CS 555: Computer Networks and Data Communication
Spring 2006
Final Examination
Points: 150
May 3, 2006 (3:45-6:45 PM)
Time allowed: 180 minutes
CLOSED BOOK, CLOSED NOTES, OPEN MIND
Answer All Questions

Turning in this exam under your name confirms your continued support for the honor code of Old Dominion University and further indicates that you have neither received nor given assistance in completing it.

Name: _____
Unix login : _____@cs.odu.edu

Question #	Possible points	Obtained points
1	40	
2	35	
3	40	
4	35	
Total	150	

SHOW ALL YOUR WORK

Question 1:

(a) An application produces a 10 Mbyte burst once in a second. The application lasts for 5 seconds. The first burst occurs at time 0. The network bandwidth allocated to this application is 50 Mbps. Suppose a 25 Mbyte leaky bucket with an output bucket rate of 5 Mbytes/sec is used to shape the traffic, show the input and output as a function of time. Also, indicate whether or not there was data loss. If so, which part of the input was lost?

(b) Suppose an organization has been allocated an IP address space indicated as 195.13.8.0/23 (i) What is the range of IP addresses assigned to it? (ii) How many IP addresses have been assigned?

(c) Suppose the Ethernet card of a computer (IP address 195.13.8.123) has been replaced with a new card. Is there any impact of this change on other computers (consider same LAN and different LAN cases separately) that have been communicating with this computer? If so, how is it resolved?

WORK AREA (Question 1):

Question 2:

(a) An application produces a 10 Mbyte burst once in a second (1 M=10^6). The application lasts for 2 seconds (i.e., two bursts). The first burst occurs at time 0. The network bandwidth allocated to this application is 50 Mbps. Suppose a token bucket with a capacity of 1 M tokens (1 token = 10 bytes) and an input token rate of 0.5 M tokens/sec is used to control traffic to network. Suppose the token bucket was full at time 0, show the input and output of the bucket as a function of time. ($S=C/(M-\rho)$)

(b) Suppose an organization with an IP address of 197.24.23.45 has 1024 computers, each of which need an IP of their own. Its system administrator has assigned the reserved IP range of 10.0.0.0-10.0.3.255 to its computers. (i) How does the computer 10.0.2.25 (using port 4500) handle a session with another external computer (IP address of 197.3.4.67 at port 3423)? (ii) How does the external computer communicate with the internal computer? (Hint: Indicate one method handling the cases providing the needed details)

(d) Why is the time to live field in IP header used as the number of hops instead of the physical time (say in milliseconds).

WORK AREA (Question 2):

Question 3:

(a) Suppose host H1 (with IP1 and Port1) wants to make a connection with H2 (IP2 and Port2) and sends a request to it with SEQ#=300. At the same time host H2 (with IP2 and Port2) sends a similar request to H1 (IP1 and Port1) with SEQ#=2000. Show the sequence of messages with appropriate parameters that are exchanged until the data flows in both directions. Assume that TCP is being used as a protocol.

(b) A client C at host H1 has established a connection with server S at H2. C sends a request of 10-bytes each second. The server reads the data as soon as it arrives. For each 10-byte request, server sends a 500-byte response. The client reads the reply at a slow speed of 200 bytes/second. If the connection lasts for 3 seconds, what should be the **minimum** window sizes at C and S so neither has to wait for lack of window size? (Client's sending requests are independent of the received replies.) All other delays are negligible.

(c) Consider the timeout maintained by H1 for its H1-H2 TCP connection. At time 0, it has the following estimates: RTT=15, D=10, $\alpha = 0.6$; (i) At t=0, it sends TPDU1 for which an acknowledgement is received at t=40; (ii) At t=45, it sends another TPDU2 for which it received an acknowledgement at t=75; Determine the timer values set for TPDU1 and TPDU2.

WORK AREA (Question 3):

Question 4:

(a) Consider a TCP connection between H1 and H2. At t=0, the threshold parameter was 8 Kbytes, the receiver window was 24 Kbytes, and the maximum segment size was 2 Kbytes. (i) The 1st transmission takes place at t=0 and was successful; (ii) Subsequently, 2nd and 3rd transmission were also successful (iii) The 4th transmission resulted in a timeout (iv) The 5th and 6th transmissions were successful. Determine the size of the congestion window at each of the 6 transmissions.

(b) Consider the TCP connection between hosts H1 and H2 under the following two conditions: (i) H1 and H2 are connected by a 100 meters line (propagation speed = 200 km/millisecond) with100 Mbps data rate (ii) H1 and H2 are separated by 3000 km and connected by a line with a data rate of 10 Mbps and propagation speed of 200 km/milliseconds. What is the bandwidth-delay product (in bits) for these two scenarios? Assuming that there is no congestion, what is the minimum size of the receiver buffer that each TCP host must offer the other so that the utilization is maximum? (There are no other errors or delays in the system. BW-delay product=BW*RTT)

(c) You are the systems administrator of a company with several networks and hundreds of computers. Some of users have complained of the low throughput they are recently noticing from the database server that is accessed over the net. Clearly state the procedure that you would adopt to diagnose the problem. Make sure to indicate the possible bottlenecks that may be the cause of this problem.

WORK AREA (Question 4):