Carnegie Mellon
Computer Science Department.
15-441 Spring 2010
Final exam

Name: ____

Andrew ID: ____

INSTRUCTIONS:
There are 20 pages (numbered at the bottom). Make sure you have all of them.
Please write your name on this cover and at the top of each page in this booklet **except the last.**
If you find a question ambiguous, be sure to write down any assumptions you make.
It is better to partially answer a question than to not attempt it at all.
Be clear and concise. Limit your answers to the space provided.

Section	Points	Score
A	21	
B	15	
C	12	
D	9	
E	18	
F	16	
G	14	
H	18	
I	14	
J	25	
K	15	
L	3	
Total	180	

A Short Answer

1. (3 points) Traditional topology based routing algorithms such as distance vector routing and link state routing don't work very well for wireless ad-hoc networks. In class, we saw alternatives such as source routing (DSR) and geographical routing (GPSR). Which type (GPSR or DSR) would you use for a vehicular network of rapidly moving cars with short-range wireless links in downtown Pittsburgh? Explain your choice briefly.

2. (4 points) Four users are sharing a common link of 4Mbps. User A is connected to the common link via a 5Mbps link and is using an application that requires xMbps. Users B, C and D are downloading large files (i.e. using as much bandwidth as possible). B is connected via a xMbps link while C and D are each connected via their own 10Mbps link.

   What is the max-min fair allocation, if....

   (a) x = 2Mbps

   (b) x = 0.5Mbps

3. (3 points) A computer is connected to a network on a 6Mbps link, regulated by a token bucket. The RTT of the network is 100ms and each packet is 1000 bytes. The token bucket is filled at a rate of 1 Mbps. It is initially filled to capacity with 8 Mbits. How long can the computer transmit at the full 6Mbps?

4. (3 points) We discussed three different routing protocols: link state routing (LS), distance vector routing (DV), and path vector routing (PV). Please answer the following questions by circling all the protocols for which the claim applies:

   (a) [LS, DV, PV] - Requires a map of the complete topology

   (b) [LS, DV, PV] - Sends its routing table to its neighbors

   (c) [LS, DV, PV] - Requires flooding

   (d) [LS, DV, PV] - Suffers the count to infinity problem

   (e) [LS, DV, PV] - BGP is this type of routing protocol

5. (5 points) Suppose you have a client C, a local name server L, and authoritative name servers A_root, A_com, and A_foo.com, where A_x is the name server that knows about the name zone x (and A_root is a root name server). Here are some possible scenarios for a name resolution:

   A. C asks L, L asks A_root, A_root tells L to ask A_com, L asks A_com, A_com tells L to ask A_foo.com, L asks A_foo.com, A_foo.com answers to L, L answers to C.

   B. C asks A_root, A_root tells C to ask A_com, C asks A_com, A_com tells C to ask A_foo.com, C asks A_foo.com, A_foo.com answers to C, C tells L the address.

   C. C asks L, L asks A_root, A_root asks A_com, A_com asks A_foo.com, A_foo.com answers to C.

   D. C asks L, L asks A_foo.com, A_foo.com answers to L, L answers to C.

   E. C asks L, L asks A_root, A_root tells L to ask A_com, L asks A_com, A_com tells L to ask A_foo.com, L asks A_foo.com, A_foo.com answers to C.

   (a) Suppose C looks up the address for www.foo.com. Assuming that all name servers initially have nothing in their cache, which of the above scenarios describes the order of events correctly? (Give the letter of the correct scenario.)

   (b) Client C performs a lookup for `ftp.foo.com` immediately after the previous request. Assuming all records have long TTLs, which of the above scenarios describes the order of events correctly? (Give the letter of the correct scenario.)

6. (3 points) Below is a plot of TCP throughput vs. time for a wired network (the line on top) and a wireless network (the line below). The link speeds of the two networks are identical.

