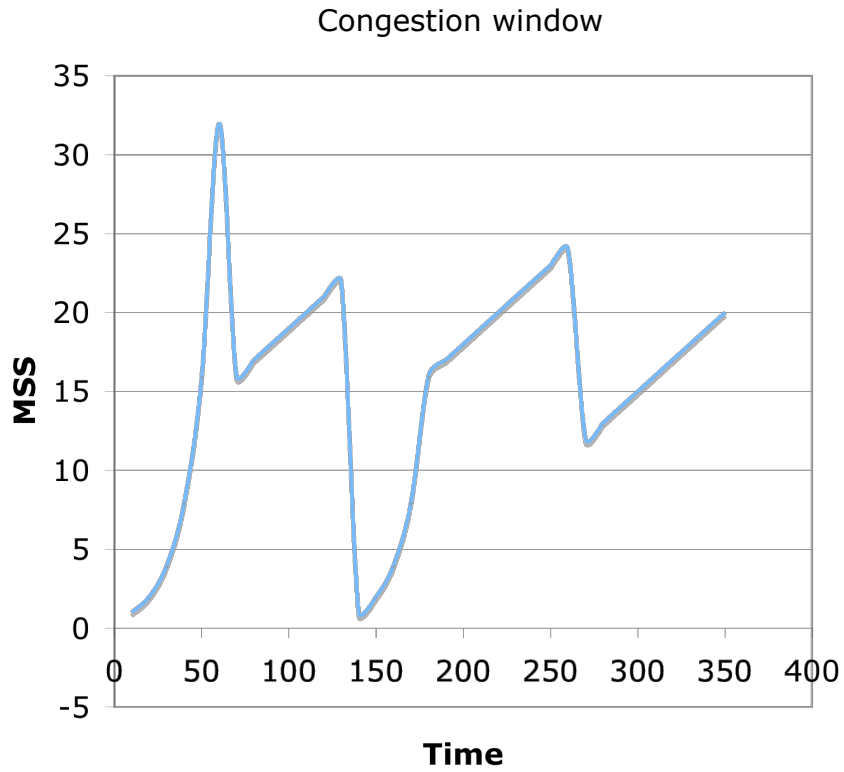
Consider the plot below, where x-axis denotes the time and y-axis denotes the TCP window size for TCP.



- Describe the slow start mechanism (from time 0 to event A), and the AIMD mechanism (between event A and C) in at most two lines each.
- Name the event that triggers the window size reduction at A.
- Does the event at A necessitate a packet drop? Why or Why not?
- Name the event that triggers the window size reduction at D.
- Does the event at D necessitate a packet drop? Why or Why not?
- Suppose the window size right before event D is 16. In the absence of congestion after event D, would the window size ever be larger than 16? Why or why not?

5

- (a) Describe what is happening in this TCP trace at the following times. (One or two word answers will suffice.)

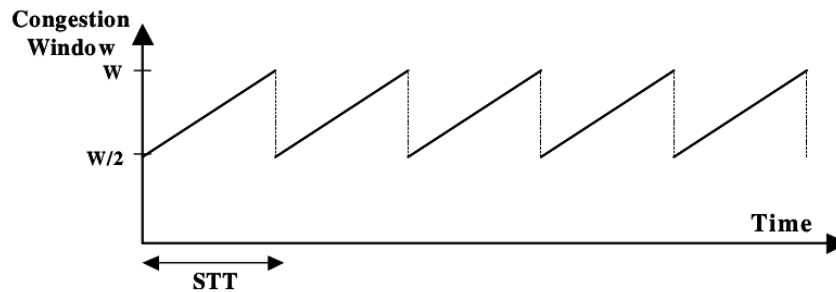


- $t = 0-60$:
- $t = 60$:
- $t = 60-140$:
- $t = 140$:
- $t = 140-180$:
- $t = 180-270$:
- $t = 270$:
- $t = 270-350$:

- (b) Recall that one reason why congestion leads to poor performance is the large queueing delays that are associated with router queues. Would the Internet perform better if router queues were eliminated (or made much smaller)? Why or why not?
- (c) In TCP, each endpoint picks a random initial sequence number (ISN). Why is this done, instead of starting each sequence number at 0?
- (d) RFC 7323 defines a **TCP Window Scale Option** that allows the specification of larger window sizes, up to 2^{30} bytes instead of the normal 2^{16} . When would this option be useful?

6

The picture below shows the famous TCP saw tooth behavior. We are assuming that fast retransmit and fast recovery always work, i.e. there are no timeouts and there is exactly one packet lost at the end of each tooth. We are assuming that the flow control window is large and that the sender always has data to send, i.e. throughput will be determined by TCP congestion control. In the picture, W represents the congestion window size at which a congestion packet loss occurs (expressed in maximum transfer units). You can assume that W is large, so feel free to approximate $(W-1)$ or $(W+1)$ by W . STT represents the saw tooth time expressed in seconds.



In this question, you will calculate the average throughput T for this connection as a function of the roundtrip time (RTT), the maximum transfer unit size (MTU), and packet loss rate p for this connection. Please use the notation suggested by the figure, i.e. W and STT, as intermediate values if you need them

- Express STT in terms of RTT and W .
- Based on the shape of saw tooth, how many units can be sent within each STT?
- Since exactly one packet (unit) loss event happens at the end of each STT, express p in terms of W ?
- TCP throughput (T) is in generally defined as number of bytes divided by the duration of time to send them. Given the results in the previous questions, please calculate the average TCP throughput as a function of RTT, MTU and p .