15-744 Computer Networks — Spring 2017
Homework 2
Due by 3/2/2017, 10:30am

(please submit through e-mail to zhuoc@cs.cmu.edu and srini@cs.cmu.edu)

Name:

A Congestion Control

1. At time t, a TCP connection has a congestion window of 4000 bytes. The maximum segment size used by the connection is 1000 bytes. What is the congestion window after it sends out 4 packets and receives ACKs for all of them...

   (a) If the connection is in slow-start?

   > **Solution:** 8000 bytes. In slow start, the sender increases its window for each byte successfully received.

   (b) If the connection is in congestion avoidance (AIMD mode)?

   > **Solution:** 5000 bytes. The sender increases its window by one segment each window.

2. In congestion avoidance (CA) mode, a TCP sender increases the size of its congestion window by one *maximum segment size* (MSS) each RTT. Suppose a TCP implementation does this by increasing its congestion window by Δ **every time it receives an ACK**, where

$$\Delta = \frac{\text{MSS}}{\text{current window size}} \cdot \text{MSS}$$

   Knowing this, how can a greedy TCP receiver get more than its fair share of the link bandwidth? (*Hint:* Remember that the sequence number acknowledged by an ACK doesn't refer to a *packet*, but rather to a *byte* in the data stream.)

   > **Solution:** The greedy receiver can send more than one ACK per MSS. For example, if the first packet from the sender contains bytes $0 - 999$, the receiver might send two ACKs: one for byte 500 and one for byte 1000. Since the sender increases its window by Δ **for every ACK**, it increases its window by 2Δ when it should only increase by Δ.

3. RED gateways can optionally be run in "byte mode," where the average queue size is measured in bytes instead of packets. Would the "SSH" application prefer routers to operate in byte mode or not? Why?

Solution: Yes. In byte mode, RED gateways take into account packet size when calculating the probability p that a packet is marked (i.e., dropped):

$$p = \frac{max_p(avg - min_{th})}{max_{th} - min_{th}} \times \frac{\text{Packet Size}}{\text{Maximum Packet Size}}$$

The SSH packets (which are small) are less likely to be dropped in the face of incipient congestion.

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

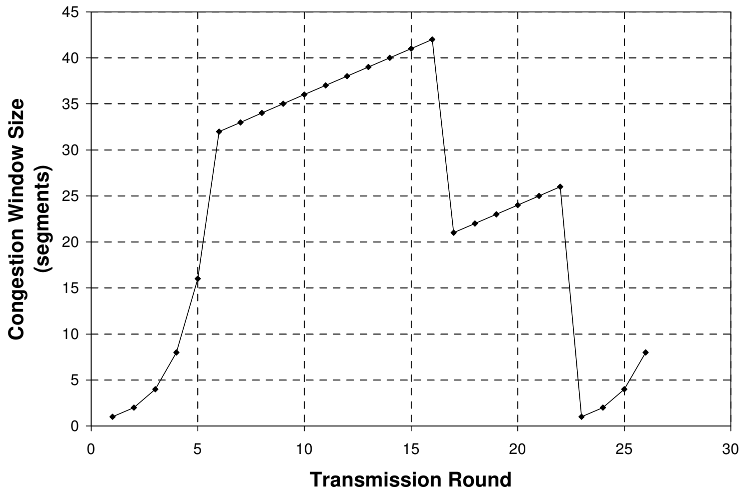

Assuming TCP Reno is the protocol experiencing the behavior shown above, answer the following questions.

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

> **Solution:** Solution: 1-6, 23-26

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

> **Solution:** 6-23

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

> **Solution:** dupack

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

> **Solution:** 32

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

> **Solution:** 21

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

> **Solution:** 13

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

> **Solution:** 7

(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?

Solution: 4, 4

B Fair Queuing

5. A network uses routers with fair queuing.

   (a) Two connections share the same congested gateway and have no other congested gateways. Connection A has RTT 5ms, connection B has RTT 10ms. Express the throughput of connection B ($tput_B$) in terms of the throughput achieved by connection A ($tput_A$), or indicate if there is no relationship between the two.

   Solution: $tput_B = tput_A$

   (b) Two connections traverse the same congested gateway, but *also* traverse some other unshared congested gateways. Express the throughput of connection B ($tput_B$) in terms of the throughput achieved by connection A ($tput_A$), or indicate if there is no relationship between the two.

   Solution: Without knowing the other flows in the network, you cannot find a relationship between the two.

6. Suppose that a router has three input flows and one output flow. It receives the packets listed in the Table below, all at about the same time, in the order listed, during a period in which the output port is busy but all queues are otherwise empty.

Packet	Size	Flow
1	100	1
2	100	1
3	100	1
4	100	1
5	150	2
6	200	2
7	80	3
8	30	3
9	120	3

All three flows share the same outbound link, on which the router can transmit one packet per time unit. Assume there is an infinite amount of buffer space.