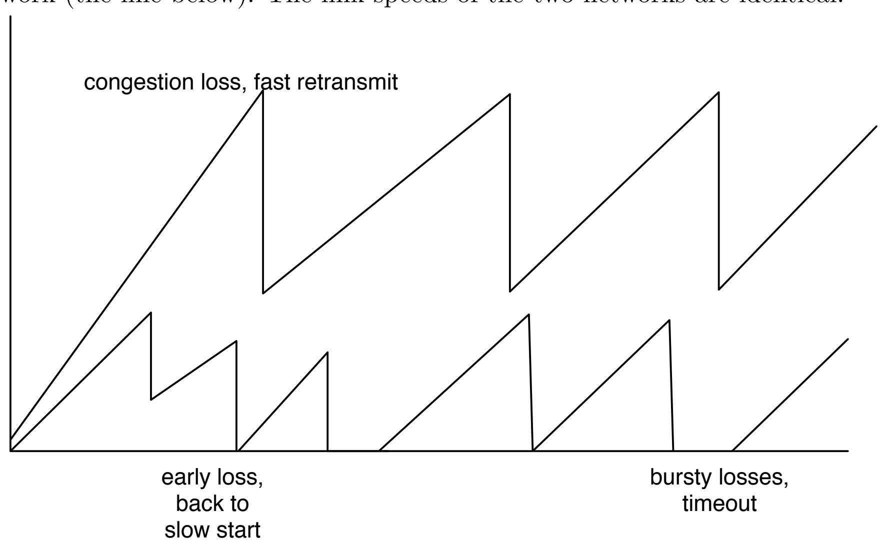

Why is the TCP performance on the wireless network so much worse? State the assumption made by TCP that breaks in the wireless case.

B MAC

7. (15 points) Consider the ad hoc network of the figure below. Two nodes are connected if they can hear each other. For this question consider two flows: flow 1 from A to B and flow 2 from C to D.

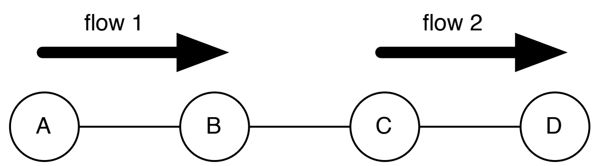

(a) You first consider using CSMA/CD from the Ethernet lecture. However, your friend says that this is impossible and that you must use CSMA/CA instead – explain why.

(b) Assume that all nodes implement CSMA/CA. Also, assume that both nodes A and C have a single packet to transmit. When node A has already begun transmitting its packet on flow 1, node C wants to transmit a packet on flow 2. What will node C do and which flows will receive their packets correctly?

(c) Assume that all nodes implement CSMA/CA with RTS/CTS. Also, assume that both nodes A and C have a single packet to transmit. When node A has already begun transmitting its packet on flow 1, node C wants to transmit a packet on flow 2. What will node C do and which flows will receive their packets correctly?

(d) Assume that all nodes implement CSMA/CA with RTS/CTS. Also, assume that both nodes A and C have a single packet to transmit. When node C has already begun transmitting its packet on flow 2, node A wants to transmit a packet on flow 1. What will node A do and which flows will receive their packets correctly?

(e) Will CSMA/CA with RTS/CTS provide a fair amount of bandwidth to each flow when both flows are backlogged? Explain briefly.

C Who's the fairest...

8. (12 points) Consider a router with three input flows and one output. It receives packets shown in the table. Assume that all packets are received at about the same time in the order listed. (Obviously two packets from the same flow don't arrive at exactly the same time, but assume they arrive very close together.) Also, assume the output port is busy during the period that the packets are received and all queues are empty initially. The output link runs at 8 Mbits/sec.

Packet ID	Flow	Size(bytes)
1	1	100
2	2	160
3	2	180
4	1	110
5	2	100
6	3	200
7	3	190
8	1	120

(a) Give the order in which the packets are transmitted when using **packet-by-packet round-robin**. Please record any calculations you do in the empty columns for partial credit.
ORDER:

SCRATCH PAD:

Packet ID	Flow	Size					
1	1	100					
2	2	160					
3	2	180					
4	1	110					
5	2	100					
6	3	200					
7	3	190					
8	1	120					

(b) Give the order in which the packets are transmitted when using **fair-queuing**. Please record any calculations you do in the empty columns for partial credit.

ORDER:

SCRATCH PAD:

Packet ID	Flow	Size					
1	1	100					
2	2	160					
3	2	180					
4	1	110					
5	2	100					
6	3	200					
7	3	190					
8	1	120					

D Hashing and Caching

Bovik has decided to set up a CDN. He is trying to figure out a method that clients can use to find cached web pages. That is, given a URL U, how should the client identify which CDN node to fetch the content from.

