15-441: Computer Networks
Homework 4

Assigned: Nov 21, 2006
Due: Dec 7, 2006

1 TCP Congestion Control

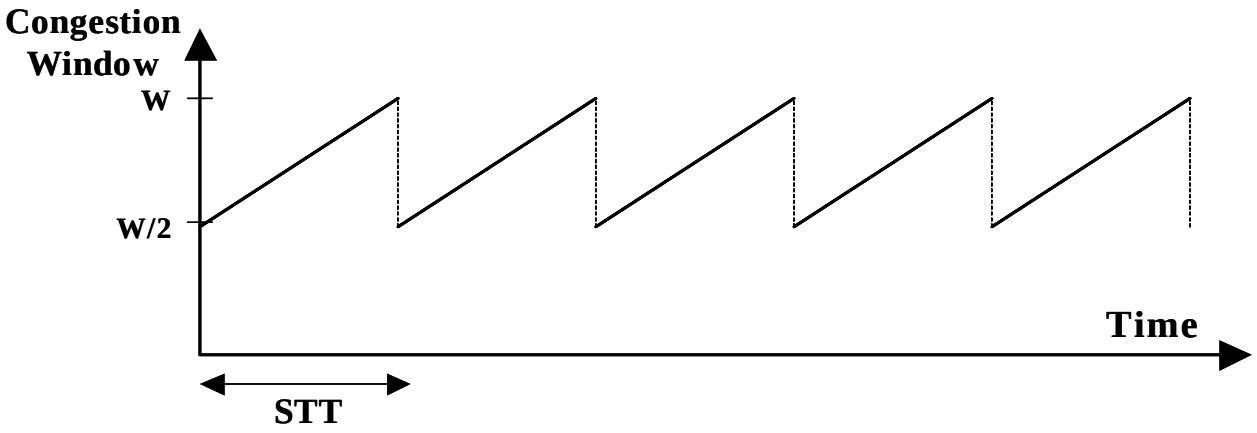

Figure 1: TCP sawtooth diagram

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.

   The aim of this exercise is to derive the average throughput of a TCP connection as a function of the roundtrip time (RTT), the maximum transfer unit (MTU), and the packet loss rate (PLR) for the connection. Please use the notation suggested by the figure, i.e. W and STT, as intermediate values if you need them.

   (a) Calculate the STT as a function of W, and the RTT. (Hint: the congestion window goes from W/2 to W in one STT, and remember the congestion window is increased by 1 MTU every RTT).

   (b) How much data is sent in one STT? (Hint: how much data is sent each RTT?)

   (c) What is the average throughput of the connection?

   (d) What is the average packet loss rate? (Hint: How many losses occur per STT?)

   (e) What is the relationship between the throughput and the packet loss rate?

2 Congestion Window

2. Consider the following plot of TCP window size as a function of time:

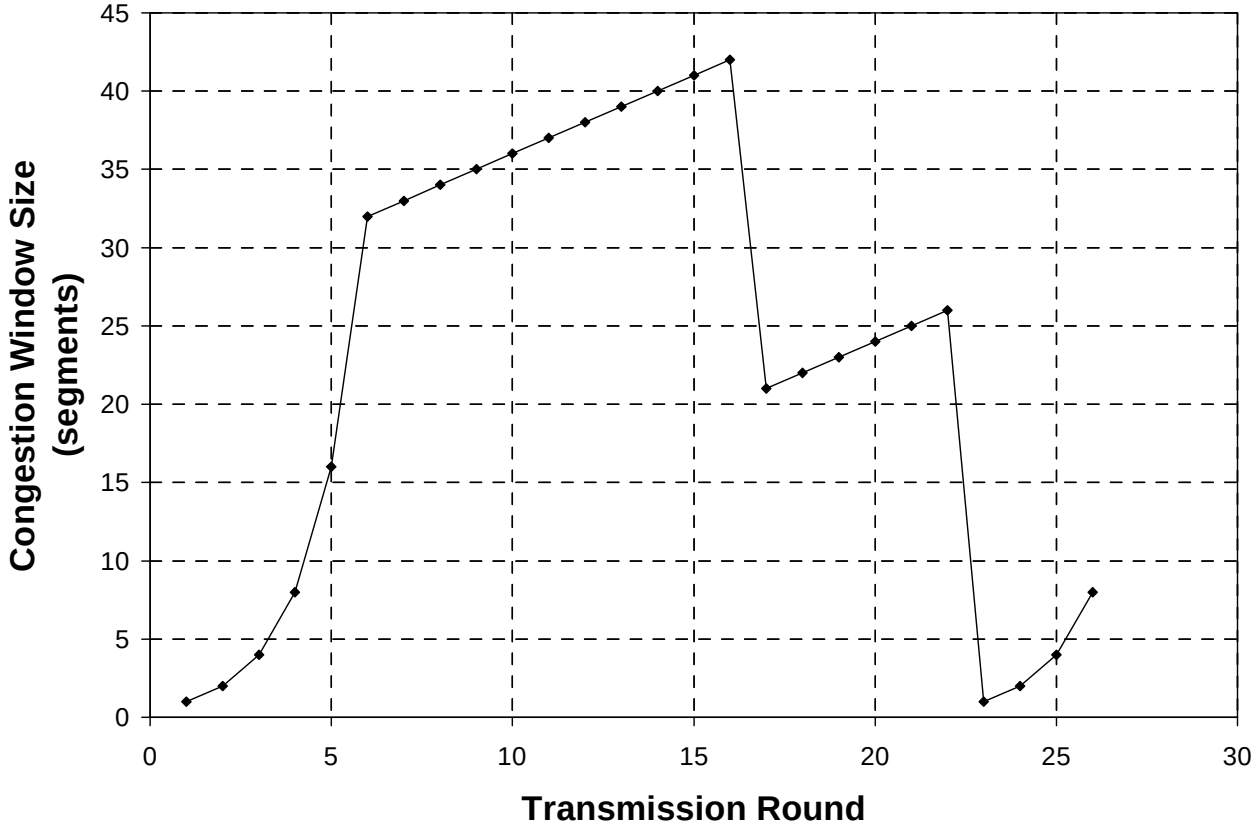

Assuming TCP Reno is the protocol experiencing the behavior shown above, answer the following questions.

(a) Identify the intervals of time when TCP slow start is operating. (2 pts)

(b) Identify the intervals of time when TCP congestion avoidance is operating (AIMD). (1 pt)

(c) After the 16th transmission round, is segment loss detected by a triple duplicate ACK or by a timeout? (2 pts)

(d) What is the initial value of ssthreshold at the first transmission round? (2 pts)

Token Rate	Bucket Size
1	10
2	4
4	1

Table 1: Token bucket parameters

(e) What is the value of ssthreshold at the 18th transmission round? (2 pts)

(f) What is the value of ssthreshold at the 24th transmission round? (2 pts)

(g) During what transmission round is the 70th segment sent? (2 pts)

(h) Assuming a packet loss is detected after the 26th round by the receipt of a triple duplicate ACK, what will be the values of the congesion-window size and of ssthreshold? (2 pts)

3 QoS

3. Consider 10 flows with arrival rates of 1,2,...,10 Mbps that traverse a link of 50Mbps. Calculate the max-min fair share on this link. What is the fair share if the link capacity is 60 Mbps?

4. Suppose a router has accepted flows with the token bucket parameters shown in Table 1. All flows are in the same direction, and the router can forward one packet every 0.1 seconds.
   (a) What is the maximum delay a packet might face?
   (b) What is the minimum number of packets from the third flow that the router would send over 2 seconds, assuming that the flow sent packets at its maximum rate of 4 packets/second uniformly?

4 Security

5. Scenario: U.S. and Russia signed a comprehensive nuclear test-ban treaty. To verify the compliance of the treaty, seismic monitoring systems must be installed to detect any underground testing of nuclear weapons, Such monitoring systems are physically secure (tamper-proof) and are installed in the host nation near certain test sites. The data they collect is transmitted back to the monitoring nation (the host nation can also listen to the transmission), and this data will include timestamps to prevent any replay attacks. Let's assume in this problem the host nation is Russia, and the monitoring nation is the U.S.

   Requirements

   - The U.S. needs to ensure that the data it receives from the monitoring system is not altered/forged.
   - Russia wants to verify in real time that only the seismic data (agreed on in the treaty) is transmitted, i.e. the U.S. is not sending any additional information using the communication channel.
   - If the U.S. finds evidence (from the seismic data) that Russia violates the treaty, it wants to convince the United Nations that such a violation occurred.

   The following verification systems are developed to address the above requirements. (Assume that algorithms required for these approaches exist). We ask you to identify the problems with each of the four approaches. Note that there may be more than one problem with each approach. In this case, if you can identify one "important" problem with an approach, you will get full credit for that approach. If, however, you only identify several "less important" problems, you may not get the full credit. *Hint: Logical flaws in algorithms/protocols are "important", while other problems are "less important". Of course, if there are no logical flaws, other problems become "important".*

   (a) The U.S. installs a secret key along with the monitoring system. When transmitting data back to the U.S., the monitoring system encrypts the transmission using a symmetric cryptographic algorithm and the secret key. How does this approach address the requirements? What are the problems with this approach (if any)?

   (b) The U.S. generates a public/private key pair and installs the private key along with the monitoring system. The public key is made available to everyone. When transmitting data back to the U.S., the monitoring system encrypts the transmission with the private key using public-key cryptography (e.g. the RSA algorithm). How does this approach address the requirements? What are the problems with this approach (if any)?

   (c) The U.S. and Russia each install a subkey along with the monitoring system (but they don't know each other's subkey). The monitoring system then constructs a private key from the two subkeys and also constructs the public key. The public key is made available to everyone, and public-key cryptography is used as in (b). How does this approach address the requirements? What are the

problems with this approach (if any)?

(d) The monitoring system automatically generates a public/private key pair. The public key is made available to everyone, and public-key cryptography is used as in (b). How does this approach address the requirements? What are the problems with this approach (if any)?

5 Firewalls

6. Sparky McFirewall has networks 10.0.2/24, 10.0.4/24 and notices that hosts 10.0.0.32, 10.0.0.252, 10.0.1.54, 10.0.3.2 and 10.0.3.129 are attacking her host 10.0.2.23. He sees that his firewall already contains:

```
deny ip to/from 10.0.0.0/25 to any
```

(a) Which hosts does this rule match?

(b) Which of the following rules would work to block the remaining hosts? (Select only one).

A. deny ip to/from 10.0.0.0/22
B. deny ip to/from 10.0.3.0/24
C. deny ip to/from 10.0.3.2/32, deny ip to/from 10.0.3.129/32
D. deny ip to/from 10.0.0.0/22, deny ip to/from 10.0.3.0/24,
E. deny ip to/from 10.0.0.0/23, deny ip to/from 10.0.3.0/24
F. deny ip to/from 10.0.0.0/24, deny ip to/from 10.0.1.0/23
G. deny ip from brain to paper

(c) After adding in this rule, someone complains that they can't get to a popular web site in the firewalled range. Sparky decides to allow internal users to browse Web sites in the firewalled range. Which rules need to be added? (You may not need all of the spaces below).

Src IP/mask	Src Port	Dst IP/mask	Dst Port	ACK set	Action

(d) A disgruntled employee tells the attacker about your new firewall rules. How could the attacker take advantage of these rules to continue attacking your hosts?