University of California at Berkeley
College of Engineering
Department of Electrical Engineering and Computer Science

EECS 122 I. Stoica
Fall 2009

FIRST MIDTERM EXAMINATION
Monday, 19 October 2009

INSTRUCTIONS—READ THEM NOW! This examination is CLOSED BOOK/CLOSED NOTES. There is no need for calculations, and so you will not require a calculator, Palm Pilot, laptop computer, or other calculation aid. Please put them away. You MAY use one 8.5" by 11" double-sided crib sheet, as densely packed with notes, formulas, and diagrams as you wish. The examination has been designed for 80 minutes/80 points (1 point = 1 minute, so pace yourself accordingly). All work should be done on the attached pages.

In general, if something is unclear, write down your assumptions as part of your answer. If your assumptions are reasonable, we will endeavor to grade the question based on them. If necessary, of course, you may raise your hand, and a TA or the instructor will come to you. Please try not to disturb the students taking the examination around you.

We will post solutions to the examination as soon as possible, and will grade the examination as soon as practical, usually within a week. Requests for regrades should be submitted IN WRITING, explaining why you believe your answer was incorrectly graded, within ONE WEEK of the return of the examination in class. We try to be fair, and do realize that mistakes can be made during the regarding process. However, we are not sympathetic to arguments of the form "I got half the problem right, why did I get a quarter of the points?"

______________________________ SID: ________________________
(Signature)

______________________________ Discussion Section (Day/Time): ________
(Name—Please Print!)

QUESTION	POINTS ASSIGNED	POINTS OBTAINED
1	10	
2	10	
3	15	
4	20	
5	10	
6	20	
7	15	
TOTAL	100	

Question 1. *Miscellaneous (10 points)*
For each of the following statements, indicate whether the statement is True or False, and provide a very short explanation of your selection (2 points each).

a. Time division multiplexing allows a connection to use unused slots of another connection. T F
Rationale:

False. Unused slots cannot be used by other connection.

b. In datagram switching networks two packets of the same flow always take the same route. T F
Rationale:

False. Each packet is forwarded independently.

c. Sliding window achieves higher throughput than Stop-and-Go. T F
Rationale:

True. Sliding window allows more in-the-fly packets during the same RTT.

d. Flow control slows down the sender when the network is congested. T F
Rationale:

False. Flow control slows down the sender when the receiver is slow.
(Congestion control slows down the sender when the network is congested.)

e. Multiplexing a large number of flows reduces burstiness. T F
Rationale:

True. As number of flows increases, aggregate_peak/aggregate_avg decreases, as there less chance that the peak utilization of all flows coincide.

Question 2. *Packet transmission (10 points)*
For the topology shown below, assume that every link has a transmission speed of 1000Mbps and a length of 250m. Assume there is no queuing or processing delay. Further, assume a propagation speed of $2*10^8$m/s.

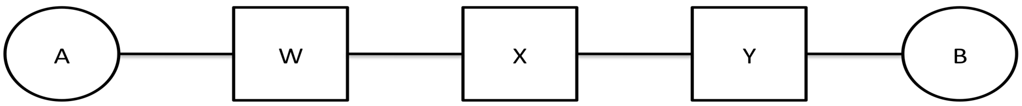

a) (5 points) How long will it take for a 125B packet to travel from node A to node B, assuming routers implement cut-through packet forwarding?

Solution:

