CS557
Spring 2011 Final Exam

PROF. DAN MASSEY

Take Home Exam
Due By 3pm on Thursday 5/12/11

- This exam has 11 pages, including this cover page. Do all your work on these exam sheets, use the backs of the pages if needed.
- Be specific and clear in your answers.
- Show *all* your work if you wish to be considered for partial credit.
- You must hand in hard copy of your exam in directly to Dr. Massey or to the Computer Science front desk (2nd floor).

Question	Points	Score
1	15	
2	15	
3	15	
4	10	
5	20	
6	10	
7	15	
Total:	100	

Name: ____________________________

EID: ________________________

<u>IMPORTANT:</u> This exam contains several problems where you must fill in a table entry. If you fill in the entry correctly, you will receive full credit for that entry. But if you fill in a table incorrectly, then you will be eligible for partial credit ***<u>only if you explain how came up with the table entries.</u>***

<u>Students who simply fill in the tables and provide no explanation of how the tables were derived will receive no partial credit if the entries are incorrect.</u>

The number of rows or columns in table does not necessarily match the number of rows or columns needed in the answer. You may need to add additional rows or columns. You may not need to use all rows or columns.

Problem 1 (15 points)

Suppose BGP router Z has learned the following potential paths to destination A.

Initial Routing Table at Node Z:

Path from neighbor B = B,A
Path from neighbor C = C,B,A
Path from neighbor D = D, E, B, A
Path from neighbor F = F, G , H, I, A

At time 0 seconds, Z selects path (B,A) as the best path and announces this path to neighbor X.
At time 1 seconds, link (B,A) fails and Z receives the following updates
At time 2 seconds: B withdraws path (B,A)
At time 35 seconds: C withdraws path (C,B,A)
At time 45 seconds: D replaces path (D,E,B,A) with path (D,C,B,A)
At time 68 seconds: withdraws path (D,C,B,A)

In all of the problems, the following damping parameters are used:

Event	Value
Penalty for an announcement	500
Penalty for a withdraw	750
Half Life	2 hours
Supression Threshold	1500
Reuse Threshold	300

A) (5 points) Assume that routers use standard BGP with a 30 second MRAI timer. List the times at which nod Z sends updates to X, list the path sent to node X, and list the damping penalty that X associates with pair (Z,A).

Time (in seconds)	Path sent from Z to X	Penalty at X for pair (Z,A)
0	Z,B,A	500

Will X damp the route from Z? If yes, when will X reuse this route?

B) (5 points) Assume that routers use Ghost Flusing BGP with a 30 MRAI timer. List the times at which nod Z sends updates to X, list the path sent to node X, and list the damping penalty that X associates with pair (Z,A).

Time (in seconds)	Path sent from Z to X	Penalty at X for pair (Z,A)
0	Z,B,A	500

Will X damp the route from Z? If yes, when will X reuse this route?

C) (5 points) Assume that routers use BGP+RCN with a 30 MRAI timer. Further assume that initially link (A,B) has an RCN sequence number of 2 and all updates received at node Z have an RCN = “Link (A,B) failed with sequence number 3”. List the times at which nod Z sends updates to X, list the path sent to node X, and list the damping penalty that X associates with pair (Z,A).

Time (in seconds)	Path sent from Z to X	Penalty at X for pair (Z,A)
0	Z,B,A	500

Will X damp the route from Z? If yes, when will X reuse this route?

Problem 2 (15 points)
Given the following routing tables, use the ***advanced*** "Routing With a Clue" method to fill in the receiver's clue table. See Figure 3 in the paper for an example of how to correctly fill in a table.

Sender's Routing Table:

Prefix	Next Hop
*	Receiver
001*	Receiver
0001*	Receiver
011*	Receiver
1*	Receiver
100*	Receiver
1001*	Receiver

Receiver's Routing Table (in other words, its Trie):

Prefix	Next Hop
*	R5
00*	R6
0000*	R7
0110*	R8
11*	R9
10*	R10
1001*	R11

Clue Table at Receiver:

Clue	FD	Pointer
*		
001*		
0001*		
011*		
1*		
100*		
1001*		

If you would like to be eligible for partial credit, explain how you came up with the above table entries.

