Name: ____

Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.02 Fall 2011

Quiz III

December 19, 2011

"×" your section	Section	Time	Room	Recitation Instructor
☐	1	10-11	36-112	Paul Ampadu
☐	2	11-12	36-112	Sidhant Misra
☐	3	12-1	36-112	Sidhant Misra
☐	4	1-2	36-112	Paul Ampadu
☐	5	2-3	26-168	Karl Berggren
☐	6	3-4	26-168	Karl Berggren

"×" your TA
- ☐ Mukul Agarwal
- ☐ Jason Cloud
- ☐ Shuo Deng
- ☐ Lyla Fischer
- ☐ Rui Li
- ☐ Ruben Madrigal
- ☐ Surapap R.
- ☐ Xiawa Wang

There are a lucky **13 questions** (with multiple parts) and **12 pages** in this quiz booklet. Answer each question according to the instructions given. You have **2 hours** to answer the questions.

If you find a question ambiguous, please be sure to write down any assumptions you make. *Please be neat and legible.* If we can't understand your answer, we can't give you credit! **And please show your work for partial credit.**

Use the empty sides of this booklet if you need scratch space. You may also use them for answers, although you shouldn't need to. *If you use the blank sides for answers, make sure to say so!*

Please write your name CLEARLY in the space at the top of this page. NOW, please!

Two two-sided "crib sheets" allowed. No other notes, books, calculators, computers, cell phones, PDAs, information appliances, carrier pigeons carrying messages, etc.!

Do not write in the boxes below

1-2 (x/20)	3-5 (x/25)	6-7 (x/18)	8-10 (x/9)	11 (x/9)	12-13 (x/19)	Total (x/100)

I It may be Little, but it's the Law!

1. [3+7=10 points]: Carrie Coder has set up an email server for a large email provider. The email server has two modules that process messages: the *spam filter* and the *virus scanner*. As soon as a message arrives, the spam filter processes the message. After this processing, if the message is spam, the filter throws out the message. The system sends all non-spam messages immediately to the virus scanner. If the scanner determines that the email has a virus, it throws out the message. The system then stores all non-spam, non-virus messages in the inboxes of users.

Carrie runs her system for a few days and makes the following observations:

1. On average, $\lambda = 10000$ messages arrive per second.
2. On average, the spam filter has a queue size of $N_s = 5000$ messages.
3. $s = 90\%$ of all email is found to be spam; spam is discarded.
4. On average, the virus scanner has a queue size of $N_v = 300$ messages.
5. $v = 20\%$ of all non-spam email is found to have a virus; these messages are discarded.

A. On average, in 10 seconds, how many messages are placed in the inboxes?

(Explain your answer in the space below.)

B. What is the **average delay** between the arrival of an email message to the email server and when it is ready to be placed in the inboxes? All transfer and processing delays are negligible compared to the queueing delays. Make sure to draw a picture of the system in explaining your answer. **Derive your answer in terms of the symbols given, plugging in all the numbers only in the final step.**

(Explain your answer in the space below.)

II Hunting in (Packet) Pairs

2. [2+1+2+5=10 points]: A sender S and receiver R are connected using a link with an unknown bit rate of C bits per second and an unknown propagation delay of D seconds. At time $t = 0$, S schedules the transmission of a **pair of packets** over the link. The bits of the two packets reach R in succession, spaced by a time determined by C. Each packet has the same known size, L bits.

The last bit of the first packet reaches R at a known time $t = T_1$ seconds. The last bit of the second packet reaches R at a known time $t = T_2$ seconds. As you will find, this packet pair method allows us to estimate the unknown parameters, C and D, of the path.

A. Write an expression for T_1 in terms of L, C, and D.

B. At what time does the **first** bit of the **second** packet reach R? Express your answer in terms of T_1 and one or more of the other parameters given (C, D, L).

(Explain your answer in the space below.)

C. What is T_2, the time at which the **last** bit of the **second** packet reaches R? Express your answer in terms of T_1 and one or more of the other parameters given (C, D, L).

(Explain your answer in the space below.)

D. Using the previous parts, or by other means, derive expressions for the bit rate C and propagation delay D, in terms of the known parameters (T_1, T_2, L).

(Explain your answer in the space below.)

III Le Big MAC

3. [8 points]: We studied TDMA, Aloha, and CSMA protocols in 6.02. In each statement below, assume that the protocols are implemented correctly.

(Circle True or False for each choice.)

