15-441: Computer Networks Fall 2010
Final Exam

Name: _____

Andrew ID: _____

INSTRUCTIONS:

There are 18 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.

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.

Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10
/ 8	/ 14	/ 9	/ 6	/ 6	/ 4	/ 4	/ 12	/ 18	/ 10

Q11	Q12	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Total
/ 9	/ 9	/ 6	/ 7	/ 8	/ 8	/ 4	/ 10	/ 8	/ 160

Stacks and Stuff

1. (8 points) Please answer each or the following short questions. Briefly explain your answer.
   (a) Why is changing the IP layer of the Internet protocol stack hard?

   (b) What challenge do NATs introduce to the Internet with respect to overcoming failures?

   (c) Most wireless access points and devices have the option of encrypting traffic. Which OSI layer is that encryption taking place?

   (d) Which principle of Internet design does the encryption in part (c) (appear to) violate?

DNS

2. (14 points) Harry Bovik is developing on a Web site that has multiple replicated servers located throughout the Internet. He plans on using DNS to help direct clients to their nearest lightly-loaded server, and comes up with the following hierarchical scheme. Harry has divided his server replicas into three groups (east, west, and central) based on their physical location. A typical query occurs as follows:

   - When a client makes a query for `www.distributed.hb.com`, the root `.com` name server is contacted first. It returns the name server (NS) record for `ns1.hb.com` (along with a corresponding A record). The TTL of this record is set to 1 day.
   - The `ns1.hb.com` name server is then queried for the address. It examines the source of the name query and returns an NS record for one of {`east-ns, central-ns, west-ns`}`.distributed.hb.com` (along with a corresponding A record). The choice of which name server is based on where ns1 thinks the query came from.
   - Finally, one of {`east-ns, central-ns, west-ns`}`.distributed.hb.com` is contacted and it returns an address (A) record for the most lightly loaded web server in its region.

   The following questions are based on this design.

   (a) Harry's name server software has only two choices for TTL settings for A and NS records - 1 day and 1 minute. What are reasonable TTLs for the following records? Briefly explain your choice.

   NS record for {`east-ns, central-ns, west-ns`}`.distributed.hb.com`:

   A record for {`east-ns, central-ns, west-ns`}`.distributed.hb.com`:

   A record returned for the actual Web server:

   Harry's Web site is especially popular among CMU students. The CMU network administrators estimate that there is one access from CMU every 6 minutes. Each access results in the application resolving the name www.distributed.hb.com. Assume the following:

   - No other queries are made from CMU
   - All CMU clients use the same local name server
   - Web browsers do not do any caching on their own.

   (b) How many accesses will be made to the following name servers each HOUR to resolve these queries? Use your answers to the previous question, and explain your calculation.

ROOT:

ns1.hb.com:

one of {east-ns, central-ns, west-ns}.distributed.hb.com:

(c) Harry finds that many people are far away (i.e. communication has high latency) from the name servers that they use. Why might this be a problem for his scheme?

Security

3. (9 points) Below are several scenarios describing simple uses of cryptographic schemes we have covered in 441. For each scenario, circle "correct" if the scenario describes a valid use of the mechanism as described in class. Otherwise circle "incorrect". In either case, provide *one sentence* explaining the vulnerability it exposes or why there is no vulnerability.

   (a) Michael wants to transmit project2 grades from his home computer to Peter at CMU. He is worried that some enterprising 441 student may have hacked a router along the path and might modify the message to improve their grade and win the project2 contest. So when Michael sends a message M to Peter, he also calculates H = Hash(M) and appends H to the message. Peter receives M and H, and calculates H' = Hash(M), only accepting the message as valid if H' = H . You can assume that Hash is a secure hash function that is one-way, collision resistant, and pre-image resistant.