Give the order in which the packets are transmitted, assuming:

(a) Round robin

Solution: 1, 5, 7, 2, 6, 8, 3, 9, 4

(b) Fair queuing

Solution: 7, 1, 8, 5, 2, 9, 3, 6, 4

(c) Suppose we use *weighted fair queuing* with flow 2 having weight 2 and the other two flows having weight 1. This means flow 2 will be assigned twice throughput compared to other flows.

Solution: 5, 7, 1, 8, 6, 2, 9, 3, 4

C Data Center Networking

7. Consider a simplified data center topology as shown below. Assume it uses the PMAC address proposed in "PortLand" to do routing.

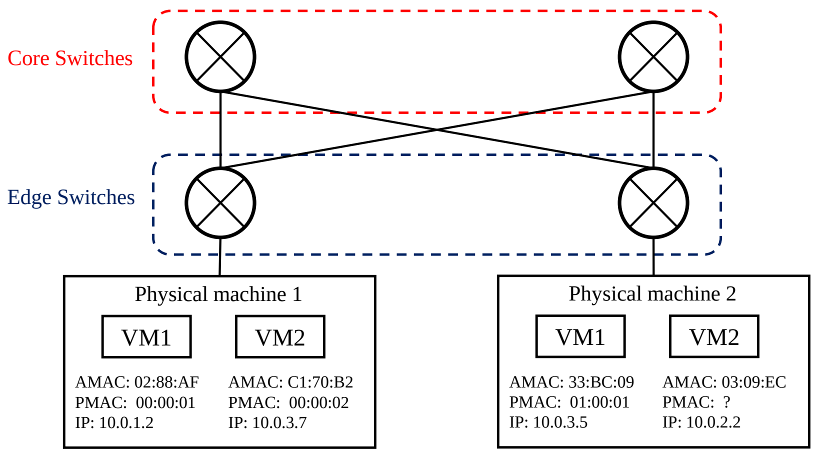

(a) Given the PMAC of some of the VMs shown in the figure, what will the PMAC of VM2 in physical machine 2 be?

> **Solution:** 01:00:02

(b) Suppose VM1 in physical machine 1 is now migrated to physical machine 2. What will its AMAC, PMAC and IP address be?

> **Solution:** AMAC will remain 02:88:AF. IP address will remain 02:88:AF. PMAC will likely be 01:00:03.

8. DCTCP (Data center TCP) tries to react to the extent, but not the presence, of congestion in the network.

(a) Assume the sender has window size of N at time t. During time t to $t + RTT$, 20% ACK packets it receives have ECN bit set. How would the window size in TCP and DCTCP react respectively?

> **Solution:** In TCP, the window size will drop by half. In DCTCP, it will update α first using equation $\alpha = (1 - g) * \alpha + g * F$ where F is 20%. Then it will update window size by $cwnd = cwnd * (1 - \alpha/2)$. If the 20% fraction is relatively stable (or g is close to 1), then α is close to 20%, then window size will drop by about 10%.

(b) What will happen if the marking threshold K is set to 0?

> **Solution:** Then all the packets will have ECN bit set. α will always be one. Window size will be zero.

(c) What will happen if the marking threshold K is set to the same as buffer size?

Solution: Then no packets will have ECN bit set. α will always be zero. DCTCP client will act as a normal TCP client.

D Basic Tools

9. Take a look at the man pages for `netstat` to answer the following questions.

(a) What does the command netstat -a show you? Explain the two parts of the output.

> **Solution:** This command shows you all the active network connections. The ”-a” flag tells netstat to display all connections, including those that are waiting for inbound connections (like TCP sockets in the LISTEN state fore example). The top half of the display shows the Internet domain activity while the bottom half of the display shows the Unix domain activity. The State field displays the state of the TCP connection (w.r.t. the TCP state machine shown in class, i.e. LISTEN, ESTABLISHED, FIN_1,..).

(b) What is the command to view the routing table of your machine using netstat? What is the command to only show IP addresses and not host names in the routing table?

> **Solution:** netstat -r, netstat -rn

(c) How can you use netstat to find out what the network interfaces of your machine are? What is the MTU of your Ethernet interface?

> **Solution:** netstat -i, 1500

(d) Start an FTP connection using the command ftp (e.g., FTP to ftp://gnu.mirror.iweb.com). After you login as anonymous, try to find information regarding the corresponding TCP connection using `netstat` in a different window. Explain the fields in the line corresponding to your ftp connection. What are the local and remote port numbers and IP addresses for that TCP connection? What is the dedicated port used by FTP?

10. In this problem, you will get experience using wireshark to do real network packet analysis. Packet traces are useful for debugging and understanding the packet-level behavior of network protocols, among other things.

You can download wireshark (http://www.wireshark.org) and install it on your local Unix machine. You may have to run the program with administrator privileges to obtain access to the network interfaces.

We would like you to do the following:

- Run wireshark and be able to capture network traffic.
- Capture the download of any suitably large file. You may use any file, but the file should take at least 5 seconds to download.

Based on the resulting packet capture:

(a) Generate a TCP sequence plot based on the traffic generated by downloading the file. Highlight where losses occur during the transfer.

(b) Generate a packet delay plot, showing the per-packet delay as a function of the sequence number.