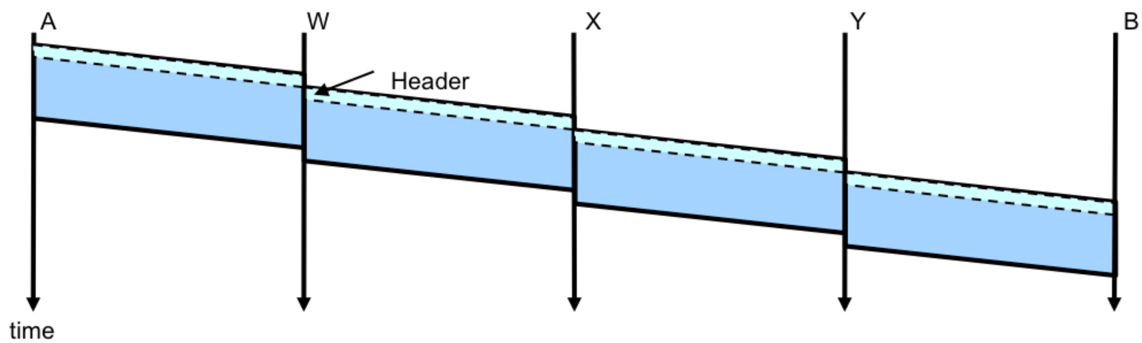

prop_delay = (250m*4)/(2*10^8m/s) = (1000m)/(2*10^8m/s) = 5us
A's transmission delay = (125B*8b/B)/(10^9b/s) = 1us
header_reception_delay = 3*(25B*8b/B)/(10^9b/s) = 0.6us
total = 6.6us

b) (5 points) How long will it take for a 125B packet to travel from node A to node B, assuming routers implement store-and-forward?

Solution:

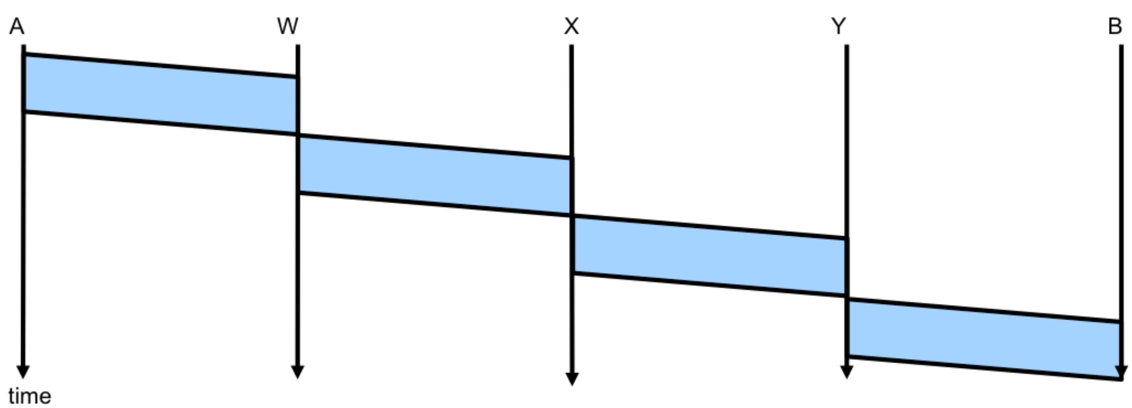

prop_delay = (250m*4)/(2*10^8m/s) = (1000m)/(2*10^8m/s) = 5us
transmission_delay = 4*(125B*8b/B)/(10^9b/s) = 4*1us = 4us
total = 9us

Question 3. *Longest Prefix Routing and CIDR (15 points)*
An IP router (which uses Longest Prefix Matching to make forwarding decisions) has the following table, which maps classful prefixes to outgoing interfaces:

Prefix	Interface
200.10.192.0/24	3
200.10.128.0/24	1
200.10.160.0/24	2
200.10.224.0/24	3

a) (5 points) After an upgrade, the router supports CIDR, so you can combine multiple entries into one and shorten some entries (fewer bits for prefix), while keeping the table equivalent (packets are forwarded to the same interface as before). Write down the equivalent forwarding table that has least number of entries and shortest entries.

Prefix	Interface
200.10.128.0/19	1
200.10.160.0/19	2
200.10.192.0/18	3

b) (5 points) A packet with destination IP 200.10.187.11 will be forwarded to interface ___2___

c) (5 points) A new entry 200.10.176.0/20 to interface 1 should be inserted into the table, write down the new CIDR-based forwarding table.