   correct / incorrect

   (b) Peter wants to send the top-secret solutions for the final exam to Wittawat, but Wittawat's email server is down. However, Wittawat promised to check the 441 bboard regularly to see if Peter posted any messages for him. During the first day of class, Wittawat gave everyone at the lecture (including Peter) his public key $K_{Wittawat}$. Only Wittawat knows his private key, $K^{-1}_{Wittawat}$. Knowing that Wittawat will recognize the correct answer key based on their past discussions, Peter encrypts the answer key with $K_{Wittawat}$ and posts it to the bboard.

   correct / incorrect

   (c) George and Vyas both share a secret key with a Key Distribution Center (KDC). We call these keys $K_{George,KDC}$ and $K_{Vyas,KDC}$ respectively. George wants to establish a shared symmetric key with Vyas, so George authenticates to the KDC using $K_{George,KDC}$ and the KDC replies with $Encrypt_{K_{George,KDC}}(K_{Vyas,KDC})$. George and Vyas then communicate using the shared secret key $K_{Vyas,KDC}$.

   correct / incorrect

4. (6 points) Bob and Jane need to communicate to decide which TA is going to grade the next homework. They have a shared secret, K_{Profs} that allows them to create unforgable message authentication codes (MAC) so that Bob can verify that Jane did in fact create any message that is received. Bob and Jane have a simple protocol: Bob sends a "Who grades HWX?" message to Jane in plain text, and Jane replies with one of three messages: $M1 = MAC_{K_{Profs}}$("TA-1"), $M2 = MAC_{K_{Profs}}$("TA-2"), or $M3 = MAC_{K_{Profs}}$("TA-3"). When Bob receives either M1, M2, or M3, he verifies the MAC using K_{Profs} and knows who will grade the next homework.

   (a) This protocol is insecure. A malicious TA on a router between Bob and Jane may be able to avoid ever having to grade a homework! In one sentence, describe the attack.

   (b) What simple change to the above protocol could defend against this attack?

Web and Peer-to-Peer

5. (6 points) Provide three reasons a company might prefer to pay Akamai to host their webpage instead of putting it onto a peer-to-peer network (such as Napster) for free.

6. (4 points) We saw no examples of of chunk-based peer-to-peer networks that use flooding. What would make such a network inefficient?

7. (4 points) While not strictly true, people tend to view hash tables as offering constant time look up in practice. (The possibility of a bad hash function leading to many collisions is why it is not strictly constant time.) Distributed Hash Tables (DHTs), on the other hand, only offer $O(\log n)$-time lookup where n is the number of nodes in the DHT. This is odd, since the two appear to be equivalent, i.e. simply map each entry in a traditional hash table onto a node in the DHT. What property or requirement of a DHT makes this approach impractical in practice.

TCP and Transport

8. (12 points) The TCP flow control window size is carried in a 16 bit field. Please answer the following questions and briefly motivation your answer.

   (a) Why is this a problem in today's networks?

   Since it is not practical to change the TCP header format, the only way of increasing the maximum flow control window size is to use TCP options.

   (b) Present a solution to increase the maximum TCP windows size (based on options) that works but may be slow in practice.

   (c) Present a solution to increase the maximum TCP windows size (based on options) that works and is also efficient in practice.

   (d) An alternative to using TCP options to extend the flow control window size would have been to add a longer window field to IPv6 when it was defined. Would this have been a good idea?

   Wrong layer: window size is specific to a transport function supported by TCP (but not all transport protocols), so it would waste header space for these other protocols.

9. (18 points) The Transmission Control Protocol uses a method called congestion control to regulate the traffic entering the network. The behavior of TCP congestion control can be represented as a graph in which the x-axis indicates the time, and the y-axis indicates congestion window size. Please use the graph shown below to the answer the following questions. Note that the graph does not explicitly show timeouts, but you should be able to figure out when timeouts happened based on the events shown.

