

Name: _____

DEPARTMENT OF EECS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.02 Spring 2012

Quiz III

May 22, 2012

<u>"×" your section</u>	<u>Section</u>	<u>Time</u>	<u>Recitation Instructor</u>	<u>TA</u>
<input type="checkbox"/>	1	10-11	Vincent Chan	Jared Monnin
<input type="checkbox"/>	2	11-12	Vincent Chan	Anirudh Sivaraman
<input type="checkbox"/>	3	12-1	Sidhant Misra	Sungwon Chung
<input type="checkbox"/>	4	1-2	Sidhant Misra	Omid Aryan / Nathan Lachenmyer
<input type="checkbox"/>	5	2-3	Katrina LaCurts	Sunghyun Park
<input type="checkbox"/>	6	3-4	Katrina LaCurts	Muyiwa Ogunnika

Please read and follow these instructions:

0. Please write your name in the space above and × your section.
1. There are **14 questions** (some with multiple parts) and **11 pages** in this quiz booklet.
2. Answer each question according to the instructions given, within **120 minutes**.
3. **Please answer legibly. Explain your answers, especially when we ask you to.** If you find a question ambiguous, write down your assumptions. Show your work for partial credit.
4. Use the empty sides of this booklet if you need scratch space. *If you use the blank sides for answers, make sure to say so!*

Two two-sided "crib sheets" and a calculator allowed. No other aids.

Do not write in the boxes below

1-3(x/20)	4-6 (x/35)	7-11 (x/29)	12-14 (x/16)	Total(x/100)

I MAC

1. **[4 points]:** We studied two ways of stabilizing Aloha: one in which a node's probability of transmission changes in response to collisions and successes, and another in which each node uses a varying contention window. Most practical systems use contention windows. Give one reason why they do that. **Explain your answer.**

2. **[6 points]:** Ben Bitdiddle runs the slotted Aloha protocol with stabilization. Each packet is one time slot long. At the beginning of time slot T , node i has a probability of transmission equal to p_i , $1 \leq i \leq N$, where N is the number of backlogged nodes. The increase/decrease rules for p_i are doubling/halving, with $p_{\min} \leq p_i \leq p_{\max}$, as described in 6.02 this term.

Ben notes that exactly two nodes, j and k , transmit in time slot T . After thinking about what happens to these two packets (don't ask us!), derive an expression for the probability that **exactly one node** (out of the N backlogged nodes) will transmit successfully in time slot $T + 1$. **Explain your answer.**


3. [10 points]: Tim D. Vider thinks Time Division Multiple Access (TDMA) is the best thing since sliced bread (“if equal slices are good for bread, then equal slices of time must be good for the MAC too”, he says). Each packet is one time slot long.

However, in Tim’s network with N nodes, the **offered load is not uniform** across the different nodes. The **rate** at which node i generates new packets to transmit is $r_i = \frac{1}{2^i}$ packets per time slot ($1 \leq i \leq N$). That is, in each time slot, the **application** on node i produces a packet to send over the network with probability r_i .

A. (4 points) Tim runs an experiment with TDMA for a large number of time slots. At the end of the experiment, how many nodes (as a function of N) will have a substantial backlog of packets (i.e., queues that are growing with time)? **Explain your answer.**

B. (6 points) Let $N = 20$. Calculate the **utilization** of this non-uniform workload running over TDMA. A possibly useful formula: if $|\alpha| < 1$, $\sum_{i=k}^N \alpha^i = \alpha^k \cdot \frac{(1-\alpha^{N-k+1})}{(1-\alpha)} \approx \frac{\alpha^k}{1-\alpha}$ when $N - k$ is large. **Explain your answer.**

II Go Ahead, Make My Route




Jack Ripper uses a minimum-cost **distance-vector routing protocol** in the network above. Each link cost (not shown) is a **positive integer** and is the same in each direction of the link. Jack sets “infinity” to 32 in the protocol. After all routes have converged (breaking ties arbitrarily), F ’s routing table is as follows:

Destination	Cost	Route
A	6	link $\langle FC \rangle$
B	4	link $\langle FC \rangle$
C	3	_____ (fill in)
D	5	link $\langle FD \rangle$
E	1	_____ (fill in)

4. [17 points]: Using the information provided, answer these questions:

- A. (2 points) Fill in the two missing blanks in the table above.
- B. (9 points) For **each link** in the picture, write the link’s cost in the box near the link. Each cost is either a positive integer or an expression of the form “ $< c, \leq c, \geq c, \text{ or } > c$ ”, for some integer c .
- C. (2 points) Suppose link $\langle FE \rangle$ fails, but there are no other changes. When the protocol converges, what will F ’s **route** (not path) to E be? (If there is no route, say “no route”.) **Explain.**
- D. (4 points) Now suppose links $\langle BC \rangle$ and $\langle BF \rangle$ **also** fail soon after link $\langle FE \rangle$ fails. There are no packet losses. In the **worst case**, C and F enter a “count-to-infinity” phase. How many distinct route advertisements (with different costs) must C hear from F , before C determines that it does not have any valid route to node A ? **Explain.**

Alyssa P. Hacker runs the **link-state routing protocol** in the network shown below. Each node runs Dijkstra's algorithm to compute minimum-cost routes to all other destinations, breaking ties arbitrarily.