A. True / False TDMA may have collisions when the size of a packet exceeds one time slot.

B. True / False In Aloha with correct stabilization, two nodes have a certain probability of colliding in a time slot. If they actually collide in that slot, then they will experience a lower probability of colliding with each other when they each retry.

C. True / False In slotted Aloha with stabilization, each node's transmission probability converges to $1/N$, where N is the number of backlogged nodes.

D. True / False In a network in which all nodes can hear each other, a correctly implemented CSMA protocol will have no collisions when the packet size is larger than one time slot.

4. [4+4=8 points]: Carl Coder implements a simple slotted Aloha-style MAC for his room's wireless network. His room has only two backlogged nodes, A and B. Carl picks a transmission probability of $2p$ for node A and p for node B. Each packet is one time slot long and all transmissions occur at the beginning of a time slot.

A. What value of p maximizes the utilization of this network, **and** what is the maximum utilization?

(Explain your answer in the space below.)

B. Instead of maximizing the utilization, suppose Carl chooses p so that the throughput achieved by A is **three times** the throughput achieved by B. What is the utilization of his network now?

(Explain your answer in the space below.)

5. [4+5=9 points]: Carl Coder replaces the "send with fixed probability" MAC with one that uses a *contention window* at each node. He configures node A to use a fixed contention window of W and node B to use a fixed contention window of $2W$. Before a transmission, each node independently picks a random integer t uniformly between 1 and its contention window value, and transmits a packet t time slots from now. Each packet is one time slot long and all transmissions occur at the beginning of a time slot.

A. What is the probability that A and B will collide the next time they each transmit?

(Explain your answer in the space below.)

B. Suppose there is no collision at the next packet transmission. Calculate the probability that A will transmit before B? (It may be useful to apply the formula $\sum_{i=1}^{n} i = n(n+1)/2$.)

(Explain your answer in the space below.)

IV Roto-Routing

Eager B. Eaver implements the distance-vector protocol studied in 6.02, **but on some of the nodes**, his code sets the cost and route to each advertised destination D **differently**:

Cost to $D = \min(\text{advertised\_cost})$ heard from each neighbor.
Route to D = link to a neighbor that advertises the minimum cost to D.

Every node in the network periodically advertises its vector of costs to the destinations it knows about to all its neighbors. All link costs are positive.

At each node, a route for destination D is **valid** if packets using that route will eventually reach D.

At each node, a route for destination D is **correct** if packets using that route will eventually reach D along some minimum-cost path.

Assume that there are no failures and that the routing protocol has converged to produce *some* route to each destination at all the nodes.

6. [3+3+2=8 points]: Circle **True** or **False** for each statement below, providing a **brief explanation** for your choice below each statement. Assume a network in which **at least two of the nodes** (and possibly all of the nodes) run Eager's modified version of the code, while the remaining nodes run the method discussed in 6.02.

A. True / False There exist networks in which some nodes will have invalid routes.

B. True / False There exist networks in which some nodes will **not** have correct routes.

C. True / False There exist networks in which **all** nodes will have correct routes.

We can't have a 6.02 quiz without our friend Ben Bitdiddle! Help Ben answer these questions about the distance-vector protocol he runs on the network shown in the figure below. The link costs are shown near each link. Ben is interested in minimum-cost routes to destination node D.

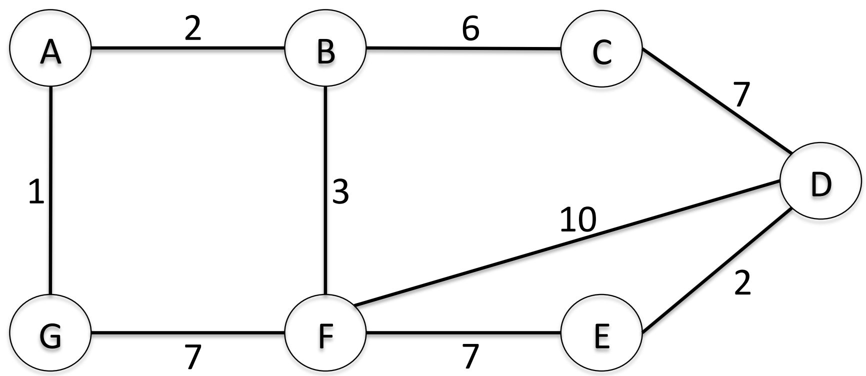

Each node sends a distance-vector advertisement to all its neighbors at times $0, T, 2T, \ldots$. Each node integrates advertisements at times $T/2, 3T/2, 5T/2, \ldots$. You may assume that all clocks are synchronized. The time to transmit an advertisement over a link is negligible. There are no failures or packet losses.

7. [5+5=10 points]: Use the definitions of a **valid** route and a **correct** route from the previous page to answer the following questions. **You must explain your answers.**

A. At what time will **all** nodes have integrated a valid route to D into their routing tables? What node is the **last one** to integrate a valid route to D? Answer both questions.

B. At what time will **all** nodes have integrated a correct (minimum-cost) route to D into their routing tables? What node is the **last one** to integrate a correct route to D? Answer both questions.

V Hyper-Routing

The hypercube is an interesting network topology. An n-dimensional hypercube has 2^n nodes, each with a unique n-bit address. Two nodes in the hypercube are connected with a link if, and only if, their addresses have a Hamming distance of 1. The picture below shows hypercubes for $n = 3$ and 4. The solid and dashed lines are the links. We are interested in link-state routing over hypercube topologies.

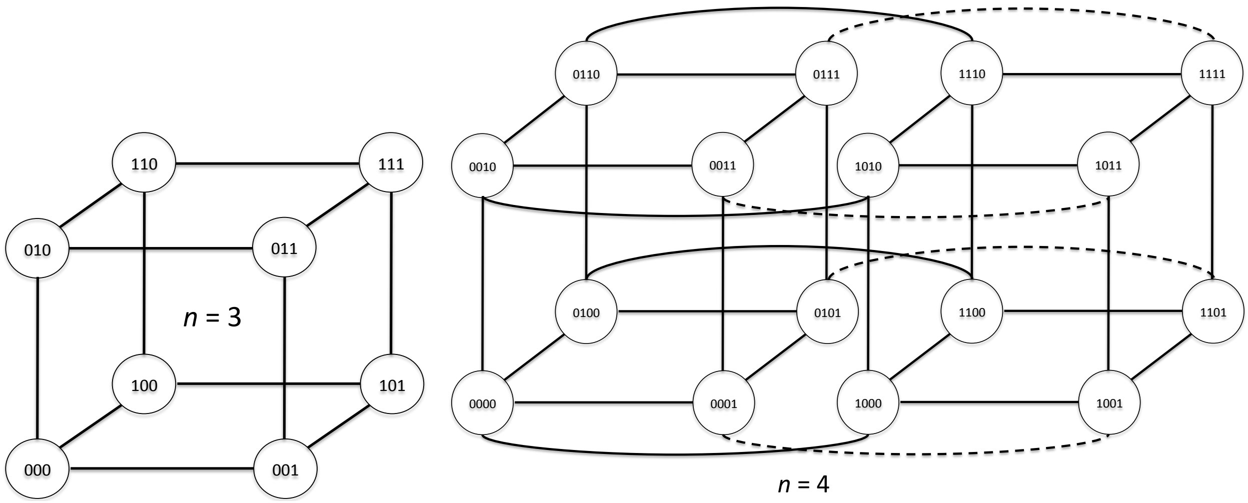

8. [2 points]: Suppose $n = 4$. Each node sends a link-state advertisement (LSA) periodically, starting with sequence number 0. All link costs are equal to 5. Node 1000 discovers that its link to 1001 has failed. There are no other failures. What are the contents of the **fourth LSA** originating from node 1000?

9. [2+2=4 points]: Suppose $n = 4$. Three of the links at node 1000, including the link to node 1001, fail. No other failures or packet losses occur.

A. How many distinct copies of any given LSA originating from node 1000 does node 1001 receive?

B. How many distinct copies of any given LSA originating from node 1001 does node 1000 receive?

10. [3 points]: Suppose $n = 3$ and there are no failures. Each link has a **distinct**, **positive**, **integral** cost. Node 000 runs Dijkstra's algorithm (breaking ties arbitrarily) and finds that the minimum-cost path to 010 has 5 links on it. What can you say about the cost of the direct link between 000 and 010? **Explain your answer.**

VI Stop, Wait, and Tell Me...

Alyssa P. Hacker sets up a wireless network in her home to enable her computer ("client") to communicate with an Access Point (AP). The client and AP communicate with each other using a **stop-and-wait** protocol.

The data packet size is 10000 bits. The total round-trip time (RTT) between the AP and client is equal to 0.2 milliseconds (that includes the time to process the packet, transmit an ACK, and process the ACK at the sender) **plus** the transmission time of the 10000 bit packet over the link.

Alyssa can configure two possible transmission bit rates for her link, with the following properties:

Bit rate	Bi-directional packet loss probability	RTT
10 Megabits/s	$1/11$	_____
20 Megabits/s	$1/4$	_____

11. [3+5+1=9 points]: Alyssa's goal is to select the bit rate that provides the higher throughput for a stream of packets that need to be delivered reliably between the AP and client using stop-and-wait. For both bit rates, the **retransmission timeout (RTO) is 2.4 milliseconds**.

A. What is the round-trip time (RTT) for each bit rate? Calculate it below (show your work) and enter the RTT values **in milliseconds** in the table above.

B. For each bit rate, calculate the **expected time**, in milliseconds, to successfully deliver a packet and get an ACK for it. **Show your work.**

C. Using the above calculations, which bit rate would you choose to achieve Alyssa's goal?

VII Slip Slidin' Away

Annette Werker correctly implements the fixed-size sliding window protocol developed in 6.02. She instruments `ReliableSenderNode` to store the time at which each DATA packet is sent and the time at which each ACK is received. A snippet of the DATA and ACK traces from an experiment is shown in the picture below. Each $+$ is a DATA packet transmission, with the x-axis showing the transmission time and the y-axis showing the sequence number. Each $\times$ is an ACK reception, with the x-axis showing the ACK reception time and the y-axis showing the ACK sequence number. All DATA packets have the same size.

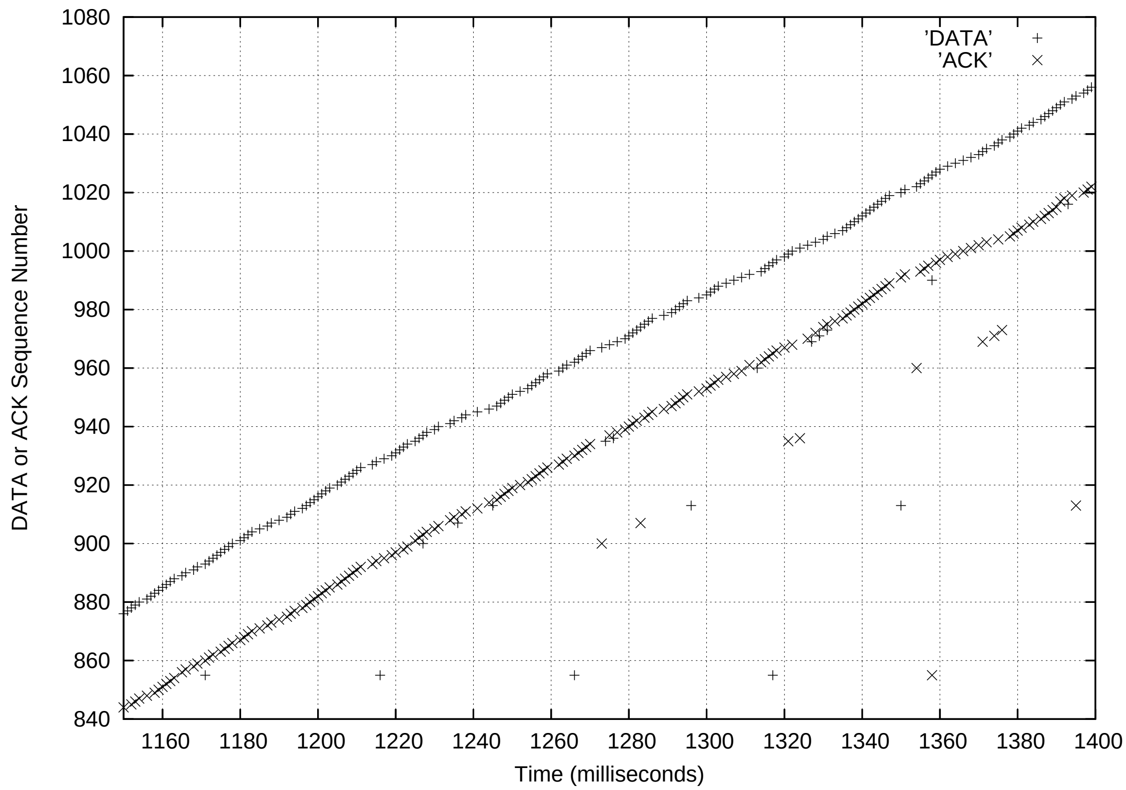

12. [11 points]: Answer the following questions, **providing a brief explanation** for each one.