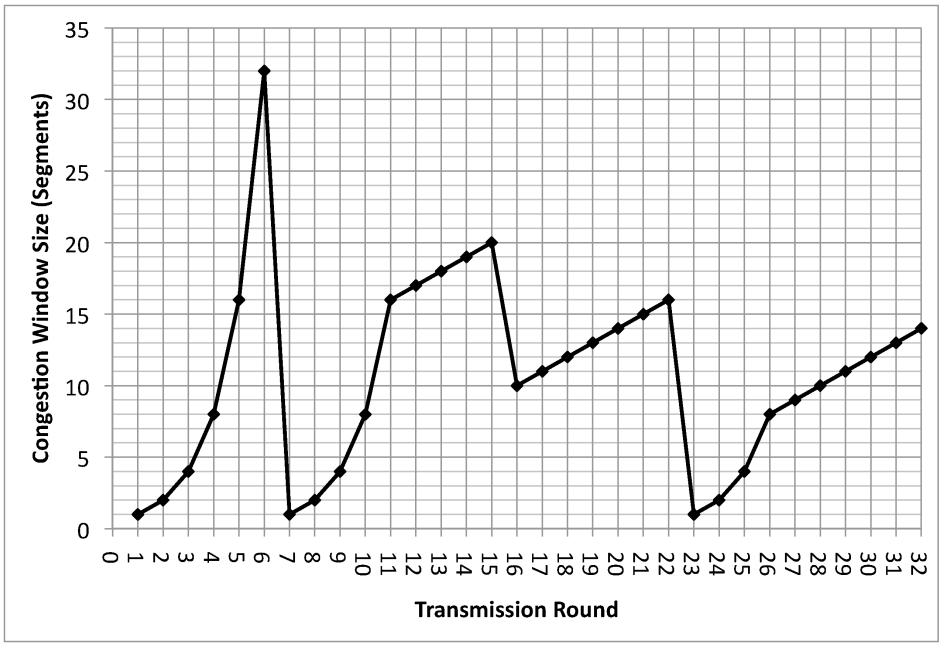

(a) **Slow Start:** give two reasons why slow start is used, and explain why it does a better job than congestion avoidance for that function.

(b) **Slow Start:** identify the intervals of time when TCP slow start is operating. For each interval, identify which of the above reasons apply and do not apply and explain why.

(c) **Congestion Avoidance:** identify the intervals of time when TCP congestion avoidance is operating. Why should congestion avoidance be used instead of slow start during these intervals? Please clearly identify one specific reason.

(d) **Fast Retransmission:** identify the intervals of time when TCP fast retransmission is used. Please explain what fast retransmission does and how it is triggered.

(e) **Fast Recovery:** identify the intervals of time when TCP fast recovery is operating. What does fast recovery do and exaplain why is it beneficial.

(f) **Lack of fast recovery:** identify the interval(s) of time when fast recovery could have happened, but did not. Identify one specific example of a circumstance that may prevent fast recovery from happening.

LANs

10. (10 points) Given the extended LAN shown in the figure below, answer the following questions.

    The LAN's in the figure are labeled LAN A-F and the bridges in the figure are labeled Bridge 1-4. Bridge i has an ID of i which is used as the tie breaker, i.e. bridges with a lower id get selected. The hosts are labeled Host A-F.