Bovik has come up with a hash function h that takes a string and maps it to a real number in the range $[0, 1)$. Assume there are 3 CDN nodes with names such that $h(node_1) = 0.1$, $h(node_2) = 0.85$, and $h(node_3) = 0.5$. When a client needs to fetch a URL and has to decide which replica to query, it picks $node_i$, such that the absolute value of the difference between $h(node_i)$ and $h(URL)$ is minimum. This scheme does *not* use circular mapping - it's just numeric closeness. This technique is called **"scheme 1"**

9. (6 points) Assuming all URLs occur with equal popularity, which node is likely to see the highest load?

Since load can be unevenly distributed in the above scheme, Bovik is not satisfied with the scheme. Instead, he thinks of a new arrangement. Let there be m CDN nodes in all; sort them using the $h(node_i)$ values. If the rank of a node is r, $(0 \leq r \leq m - 1)$, it is responsible for storing all URLs that map to the interval $[r/m, (r + 1)/m)$. This new scheme is called **"scheme 2"**.

10. (3 points) Why might a CDN with a large number of nodes (that occasionally crash and are later repaired) choose scheme 1 over scheme 2?

E Mobile IP

11. (18 points) A sender, CH (correspondent host) is sending TCP data to a mobile node, MN@home (see figure below) which is initially in its home network (HA represents the home agent). Later on, it moves to a different network (MN@away) and uses basic Mobile IP (as described in class) in order to continue receiving the TCP data from CH. In this case, the mobile node acquires a temporary address dynamically in the new (away) network (e.g. through DHCP). The IP address of a node is represented as IP-<name> (e.g. IP-CH) and the TCP port is represented as TCP-<name> (e.g. TCP-CH).

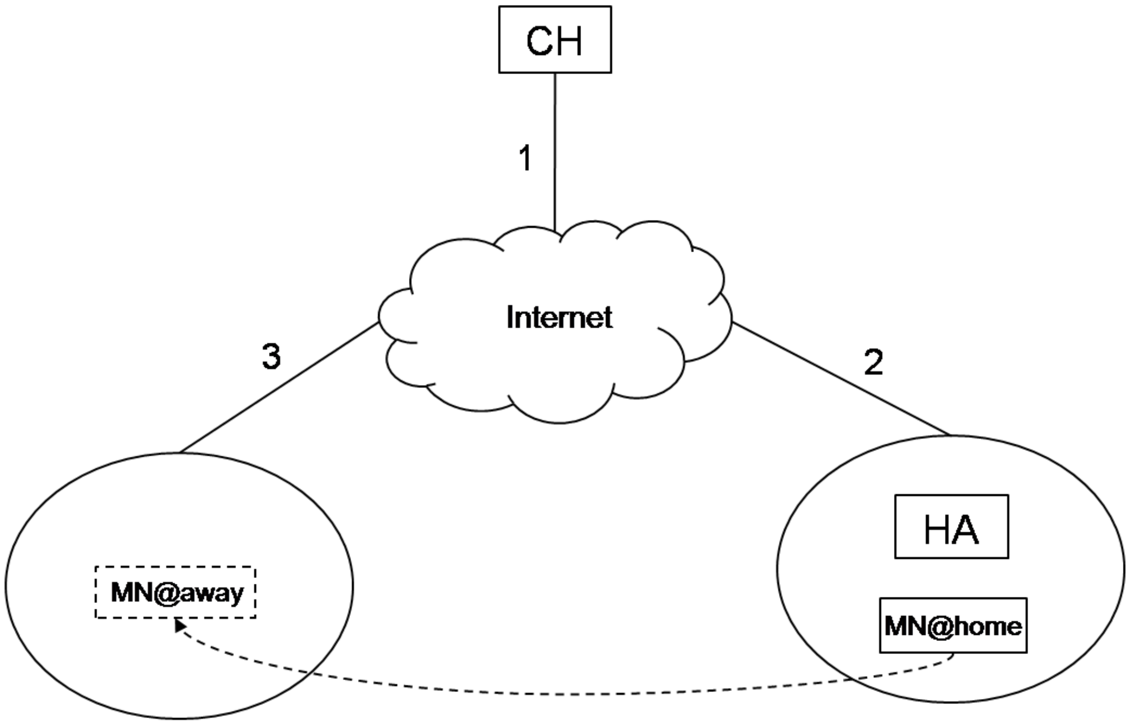

In a setting were node A is sending a TCP packet to node B, the table below shows the network and transport headers:

Header Type	Src Addr/Port	Dest Addr/Port
IP	IP-A	IP-B
TCP	TCP-A	TCP-B

(a) Consider a packet from CH to MN@home (while the mobile node is in the home network) when the packet leaves CH (on link 1). Show the network and transport layer source and destination address/port as the case may be (outermost header first – include *ALL* sets of headers including any encapsulated headers).