Prefix	Interface
200.10.128.0/19	1
200.10.176.0/20	1
200.10.160.0/20	2
200.10.192.0/18	3

Question 4. *Queueing (20 points)*
Consider two nodes A and B connected by a store-and-forward router R, which has an infinite queue. Assume the link from A to R has an infinite capacity, while the link from R to B is 1Mbps. Unless otherwise specified, there is no other traffic in the network besides the traffic sent from A to B.

Assume node A sends to B 100 packets, every 10ms each. The size of each packet is 125bytes

a) (3 points) What is the maximum number of packets in R's queue?
b) (3 points) What is the average queueing delay experienced by A's packets at R?

Assume now, that node A sends 250byte packets, instead of 125byte packets.

c) (4 points) What is the maximum number of packets in R's queue?
d) (4 points) What is the average queueing delay experienced by A's packets at R?

Finally, assume that there is an intermittent failure on the link from A to R, so the 31st, 32nd, 33rd, …, and 70th packets are all lost, that is, they never reach R. The size of all packets is 250bytes, and none of the packets is retransmitted.

e) (3 points) What is the maximum number of packets in R's queue?
f) (3 points) What is the average queueing delay experienced by A's packets at R?

Solution: In all cases it takes les than 10m to transmit the packet on the bottleneck link, R-B. In particular, it takes 125bytes*8/1Mbps = 1ms to transmit an 125byte packet, and 250bytes*8/1Mbps = 2ms to transmit a 250byte packet.

Also, since R is a store-and-forward router, the packet is stored (queued) before being send out. Thus the answers for about points are:

a) 1 packet;
b) 0s;
c) 1 packet
d) 0s
e) 1packet
f) 0s

Note: There was only 1 point deduction if you answered 0 at questions (a), (c) and (e).

Question 5. *DNS ans Link Layer (10 points)*
Consider the following network where Local DNS Server A receives a request to resolve host C's address. A already has the IP address of the appropriate TLD server (B) cached. The Authoritative DNS server for C is D. Circle R represents a router connecting two network segments. The MAC address and the IP address of host X are denoted as MAC_X and IP_X, respectively.

Node R has two network adaptors. (R1 is attached to the left network while R2 is connected with the right network.) Assume the Address Resolution Protocol (ARP) table at each host and router is pre-populated with the MAC, IP addresses of hosts in the same local area network.

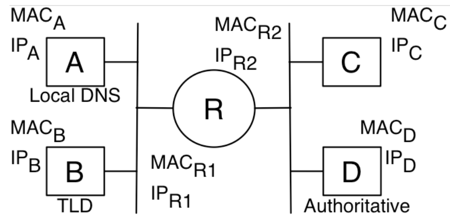

a) (5 points) If B does not have a record for C, and it does not support DNS recursive queries, write down the sequence of Ethernet frames during the DNS lookup process by filling the following table. (The number of rows is not necessarily equal to the number of frames during the process.)

Frame Number	Destination MAC address	Source MAC address	Destination IP address	Source IP address	Transport protocol
1	MAC_B	MAC_A	IP_B	IP_A	UDP
2	MAC_A	MAC_B	IP_A	IP_B	UDP
3	MAC_{R1}	MAC_A	IP_D	IP_A	UDP
4	MAC_D	MAC_{R2}	IP_D	IP_A	UDP
5	MAC_{R2}	MAC_D	IP_A	IP_D	UDP
6	MAC_A	MAC_{R1}	IP_A	IP_D	UDP
7					

b) (5 points) Repeat question (a), assuming B now supports DNS recursive queries.

Frame Number	Destination MAC address	Source MAC address	Destination IP address	Source IP address	Transport protocol
1	MAC_B	MAC_A	IP_B	IP_A	UDP
2	MAC_{R1}	MAC_B	IP_D	IP_B	UDP
3	MAC_D	MAC_{R2}	IP_D	IP_B	UDP
4	MAC_{R2}	MAC_D	IP_B	IP_D	UDP
5	MAC_B	MAC_{R1}	IP_B	IP_D	UDP
6	MAC_A	MAC_B	IP_A	IP_B	UDP
7					

Question 6. *Flow control (20 points)*
Assume host A sends traffic to B using a sliding window protocol, with a window size of 4 packets. Assume the RTT between A and B is 4ms, and that host A can send 2 packets every 1ms, while host B can deliver at most 1 packet every 1ms to the local application.