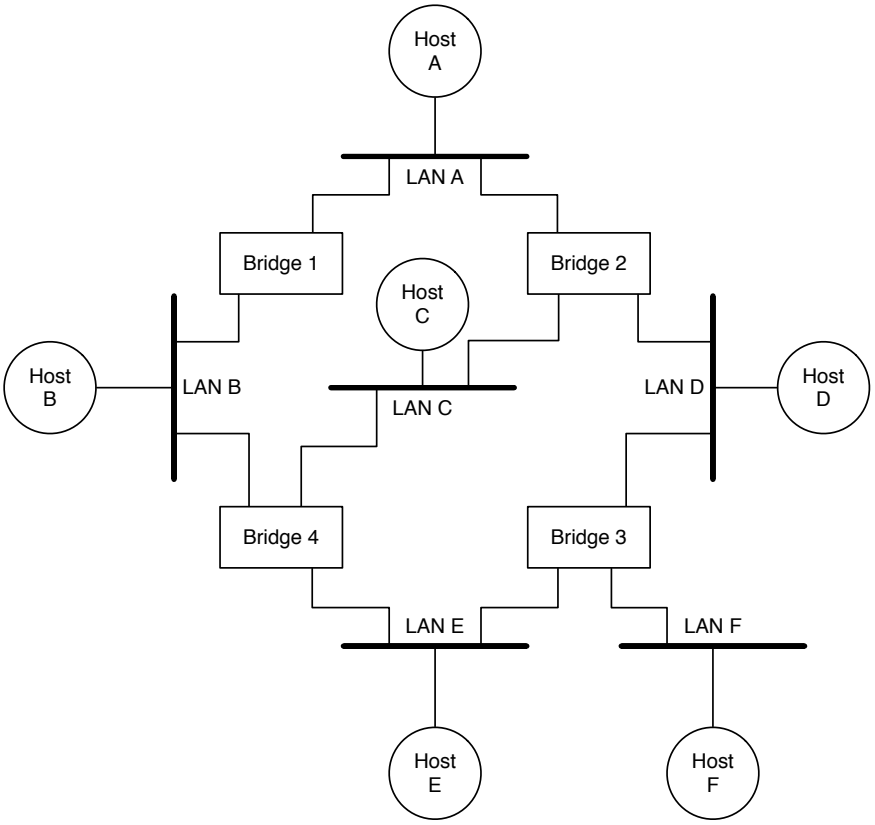

(a) Cross out ports (in the figure above) which are not selected by the spanning tree algorithm.

Suppose that the forwarding table for each bridge is empty. Let each host runs its Ethernet card in promiscuous mode where each host listen to all frames instead of a frame sent to it.

(b) If Host F sends to Host A, who can get that message?

(c) Quickly after Host F sends a message. If Host B sends to Host F, who can get that message?

11. (9 points) Five prisoners are locked up in adjacent cells in a prison. They would like to communicate with each other but the walls and doors are too thick. One day, one of the prisoners discovers that if he hits the water pipe in his cell with a metal spoon, the sound travels to two cells in each direction, i.e. the sound from cell i can be heard in cells i-2, i-1, i+1, and i+2, assuming these cells exist. After some experiments, they discover this is true for all the cells.

    Over lunch, they decide to define a protocol that will allow efficient communication. One of the prisoners has taken 441 and argues that this is very much like an Ethernet so they decide to use the Ethernet protocol over their Water Pipe Network. The prisoners planning the break are in five adjacent cells that are lined up in a row. Unfortunately, there are some problems. Can you help them?

    (a) Ethernet uses CSMA/CD as its medium access mechanism. Can you explain how the three concepts that are used in CSMA/CD (CS, MA, and CD) map onto specific aspects of this network?

    CS:

    MA:

    CD:

    (b) In the Water Pipe Network, not all cells can hear each other. What mechanism could you use so all inmates can talk to each other?

    (c) As they get closer to the night they plan to escape, traffic on the Water Pipe network increases. Unfortunately, they discover that using CSMA/CD over the Water Pipe Network results in a significant packet loss rate. Can you identify the problem responsible for the packet losses and propose a solution?

Remember Me?

12. (9 points) Suppose that computer A is sending a file to computer B using a private Ethernet with no other computers using it. They are connected by 100m of wire. Bits travel at the rate of 2×10^8m/s in this wire. Suppose the Ethernet has a bandwidth of 10^9 bits per second ("gigabit Ethernet").

    Answer the following questions. For answers that are numbers, you must use units. However, you do not need to carry out the calculations to write it as a single number. You may provide a mathematical expression that yields the needed number instead.