Header Type	Src Addr/Port	Dest Addr/Port

(b) Now, the mobile node decides to travel to the new location as shown in the figure. What is the path taken by a packet sent to the mobile node from CH?

Show the packet headers as above, when the packet leaves the home agent (i.e., on link 2) Show outermost header first – include *ALL* sets of headers including any encapsulated headers.

Header Type	Src Addr/Port	Dest Addr/Port

(c) Routers are now often configured to perform reverse path forwarding checks. That is, they check the *source* address on a packet and see if the packet came in on the interface that they would use for delivering packets to that source address. If this check fails, the packet is discarded. This is used to prevent nodes from spoofing their source address (i.e., masquerading as some other IP address). Why does this cause problems for Mobile IP forwarding as described in class?

(d) Suggest an alternative way packets from the mobile node (in the away network) could be sent to CH without breaking other parts of Mobile IP, while passing revert path forwarding checks at intermediate routers.

F Congestion Window

12. (16 points) Consider the following plot of TCP window size as a function of time:

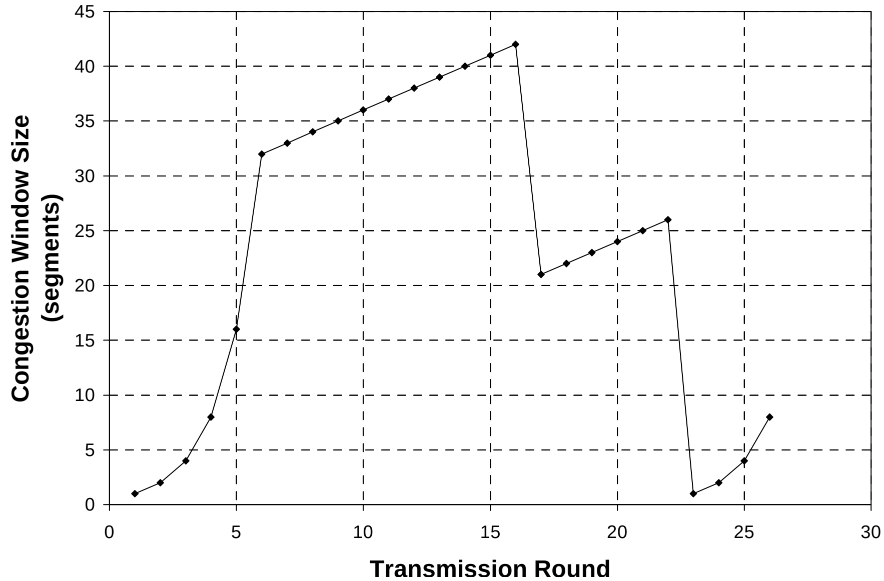

Assuming TCP Reno is the protocol experiencing the behavior shown above, answer the following questions.

(a) Identify the intervals of time when TCP slow start is operating.

(b) Identify the intervals of time when TCP congestion avoidance is operating (AIMD).

(c) After the 16th transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?

(d) What is the initial value of ssthreshold at the first transmission round?

(e) What is the value of ssthreshold at the 18th transmission round?

(f) What is the value of ssthreshold at the 24th transmission round?

(g) During what transmission round is the 70th segment sent?

(h) Assuming a packet loss is detected after the 26th round by the receipt of a triple duplicate ACK, what will be the values of the congesion-window size and of ssthreshold?

G TCP Buffering

Charlie gets a job tuning TCP to perform better. At his new workplace, they are using 10Mbit/s links and the propagation RTT between communicating nodes is 100ms.

When Charlie benchmarks TCP, he gets over 7.5Mbit/s but not 10Mbits/sec. He argues that this is great! He says TCP should only get about 7.5Mbit/s since: 1) TCP would increase its transmission rate to 10Mbit/s, 2) this would cause network congestion and packet losses, 3) TCP would drop its window by half after retransmitting data, 4) this would cause TCP to transmit data at half speed (at 5Mbit/s). Since TCP would oscillate between using the network 100% and 50%, 75% would be the average.