Problem 3 (15 points)

These problems use the fairness definitions and fair queuing algorithm discussed in "Analysis and Simulation of a Fair Queueing Algorithm [DKS89]".

A) (5 points) Suppose 3 flows are sharing a single 12 Mbps link. Flow 1 requests 2 Mbps. Flow 2 requests 20 Mbps. Flow 3 requests 80 Mbps. In a fair assignment, how much bandwidth should be assigned to each flow?

Flow 1 should receive _______ Mbps.
Flow 2 should receive _______ Mbps
Flow 3 should receive _______ Mbps

B) (10 points) Suppose 4 flows are sharing a single 1 bps link. Each flow has 4 packets that are queued and waiting to be sent:

Flow Number	Packet 4 (End of Queue	Packet 3	Packet 2	Packet 1 (Front of Queue)
Flow 1	100 bit packet	30 bit packet	20 bit packet	70 bit packet
Flow 2	20 bit packet	100 bit packet	20 bit packet	40 bit packet
Flow 3	40 bit packet	40 bit packet	40 bit packet	40 bit packet
Flow 4	10 bit packet	80 bit packet	20 bit packet	100 bit packet

At time 0, some packet is selected as the first packet to send. After all bits in the first packet are sent, a second packet is selected and all bits in that packet are transmitted. The process continues until all 16 packets have been sent. A fair queuing algorithm is used to select the packet order. Fill in the completion time (the time when the last bit of the packet has been transmitted) for each packet.

Flow Number	Packet 4	Packet 3	Packet 2	Packet 1
Flow 1				
Flow 2				
Flow 3				
Flow 4				

If you would like to be eligible for partial credit, explain how you came up with the above table entries.

<u>Problem 4</u> (10 points)

Suppose a TCP session starts with an initial congestion window of 1 and threshold of 8. Define an epoch as the time it takes to send a window's worth of packets and receive the acks for all packets that arrived during the epoch.

During the course of the TCP session, a triple duplicate ACK occurs after the 3rd epoch and a timeout occurs after the 12th epoch. Calculate the TCP congestion window size and threshold size for each epoch. Also identify whether the connection is in slow start (SS) or congestion avoidance (CA). (the first two rounds have been filled in for you.)

Epoch	Congestion Window Size	Threshold Size	Slow Start (SS) or Congestion Avoidance (CA)
1	1	8	SS
2	2	8	SS
3			
Triple Duplicate ACK	Triple Duplicate ACK	Triple Duplicate ACK	
4			
5			
6			
7			
8			
9			
10			
11			
12			
Timeout	Timeout	Timeout	
13			
16			
17			
18			
19			
20			

If you would like to be eligible for partial credit, explain how you came up with the above table entries.

Problem 5 **(20 points)**

Consider the following network from [ZSC91]:

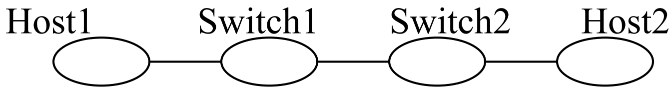

You may assume the following:

1) Host1 is sending data packets to Host2 and Host2 is sending data packets to Host1
2) The Host1-Switch1 and Host2-Switch2 links run at 10 Mbps
3) The Switch 1- Switch2 link runs at 50Kbps.
4) Data packets (in both directions) are 500 bytes; ack packets are 50 bytes.
5) The window size on both Host1 and Host2 is fixed.
6) The queues at Switch 1 and Switch 2 have are infinitely big (they never drop packets).

The graph below shows the resulting queue behavior at Switch 2.