    (a) (3 points) Based the on the provided information, what is the latency (i.e., propagation delay) of the connection? (Estimate if need be.)

    (b) (3 points) How long would it take for 10^6 bits to finish traveling from computer A to computer B? (Estimate if need be.)

    (c) (3 points) Suppose you measured the time it takes to transmit a 10^6 bit file from computer A to computer B. To not include the time it takes for either computer to process the file, you start measuring from time the first bit of the file leaves computer A until last bit of the file reaches computer B. Furthermore, you make sure that the computers are fast enough that they do not limit the speed of transmission. Nevertheless, the time you measure is longer than time you calculated above. What factors could have resulted in this?

Routers and Routing

13. (6 points) For each part below, select the one answer that best matches the question:

(a) Crossbar switching and virtual output queues are *both* good examples of

A. Techniques that, when combined, support latency guarantees
B. A trade off between space (memory, wires, etc.) and time
C. A solution that is too expensive to be widely adopted
D. NP problems being solved in practice

(b) Ternary content addressable memory is a good example of

A. Running into the limits of optical-electric conversion
B. A violation of fate sharing
C. A solution being too expensive to widely adopt
D. Hardware giving off too much heat

(c) Why do routers use binary search *tries* instead of binary search *trees*?

A. Tries are easier to implement in hardware
B. Tries are easier to implement in software
C. To take advantage of the hierarchical nature of IP addresses
D. To avoid head-of-line blocking

14. (7 points) Which of the following is true about BGP? (Circle all letters that apply)

T F BGP uses a distance vector protocol to improve route stability

T F A BGP router always picks the path with the least number of router hops to the destination.

T F A BGP router always picks the path with the least number of AS hops to the destination.

T F If we denote a Customer → Provider link by a -1, a Peer → Peer link by a 0, and a Provider → Customer link by a +1, the following sequence is a valid BGP path (-1,-1,-1,0,+1,+1)

T F An Autonomous System will announce routes learned from its customers to its peers.

T F An Autonomous System will announce routes learned from its peers to other peers.

T F If an Autonomous System learns of 5 different routes to a destination prefix, it will announce all 5 routes to its neighbors.

15. (8 points) Which of the following is true about MPLS/ATM:

T F ATM and MPLS has a much simpler route lookup than IP since they use fixed length tags.

T F During circuit setup, a globally unique connection identifier is created to route the circuit through the network.

T F ATM and MPLS can force packets to follow a pinned path in the network.

T F ATM uses 53 bytes cells since prime numbers provide strong security.

16. (8 points) We have 5 routers labeled A-E. Suppose we have the forwarding tables shown below after RIP is stable. Let all links have cost 1.

Forwarding Table for A

Destination	Cost	Next Hop
A	0	-
B	1	B
C	2	B
D	1	D
E	2	D

Forwarding Table for B

Destination	Cost	Next Hop
A	1	A
B	0	-
C	1	C
D	1	D
E	1	E

Forwarding Table for C

Destination	Cost	Next Hop
A	2	B
B	1	B
C	0	-
D	2	E
E	1	E

Forwarding Table for D

Destination	Cost	Next Hop
A	1	A
B	1	B
C	2	B
D	0	-
E	1	E

Forwarding Table for E

Destination	Cost	Next Hop
A	2	D
B	1	B
C	1	C
D	1	D
E	0	-

(a) If a message is originated from A and a destination is E. Which path does it take?

(b) If a message is originated from C and a destination is D. Which path does it take?

(c) Give a diagram of a possible network consistent with these tables.

17. (4 points) We use CIDR. Without using longest prefix matching, a forwarding table looks like this.

If we use longest prefix matching, we can combine a few entries together. What is a table with a minimum number of entries that still be able to forward packet correctly?

Prefix	Outgoing Interface
128.0.0.0/11	eth1
128.16.0.0/12	eth1
128.24.0.0/12	eth2
128.32.0.0/12	eth2
128.40.0.0/12	eth1
128.48.0.0/11	eth1
128.64.0.0/9	eth0
128.128.0.0/10	eth0
128.160.0.0/11	eth1
128.176.0.0/11	eth0
128.192.0.0/9	eth0
default	eth3

$\Rightarrow$

Prefix	Outgoing Interface
default	eth3

Mobile IP

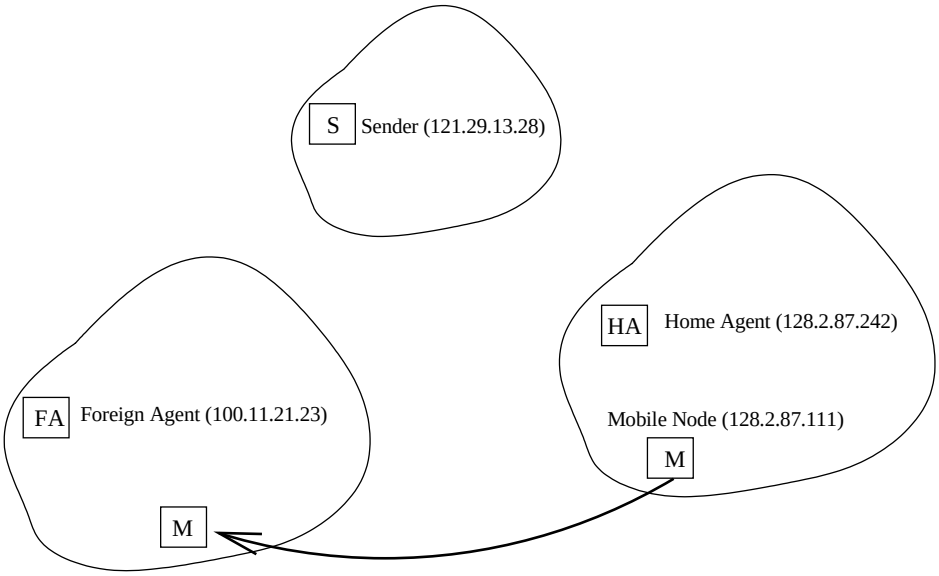

18. (10 points) A sender S is sending TCP data to a mobile host M (see Figure). Initially the mobile host is in its home network. Later on it moves to a different network and needs to use Mobile IP in order to receive data from S. All local area networks are Ethernets.

Part 1: The sender S sends TCP data to the mobile node while is in its home network.

(a) Each packet has more than one header as more than one protocol is being used to send it. Name the protocols that are contributing headers to the packets starting with the layer 2 protocol going up to the transport layer protocol.

(b) What are the source and destination IP addresses in the packet?

Part 2: The correspondent host is sending TCP data to the mobile node which has moved to the foreign network.

(c) What headers does each packet have (names only), starting with the layer 2 header and up to the transport layer header, as the packets arrive at the mobile's home agent?

(d) What headers does each packet have (names only), starting with the layer 2 header and up to the transport layer header, as the packets arrive at the mobile's foreign agent?

(e) What are the source and destination IP addresses in the packet in (d)?

QoS

19. (8 points) You are in charge of doing traffic enforcement for a large ISP and a customer gives you the traffic pattern that consists of a bursts of traffic sent at rate R and of length T1, separated by periods of length T2 where no traffic is sent. You are asked to specify the tightest token bucket parameters that will let this traffic stream through (i.e. no packets will be dropped). "Tightest" means that you first minimize the token bucket rate and, for the minimum token bucket rate, you then minimize the bucket size.

    (a) That is the minimum token bucket rate Rb?

    (b) What is the minimum token bucket size Sb?

    (c) The customer complains that some packets get dropped but he also admits that they may occasionally have shorter bursts that are sent at a slightly higher rate. What would you suggest to accommodate this?