5. [8 points]: Answer the following questions, **explaining each answer**.

- A. (3 points) In what order does C add destinations to its routing table in its execution of Dijkstra's algorithm? Give all possible answers.

- B. (3 points) Suppose the cost of link $\langle CB \rangle$ increases. What is the largest value it can increase to, before **forcing** a change to any of the routes in the network? (On a tie, the old route remains.)

- C. (2 points) Assume no link-state advertisement (LSA) packets are lost on any link. When C generates a new LSA, how many copies of that LSA end up getting flooded in total over all the links of this network, using the link-state flooding protocol described in 6.02?

Picture reproduced here for convenience.



The links in Alyssa's network are unreliable; on each link, any packet sent over the link is delivered with some probability, p , to the other end of the link, independent of all other events ($0 < p < 1$). Suppose links $\langle CE \rangle$ and $\langle BD \rangle$ fail.

6. [10 points]: Answer the following questions, **explaining each answer**.

A. (2 points) How do C and E discover that the link has failed? How does the method work?

B. (5 points) Over this unreliable network, link state advertisements (LSAs) are lost according to the probabilities mentioned above. Owing to a bug in its software, E **does not originate any LSA of its own or flood them**, but all other nodes (except E) work correctly. Calculate the probability that A learns that link $\langle CE \rangle$ has failed from the **first** LSA that originates from C after C discovers that link $\langle CE \rangle$ has failed. **Note that link $\langle BD \rangle$ has also failed.**

C. (3 points) Suppose only link $\langle CE \rangle$ had failed, **but not $\langle BD \rangle$** , which like the other surviving links can delivery packets successfully with probability p . Now, would the answer to part B above **increase, decrease, or remain the same**? Why? (No math necessary.)

III Seek-Transmit, by Gloria Mundi

Gloria Mundi, a recent alum, has just started an open source project she calls *Seek-Transmit* to help users find movie clips (legally, of course!) over a best-effort network. Her system has two parts:

1. **Seek:** The client running on the user's device sends a short search string, called a *seek*, to a seek server, which responds with a list of zero or more *transmit servers* that have the desired content.
Seek uses the 6.02 stop-and-wait protocol. The ACKs are replaced by the server's response to a seek; if the client does not get a response within a timeout period, T seconds, it re-tries the seek. It keeps retrying every T seconds until it gets a response. There are no spurious retries.
2. **Transmit:** The client contacts one of the transmit servers to stream the content. Transmit uses the 6.02 sliding window protocol for reliable delivery. There are no spurious retransmissions.

Gloria's goal is for each client to obtain a response to the seek as quickly as possible, and then for the transmit phase to achieve as high a throughput as possible.

Each packet includes in its headers a **source** field and a **hop limit** field.

7. [4 points]: Answer the following questions about these packet header fields:

A. (3 points) What is the purpose of the **hop limit** field? How does it achieve this purpose?

B. (1 point) State one use of the **source** field in Gloria's Seek-Transmit system.

Gloria deploys multiple seek servers around the world. The video clips are also on a number of transmit servers around the world. She writes some monitoring software to help select the best seek and transmit servers, from among the set of possible servers. The monitoring software measures the following and provides the results to the client:

- R_i , the average round-trip time (RTT), in seconds, along the **network path** to the i^{th} server, **excluding** the server's processing delay, but **including** the transmission delay of a packet and its response.
- D_j , the average delay (in seconds) to process a seek request at the j^{th} seek server.
- p_i , the probability that the client successfully gets a response from the i^{th} server to a packet it sends. (The packet/response is a seek/seek response for a seek server, and data/ACK for a transmit server.)
- B_i , the rate of the bottleneck link for transfers between the i^{th} server and client, in bytes per second.

8. [5 points]: Recall that the goal is for the client to select **one** seek server that will respond quickest (among the set of choices) to the client. Recall also that the seek protocol is stop-and-wait, with a fixed timeout value of T seconds. For each seek server i , the client computes some function, f_i (in terms of the parameters defined above), and then minimizes or maximizes it, to accomplish the desired goal. Write down f_i , state whether it should be maximized or minimized, and **explain your answer**.

9. [5 points]: Recall that Gloria uses the 6.02 sliding window protocol to deliver video from a transmit server to the client. Assume that the window size for the data transfer is picked to achieve the maximum possible throughput of the transfer, subject to the parameters of the network defined above. Assume also that each transmit server produces data extremely fast, so the server itself is not the bottleneck. The client's goal is to pick **one** transmit server that will give it the highest video transfer throughput. For each transmit server i , the client computes some function, g_i (in terms of the parameters defined above), and then minimizes or maximizes it, to accomplish the desired goal. Write down g_i , state whether it should be maximized or minimized, and **explain your answer**.

10. [5 points]: Annette Werker implements an exponentially weighted moving average (EWMA) method to estimate the smoothed RTT (srtt) and the RTT deviation (rttdev) of the network path between a server and client. These estimators have the form

$$\text{srtt} \leftarrow \alpha r + (1 - \alpha)\text{srtt} \quad \text{and} \quad \text{rttdev} \leftarrow \beta(|r - \text{srtt}|) + (1 - \beta)\text{rttdev}.$$

Fill in the blanks (if choices are provided, pick the best one):

- A. The variable r in the above expressions is _____.
 - B. These EWMA estimators help the _____ (sender / receiver / switch / MAC) to set its _____ (window size / timeout / transmit probability / queue size).
 - C. The larger the value of α and β , the _____ (greater / lower) the contribution of any past samples to the srtt and rttdev estimates.
 - D. The srtt and rttdev estimates are updated when _____ (what event?).
- 11. [10 points]:** Annette Werker conducts tests between a transmit server and a client using the sliding window protocol. There is no other traffic on the path and no packet loss. She finds that:

- With a window size $W_1 = 50$ packets, the throughput is 200 packets per second.
- With a window size $W_2 = 100$ packets, the throughput is 250 packets per second.

Annette finds that even this small amount of information allows her to calculate several things, assuming there is only one bottleneck link. Calculate the following, **explaining your answers**:

- A. (2 points) The minimum round-trip time between the client and server.
- B. (3 points) The average **queuing** delay at the bottleneck when the window size is 100 packets.
- C. (5 points) The average queue size when the window size is 100 packets.

Gloria realizes that packets that arrive too late at the receiver cause the video to stall, and that the application does not need *all* the packets to play properly. Missing packets are acceptable, but the packets must be delivered to the application in **increasing sequence order**. She **removes all retransmissions from the protocol**, so now the sender just sends data packets with incrementing sequence numbers at some constant rate. Packets may show up in arbitrary order, be delayed arbitrarily, and be duplicated in the network.

Gloria would like to modify the receiver transport layer to apply the following rules:

- R1. Deliver packets to the application only in strictly increasing sequence order (perhaps skipping some).
- R2. Every t milliseconds, deliver exactly one packet to the application. This packet should be the one with the smallest sequence number in the receiver's buffer that is strictly greater than the one previously delivered to the application. You may assume that such a packet always exists in the receiver's buffer.

12. [3 points]: The receiver transport protocol receives the following packets at the times shown:

Time (ms): 0, 8, 9, 13, 14, 19, 44, 48, 56

Packet: 5, 3, 1, 2, 10, 8, 5, 7, 4

Let $t = 10$ ms. The receiver transport delivers the first packet to the video application at time 10 ms. Write the time and packet sequence at which various packets are delivered to the application.

Time (ms): 10,

70

Packet:

13. [12 points]: Write Python code (don't worry about perfect syntax!) to implement Gloria's receiver transport protocol according to the rules above. Write two functions: `receive` and `deliver2app` (next page). The specifications of these functions are given below. Name variables descriptively, or define them. Use the blank side if you need more space, but tell us if you do so.

```
'''
receive() is called whenever a data packet arrives. Maintain
self.rcvbuf, a buffer of packets that haven't yet been delivered
to the app. Delete packets that will never be delivered to the app
from rcvbuf. You may use standard Python library modules and functions,
e.g., foo.sort() if you wish to sort list foo in ascending order.
'''
def receive(self, pkt):
    # Notation: pkt.seq is the sequence number of pkt
```

```
'''
Assume that deliver2app is called every t seconds by the receiver
transport protocol's main loop. Call app_receive(pkt) to deliver
a pkt from self.rcvbuf to the app, following rules R1 and R2.
'''
def deliver2app(self):
```

14. [1 points]: This elaborate set up with Seek-Transmit and Gloria Mundi has nothing to do with networking but everything to do with classical education. What does *sic transit gloria mundi* mean?
(Circle ALL that apply.)

- A. Gloria fell sick on Monday while traveling.
- B. Glory is fleeting, but obscurity is forever.
- C. Thus passes the glory of the world.
- D. You *#@\$*!#s. The "T" in MIT is "technology", not "torture". This truly *is* the last straw; I'm reporting you to the new President (the complaints to the Dean didn't work the last time).

FIN

Have a great summer!