A. Estimate any one sample round-trip time (RTT) of the connection. **Also show it on the picture.**

B. On the picture, circle DATA packet **re**transmissions for four different sequence numbers.

C. Some DATA packets in this trace may have incurred **more than one re**transmission? **On the picture**, draw a square around one such retransmission.

Picture from the previous page copied below for your convenience.

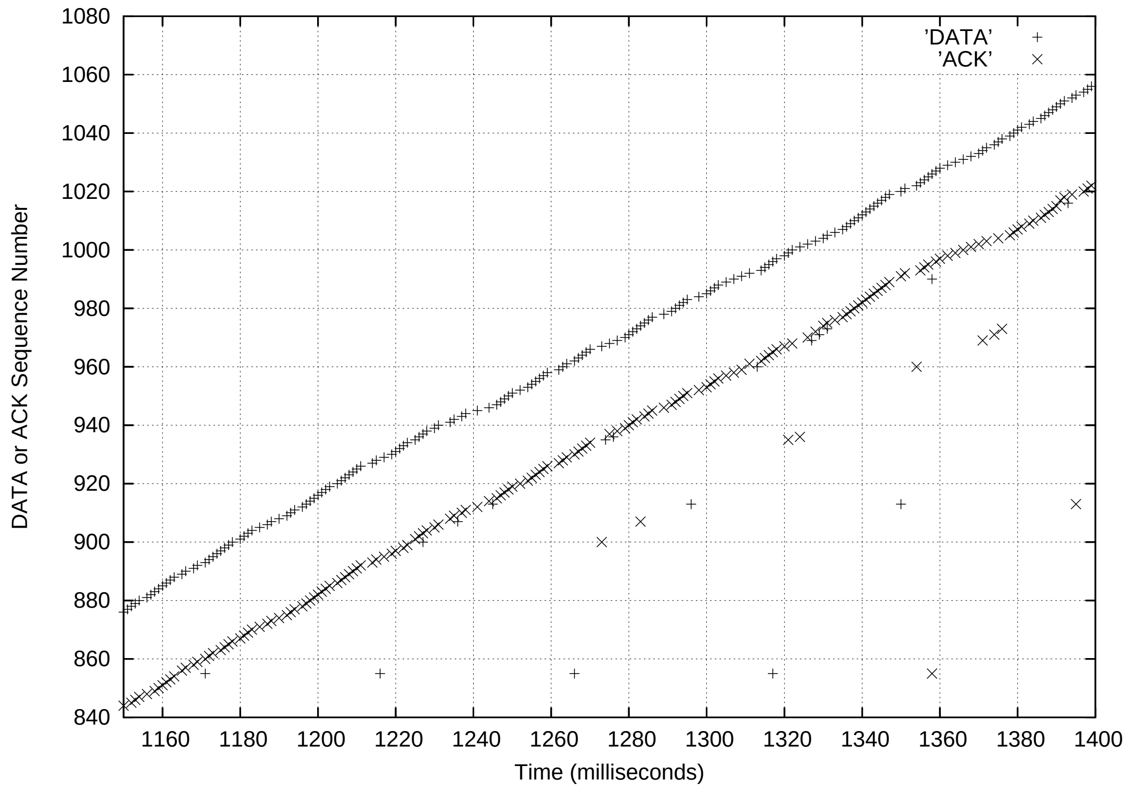

D. Which of these is the **best estimate** of the sender's window size? **Also show it on the picture.**

(a) 20.

(b) 28.

(c) 36.

(d) 44.

E. Which of these is the **best estimate** of the throughput of the connection in packets per second? Give a one-line explanation of how you estimated it.

(a) ≈ 0.7.

(b) ≈ 500.

(c) ≈ 750.

(d) ≈ 1500.

F. Considering only sequence numbers > 880, which of these is the **best estimate** of the packet loss rate experienced by DATA packets? Give a one-line explanation of how you estimated it.

(a) $\approx 1.75\%$.

(b) $\approx 3.5\%$.

(c) $\approx 7\%$.

(d) $\approx 14\%$.

Consider the same setup as in the previous two pages. Suppose the window size for the connection is equal to twice the bandwidth-delay product of the network path.

13. [2+3+3=8 points]: For each change to the parameters of the network path or the sender given below, explain if the connection's **throughput** (not utilization) will **increase**, **decrease**, or **remain the same**. In each statement, nothing other than what is specified in that statement changes. **Explain your answers.**

A. The packet loss rate, ℓ, decreases to $\ell/3$.

B. The minimum value of the RTT, R, increases to $1.8R$.

C. The window size, W, decreases to $W/3$.

FIN

Have a great holiday season!