Host B assume a packet is lost if it sees a gap in the sequence numbers of the packets it receives. For example, if host B receives packet 1 and then packet 3, it assumes that packet 2 has been lost.

B sends and *ack* message when it delivers a packet to the local application, and a *nack* message when a packet has been lost. B sends only one *nack* message for each lost packet. Upon receiving a *nack* packet, A resends that packet.

Assume A sends 6 packets to B, and that the 2^{nd} and the 3^{rd} packets are lost (these are the only packets that are lots). Using the time diagram on the next page answer the following questions:

a) (10 points) Depict each transmission of a packet, a *nack*, and an *ack* messages.
b) (6 points) Depict the packets queued at both the sender (i.e., host A), and the receiver (i.e., host B).
c) (4 points) When is the 6th packet delivered to the local application by host B?

Note: The followings three clarifications (corrections) were made during the exam:

- **The RTT should be 2ms (as shown in the time diagram) not 4m.**
- **Host B can send to *nacks* in the same message.**
- **Host performs in-order delivery.**

Solution: (See next page)

Note: The following solutions also received full score:

- Solutions assuming that Host B can send only one nack at a time;, i.e., only one nack per 0.5ms time slot.
- Solutions assuming the sender can send two packets at once (during a 0.5ms time slot) as long as Host B doesn't send any other packet for the next 0.5ms slot.

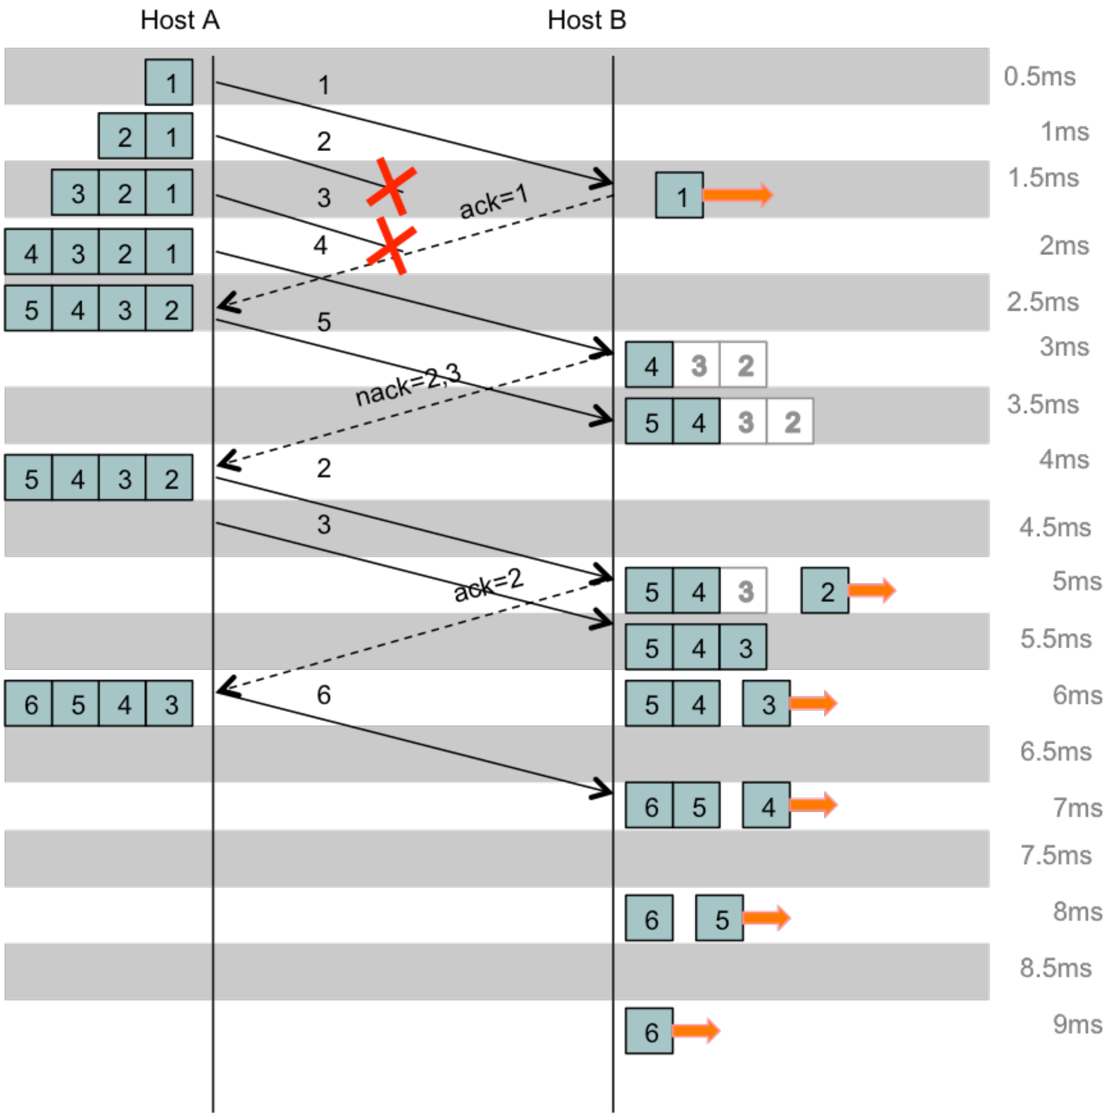
Host A
Host B
0.5ms
1ms
1.5ms
2ms
2.5ms
3ms
3.5ms
4ms
4.5ms
5ms
5.5ms
6ms
6.5ms
7ms
7.5ms
8ms
8.5ms
9ms
ack=1
nack=2,3
ack=2

Question 7. *Media Access Protocol (15 points)*
Consider N hosts connected to a local are network (LAN), and consider two possible media access protocols:

- A random access protocol, such as Sloted Aloha.
- A Token Passing protocol where a token is passed sequentially to each host. A host can send a packet only when it has the token. A host can send only one packet at a time. If it doesn't have a packet to send, the host simply passes the token to the next host.

In each of the following two cases, which strategy would you pick and why?

a) (5 points) Assume there is only one host sending data. Which one of the two media access protocols will achieve a higher throughout and why? (no more than three sentences)

The random access protocol as the host can send back-to-back packets. In contrast, with the Token Passing protocol, the host needs to wait for the token to visit every other host (which have no packets to send) before sending a new packet.

b) (5 points) Now assume that every host in the LAN sends data. Which of the two media access protocol will achieve higher throughput in this case, and why? (no more than three sentences)

The Token Passing protocol since every host will send a packet when it gets the token. In contrast, with the random access protocol, the more senders are the more collisions and packets loss are.

a) (5 points) Assume exactly two hosts send data. In the case of the random access protocol the probability of collision is $p = 0.5$ and all transmissions are synchronized. In the case of the Token Passing protocol assume that it takes one packet transmission time for the token to traverse all of the idle hosts. For both protocols assume that all packets are the same size. Which of the two media access protocols will achieve a higher network utilization and by how much?

Let T be the packet transmission time. With the Token Passing protocol every host sends one packet every 3*T sec (it needs to wait 2*T sec to get the token and it needs another T sec to send the packet). Since there are two hosts, the network utilization is 2/3.

In the case of the random access protocol, every T sec a packet is successfully sent with probability $1 - p = 0.5$. Thus, the network utilization is 1/2.

Thus, the Token Passing protocol achieves a 33% higher network utilization than the random access protocol.

Note: Solutions that stated that the token-passing improves the utilization by 1/2 - 1/6 = 1/6 also received full points.