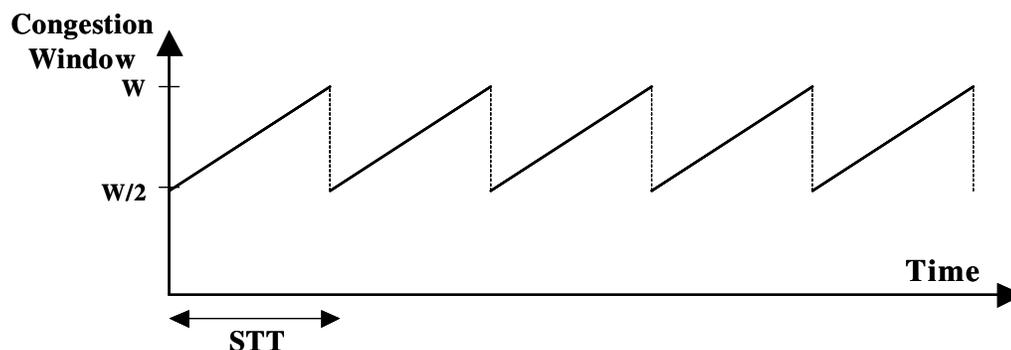


Homework 3

Due: 12 Nov 2015, 11:59pm

Problem 1



The picture above 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.

- 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 PLR for this connection. Please use the notation suggested by the figure, i.e. W and STT , as intermediate values if you need them. Explain each step of your answer. *Hint: Calculate STT and amount of data sent in one STT separately and then divide.*
- The above TCP connection is used in an idle network, and the link capacity of the bottleneck link is B . What throughput will the user observe? Please explain briefly.
- What packet loss rate will the connection suffer? Explain.

Problem 2

The window-scale option of TCP provides a means for supporting a large advertised window. PAWS provides a means for extending the 32-bit sequence number of TCP.

- a. Suppose we have a one-gigabit/second network running cross-country. It has a one-way delay of 60ms. What is the maximum throughput we can attain if we use 16 bits for the window size? Explain. You may assume whatever you'd like for other parameters, such as MTU.
- b. With the window-scale option we can have windows up to $2^{30}-1$ in length (to keep things simple, let's say the maximum window size is 1 gigabyte). With a gigabyte window, what's the maximum throughput we can attain over the network of part 1?
- c. With 32-bit sequence numbers and no PAWS, what problems will we have if we have sustained maximum-speed transfers using the parameters of part 2? Explain, showing an example.

Problem 3

Suppose we have a ten-gigabit/second network with a roundtrip time of 100 milliseconds. The maximum segment size (MSS) is 1500 bytes.

- a. How long will it take for a TCP connection to reach maximum speed after starting in slow-start mode?
- b. Assume there is some probability that a packet is lost in transmission, even though there is no congestion. Give a scenario in which the congestion window becomes 1 MSS (i.e., 1500 bytes) and TCP begins growing the window linearly rather than exponentially (i.e., it's not in slow-start mode).
- c. How long will it take the TCP connection to achieve maximum speed after the scenario of part 2, assuming no further packet loss?
- d. Suppose (exactly) every 100th packet is dropped. Approximate the throughput of TCP on this network. How about every 1000th packet?

Problem 4

We have learned that TCP provides a “reliable” byte stream. Consider a simple client/server pair. The client has 1 kilobyte of data it wishes to send to the server. It calls `connect`, `write`, and `close` appropriately, and each returns without error. Explain how:

- a. The server might not receive the data.
- b. The server might receive corrupted data.

You should assume this is the only TCP connection between the two hosts, ever. More plausible answers are preferred.