EE 122 Fall 2001 1st Midterm

Professor: Lai Stoica

Question 1 (15 pt)
Layering is a key design principle in computer networks. Name **two** advantages, and **one** disadvantage to layering. Explain. Use no more than three sentences to describe each advantage or disadvantage.

Answer 1
Any two of the advantages and any of the disadvantages listed below:

Advantages:
(1) modularity: protocols easier to manage and maintain
(2) abstract functionality: lower levels can be changed without affecting upper layers
(3) reuse: upper layers can reuse the functionality provided by lower layers

Disadvantages:
(1) information hiding: make hard to optimize higher level protocols. For example, IP does not reveal the exact cause of dropping a packet, which makes it very difficult for the TCP to work across different communication media. As a result TCP works well on top of wired networks, but not on top of wireless networks.

(2) inefficiency: each layer adds his own processing overhead; some applications may not need some functionality provided by a particular layer, case in which that functionality is pure overhead (e.g., TCP reliability for video/audio traffic).

Question 2 (20 pt)
Let l = 1000 bit be the minimum size of an Ethernet packet. Let c=2*10^8 m/sec be the speed of light in the Ethernet cable. The Ethernet has a bandwidth of 10 Mbps. What is the maximum distance d between two end-hosts on an Ethernet cable? Justify your answer.

Answer 2
Suppose hosts A and B are at the opposite ends of the Ethernet segment, and that host A sends a packet of minimum size l to B at time t. Host B receives A's packet at time $t+d$. Suppose that just before it receives the packet, host B sends its own packet. This packet reaches A at time $t+2d$. At this point A should still be in the process of sending the packet; otherwise it will not be able to detect the collision. Thus, the transmission time of the packet, l/B, should be such that

l/B = 2*d;

where B is the Ethernet bandwidth. Let L be the distance between A and B. Then, we have

l/B = 2*L/c --> L = l*c/(2*B) = 10^3 *2 * 10^8 / (2*10^7) = 10,000 m

Question 3 (20 pt)
(a) Explain the TCP slow-start. What is its goal and how does it work? (8pt)

(b) Assuming that the available link capacity and the receiver window are infinite how many round-trip times does it take in TCP to send the first 10 packets? In general, how many round-trip times does it take to send the first *k* packets? (12pt)

Answer 3
The goal of the TCP slow-start is to find the available bandwidth on the congested link as soon as possible. With the TCP slow-start, a sender starts with a congestion window (cwnd) of 1, and then increments it by 1 each time it receives an acknowledgement. This results in effectively doubling cwnd every round-trip time RTT.

After 2 RTTs the TCP flow sends 1+2+4=7 packets. After the third RTT, the TCP's cwnd becomes 8. As a result, the TCP will send the remaining 3 packets. THus the TCP needs 3 RTTs to send 10 packets by using the slow-start protocol. In general, given k packets, it takes m RTTs to send them, where m is such that

$1+2^1 + 2^2 + ...2^{(m-1)} < k <= 1 + 2^1 + 2^2 + ...2^m$,
$(2^m) - 1 < k <= 2^{(m+1)} - 1$,
$m-1 < \log2(k+1) -1 <= m$,
$m = \text{ceiling}(\log2(k+1)-1) = \text{ceiling}(\log2(k+1)) - 1$
where ceiling(x) denotes the smallest integer that is greater or equal to x.

Question 4 (15 pt)
The sliding window algorithm is used in TCP for error control and flow control. Due to the finite space of sequence numbers, even with infinite link speed and host processing power, the maximum throughput between two hosts is still limited. What is the maximum throughput between two TCP hosts with a round-trip time of 10ms?

(Note: Recall that the TCP sequence numbers are 32-bit long.)

Answer 4
MaximumWindowSize <= (MaxSeqNum+1)/2 (see textbook, Sec. 2.2, pages 108-109)

Throughput = MaximumWindowSize/RTT = (MaxSeqNum+1)/(2*RTT) =
(2^32 + 1) bytes / (2*10ms) = 1.6 * 10^12 bps approx.

Note: This problem ignores the advertised window size, which is only 16 bits (though it is possible to extend it; see page 395 in textbook). Solutions that have considered the advertised window the limiting factor received full credit.

Question 5 (20 pt)
Consider a sliding window flow control protocol between hosts A and B. Suppose the link propagation delay is 1 time unit, the retransmission timeout is 3 time units, and the window size is 3. Assume the link drops every third data packet, i.e., the link drops the 1st, 4th, 7th, ... data packets. (Note that here "kth packet" means the kth packet transmitted on the link, and **not** the sequence number of the packet.) How long does it take to transmit 6 packets between A and B?

(Note: Ignore the transmission times and the queueing delay, and assume that no acknowledgements are lost.)

Answer 5
(a) represents the time diagram when the receiver sends **only** 6 data packets;
(b) represents the time diagram when the receiver sends more than 6 data packets. Both solutions received full credit. (In addition, due to some apparent misunderstandings during the exam, solutions that assumed that the window size represents the number of un-acknowledged packets at the receiver were also considered, despite the fact that this assumption is wrong.)

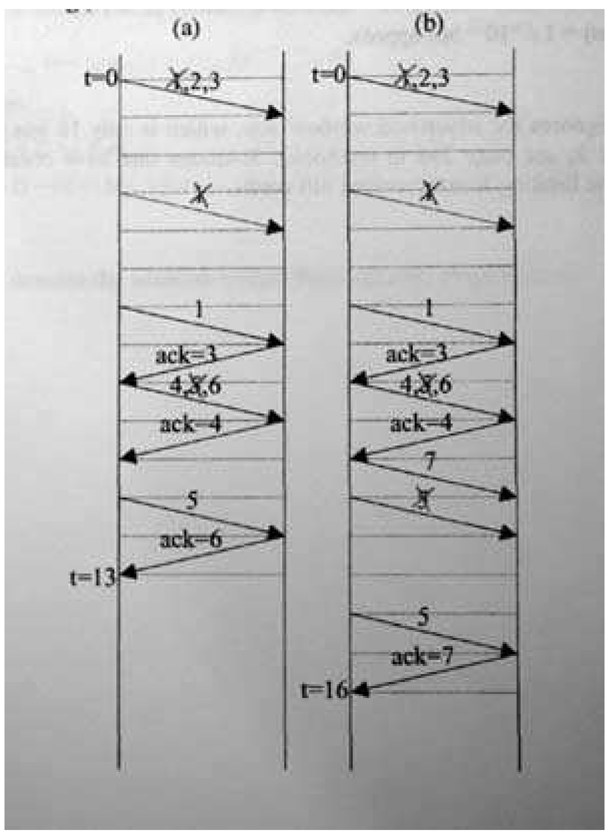

Question 6 (20 pt)
Token ring is an alternative to Ethernet to arbitrate the access to a shared medium. With token ring, a special packet (called "token") is circulated in the clock-wise direction around a ring to which N hosts are connected. When it receives the token, a node can send one packet, if it has any packets to send. After it sends a packet, the node passes the token to the next node on the circle. If it doesn't have any packets to send, the node passes the token immediately to the next node.

Consider the following scenario based on the figure below. Initially, A has the token and has a packet to send. Thus, A sends its packet and then passes the token to B. Suppose B has several packets to send when it receives the token. B sends only one of its packets and then passes the token to C. Suppose C does not have any packet to send, when it receives the token. As a result, C passes the token immediately to D. If D has at least a packet to send when it receives the token, D sends one packet and passes the token to A. This process repeats indefinitely.

Suppose the propagation delay along the ring is d, and that the time to process the token is negligible. Thus, when no node has packets to send, it takes the token time d to make a complete round. Let n be the number of hosts connected to the ring that have packets to send, let C be the link capacity, and let l be the packet size (assume all packets have the same size). What is the utilization of the token ring?

(Note: Define the utilization as the ratio between the maximum total number of bits that were sent by all hosts during one round and the total number of bits that can be carried by the ring during one round, i.e., C*TRound, where TRound is the time it takes the token to complete a round).

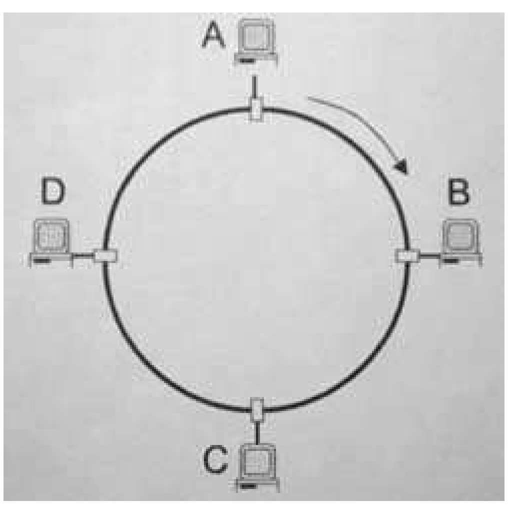

Answer 6
According to the note, we have

utilization = $n*l/C*TRound$

The time it takes to transmit a packet is l/C. Thus, when all n hosts have packets to send,

$TRound = d + n*(l/C)$,

utilization = $(n*l)/(C*d + n*l)$