Charlie obviously never took 441. Your job is to help Charlie out.

13. (4 points) How much buffer space should the routers in Charlie's network have for optimal performance?

14. (4 points) You help Charlie make some detailed measurements of the network. You find that the congestion window of the connection varies between 80KB and 160KB. Estimate how much buffer space the routers in the network actually have.

15. (6 points) Charlie can't upgrade his routers. However, his crazy bosses will allow him to modify the TCP implementations. He decides to play with how TCP reacts to congestion. Instead of cutting the congestion window by 50%, he cuts it by only 33%. This seems to help, but not enough. How should he reduce the congestion window to get the full 10Mbit/sec

H Link Probing

For the following three questions assume that all network routes are symmetric, and that there are no queuing or processing delays.

16. (6 points) Ulysses dreams up a cute new idea for a networking tool which would measure the latency of each link along a route. In traceroute, a source sends packets with increasing TTL towards a destination. Traceroute reports the distance and round-trip time (RTT) to each router on the path to the destination. How would Ulysses modify traceroute to report the latency (propagation delay + transmission delay) of a particular link in the network (i.e., not the latency of the entire route to that distance)?

17. (6 points) Ulysses also decides to improve the ping tool as well. He notices that ping can be used to send packets of different sizes. He uses the ping tools on the first hop router for his computer. He notices that the router has a ping time of 10ms for small packets (64byte) and 12 ms for large packets (1500bytes). What is the propagation delay and transmission rate of this network link?

18. (6 points) Ulysses is on a roll! He combines his two tools and has traceroute send small (64 byte) and large (1500 byte) packets. How could he use the output to measure the bandwidth of a specific link along the path?

I The ACME Transport Factory

Andrew decides that computers and networks have changed a great deal since TCP was designed. He decides to create a new reliable byte-stream transport protocol called Andrew's Computer MEssaging (ACMe) Protocol.

19. (4 points) ACMe uses a sliding window (like TCP). Andrew's target is to support full channel utilization over a 100 Mbps network, with an RTT of 100 ms, and a maximum segment lifetime of 60s. what is the MINIMUM number of bits you would need for the **Advertised Window** field of your protocol header?

20. Andrew decides that slow start is too conservative for today's networks. He decides to modify slow start by changing the initial window from 1 packet to 4 packets You should assume that ACMe ACKs every packet (i.e. delayed ACKs are not supported).

    (a) (3 points) Assuming that the available link capacity and the receiver window are infinite, how many roundtrip times does it take in ACMe to send the first 10 packets and receive their ACKs? Ignore connection setup time.

    (b) (4 points) Assuming that the available link capacity and the receiver window are infinite, how many roundtrip times does it take in ACMe to send the first k packets and receive their ACKs? Ignore connection setup time.

    (c) (3 points) At most, how much time does ACMe save over TCP?

.

J Security at SA associates

For some obscure reason your pointy-haired boss asks you to make the company local area network IP broadcast-enabled, i.e., a packet sent to an IP broadcast address is sent to every node on the LAN. Since you have to keep your job, you do this, but point out to your boss that if you don't put a firewall on the system your network might be used by someone else to attack an unsuspecting victim and in the process give the entire company a black eye.

21. (4 points) Name and briefly describe an attack that would take advantage of a broadcast-enabled network.

22. (6 points) Once you have a firewall there are all kinds of things you can do to help improve the security of your network. Indicate a set of rules (one or more) that would deny outsiders from using the broadcast address on your LAN.

Src IP/mask	Src Port	Dst IP/mask	Dst Port	ACK set	Action

SA associates is developing a mobile application for playing Sudoku contests on a new phone, the Not Very Smart™phone. The NVS™has a very slow, power hungry processor. When the battery gets low it runs even slower than its usual slow pace. When a person buys the NVS™they are given a unique public/private key pair. The application works by delivering a game board, a 9x9 grid with some cells containing a blank and others containing a digit from 1-9, to the user. The user fills in the blank squares with digits from 1-9 such that each row, column, and 3x3 sub-grid contains contains no duplicate digits. The winner of the contest is the person who finishes the game the fastest without making any errors in filling out the board. (Assume that users will not show each other boards and will play the game solely on their phones.)