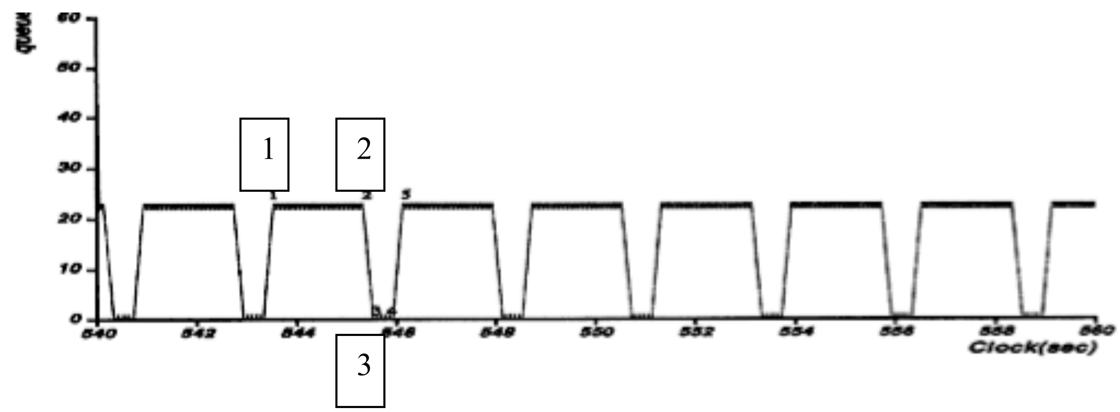

a) (10 points) Consider the queue behavior between time 1 and 2.

a. Which packets are arriving in the queue?

b. Which packets are leaving the queue?

c. During this time, the queue remains roughly fixed at around 22 packets. Why does the queue size remain roughly fixed?

b) (10 points) Consider the queue behavior between time 2 and 3

a. Which packets are arriving in the queue?

b. Which packets are leaving the queue?

c. During this time, the queue rapidly drops from 22 packets to 0 packets. Why does the queue size drop quickly?

<u>Problem 6</u> (10 points)

Consider the following network:

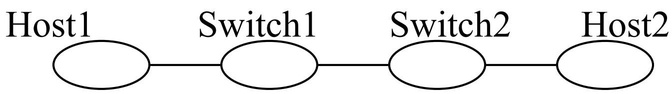

You may assume the following:

7) Host1 is sending data packets to Host2 and Host2 is sending data packets to Host1
8) The Host1-Switch1 and Host2-Switch2 links run at 10 Mbps
9) The Switch 1- Switch2 link runs at 50Kbps.
10) Data packets (in both directions) are 500 bytes; ***ack packets are also 500 bytes***.
11) The window size on both Host1 and Host2 is fixed.
12) The queues at Switch 1 and Switch 2 have are infinitely big (they never drop packets).

A) (5 points) Will clustering occur in this scenario? Note that unlike the Two-Traffic paper, the ACK size is equal to the packet size in this problem. Explain your answer. No credit will be given without a detailed discussion of why clustering does (or does not) occur.

B) (5 points) Will ACK compression occur in this scenario? Note that unlike the Two-Traffic paper, the ACK size is equal to the packet size in this problem. Explain your answer. No credit will be given without a detailed discussion of why ACK compression does (or does not) occur.

<u>Problem 7</u> (15 points)

The following entries are present in the "com" DNS zone.

bar.com.	NS	a.bar.com.
bar.com.	NS	b.bar.com.
bar.com.	NS	a.foo.com.
foo.com.	NS	a.bar.com.
foo.com.	NS	b.bar.com.
foo.com.	NS	a.foo.com.
a.bar.com.	A	2.2.2.2

You may assume there are no other NS or A records listed in the com zone. Resolvers always pick the first NS record listed in a NS record set and servers always rotate the order after each query for the NS record set For example, the first query for (foo.com, IN, NS) will return records (a.bar.com, b.bar.com, and a.foo.com) in that order. The query for (foo.com, IN, NS) will return records (b.bar.com, a.foo.com, and a.bar.com,) in that order. And so forth.

A) (5 points). If the server a.foo.com fails (stops responding to queries for any reason), will DNS data from foo.com (such as www.foo.com) still be available? Why or why not.

B) (5 points). What is the minimum number of server failures needed to make the foo.com DNS data unavailable? Identify which servers must fail to achieve this minimum.

C) (5 points). If 90 different resolvers each attempt to lookup www.bar.com, what will be the expected query load on each bar.com server?

_______ queries will be sent to server a.foo.com

_______ queries will be sent to server a.bar.com

_______ queries will be sent to server b.bar.com

Explain how you determined the expected number of queries.

D) (EXTRA CREDIT: 5 points). Over a period of 5 days, suppose one resolver sends 90 queries for www.bar.com. You may assume that NS records and their associated A records have a TTL value of 7 days. The www.bar.com A record has a TTL of only 1 second. Will the query load be different than in part c? Why or why not?