Describe a hybrid cryptographic scheme which uses public-key cryptography and some other cryptographic method which will allow SA associates to determine the winner of the contest. The challenge is two-fold. First, the NVS™is so slow that during game-play, the only crytographic method allowed is the XOR operation. Any other method would introduce too much delay. Second, each time the user enters a new number on the board, the application will send this number back to SA associates. The information sent back (not counting any packet headers) should be 3 bytes long, i.e., just long enough to contain the X coordinate in ascii, Y coordinate in ascii, and the number in the cell in ascii.

23. (2 points) When users purchase the game from SA associates it will contain one half of the key pair for SA associates. Will this be their public key or their private key?

24. (10 points) It can take anywhere between 100ms and 10 seconds (depending on the battery strength of the users phone) to perform a public key decrypt or encrypt operation. However, XOR always takes the

same amount of time. What should the protocol be between the user's phone and SA associates to ensure the user gets a board no one else can read and also ensure that the latency between decrypting the board be constant between all users? Notice, we are only asking you to provide confidentiality, **not** asking you to gaurantee authenticity or integrity. (Hint: SA associates can generate a string of perfectly random numbers very easily.) (Hint: You should look for a procedure which has no assymetric cryptographic operations after the board is sent to the user's phone.)

Show each step in the procedure including any information generated by SA, what is sent to whom and whether or not it is encrypted, and if encrypted how it is encrypted or decrypted. We listed the first 2 steps, please finish the procedure.

1. User sends request for board with username $\rightarrow$ SA
2. SA looks up user's public key - U^{public}

25. (3 points) You may assume that no packets are ever lost in this system. What should the application do to each entry made by the user to ensure that their answers can't be understood by anyone else? I.e., show the protocol for sending back the answers. We listed the first step in the procedure, please finish listing the procedure.

    1. User presses a key.

K P2P

26. Srini set up a P2P network to share poor jokes with other professors, but he's worried that the PJAA will shut down his centralized server, just like they did to Napster. So he set up a Chord ring for lookups and routing in his peer to peer network. Sadly, Srini's network is not popular, consisting of only four peers. The peers have successor tables as shown (the $id + 2^i$ column is there to remind you how the successor table is set up).

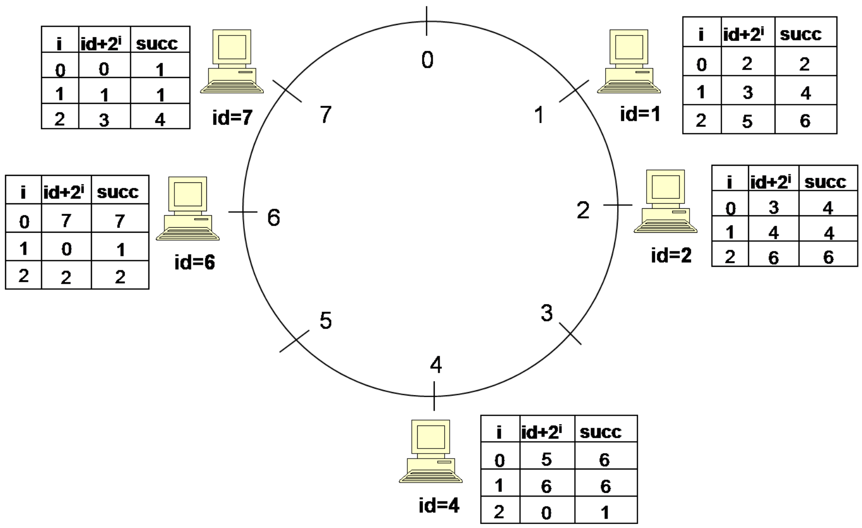

(a) (3 points) Now, jokes (items) are added to the DHT. Which node(s) would store the jokes with the following *id*'s: 3, 5, 0, 2, 6, 1?

(b) (3 points) List the nodes that will receive a query from *node 4* for *item 2*.

(c) (3 points) List the nodes that will receive a query from *node 2* for *item 5*.

(d) (6 points) Rui thinks that these jokes are awful, so he launches a DDoS attack and takes out node 2. Time passes, and the nodes converge on new routing tables that don't involve *Node 2*. Later, *Node 7* queries for *Item 5*. List the nodes that will receive this query.

The End – Phew!

L 3 Free Points for Tearing Off Page: Anonymous Feeback

List one thing you liked about the *class* and would like to see more of or see continued (any topic - lectures, homework, projects, bboards, topics covered or not covered, etc., etc.):

List one thing you would like to have changed or have improved about the class: