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

EE122
MIDTERM EXAMINATION
Monday, 18 October 2010
Scott Shenker

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, iphone, laptop computer, or other calculation aid. Please put them away right now! You MAY use one 8.5x11 double-sided crib sheet, as densely packed with notes, formulas, and diagrams as you wish. All work should be done on the attached pages. Don't be scared off by the number of questions; many of them are simple warm-up exercises, and you should have plenty of time to finish them all the questions.

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.

Please write your SID on each page!

________________________________ (Name—Please Print!)

________________________________ (Signature)

SID: ________________________________

Discussion Section (Day/Time): ________________

Question	Possible Points	Points Obtained
1	7	
2	5	
3	5	
4	5	
5	10	
6	5	
7	10	
8	9	
9	7	
10	10	
11	9	
12	5	
13	10	
14	3	
Total	**100**	

1. General Multiple Choice [7 points]

1. There are three standard techniques in computer science:
(a) inserting a layer of indirection
(b) creating a hierarchy, and
(c) using caching.
Use each of the concepts once to answer the following three questions:

i) Which of these is most useful for achieving better scalability? **B**

ii) Which of these is most useful for providing more flexibility? **A**

iii) Which of these is most useful to improve performance? **C**

2. There are two general guidelines in networking:
(a) The End-to-End Principle, and (b) Fate Sharing.

i) Which of these guides the placement of functionality? **A**

ii) Which of these guides the placement of state? **B**

3. Consider the following four elements of a packet header:
(a) source address,
(b) flow label,
(c) fragmentation fields, and
(d) source port.

i) Which of these header elements is in IPv4 but not IPv6? **C**

ii) Which of these header elements is in IPv6 but not IPv4? **B**

iii) Which of these header elements is in both IPv4 and IPv6? **A**

iv) Which of these header elements is not in IPv4 or IPv6? **D**

4. Consider the following:
(a) Ethernet header, (b) IP header, and (c) TCP header

i) Routers use the contents of which header? **B**

ii) Switches use the contents of which header? **A**

iii) Hosts direct packets to the correct process based on which header? **C**

5. When IPv4 packets are fragmented, where does reassembly happen?

C

Choose one answer from the list below:

(a) In the next-hop router
(b) In the next-hop switch
(c) In the receiving end host
(d) In the textbook

6. Host A is transferring a large file to host B over a 10gigabit-per-second link (that is, 10^{10}bits/sec). This link has a one-way latency of 10msec. How big would the sliding window have in order to "fill the pipe"? That is, how large does the window have to be in order to take advantage of the full link bandwidth, assuming there is no other traffic on the link?

D

Answer this question using one of the letters below:

(a) 10^8 bytes
(b) 10^8 bits
(c) $2x10^8$ bytes
(d) $2x10^8$ bits
(e) $8x10^8$ bytes
(f) $8x10^8$ bits

2. Acronyms [5 points]

Match the acronyms to the questions, using each acronym once: BEB, CDN, CIDR, CSMA, CSMA/CD, ICANN, MTU, NAT, NIC, RFC, RTT, TCP, TLD, TTL, UDP, WTF

1. The documents used by the IETF to describe protocol standards: **RFC**
2. Two transport-level protocols: **TCP, UDP**
3. A term for a host's network adaptor: **NIC**
4. The largest packet that can be sent over a link without requiring fragmentation: **MTU**
5. Style of multiple access used by slotted Aloha: **CSMA**
6. Style of multiple access used by Ethernet: **CSMA/CD**
7. The organization responsible for allocation of Internet address blocks to the regional Internet registries (RIR): **ICANN**
8. The term for a primary DNS zone such as .com or .org or .fr: **TLD**
9. The time it takes a packet to reach the destination and have a response return: **RTT**
10. A counter in the IP header that keeps packets from circulating forever: **TTL**
11. The backoff algorithm used by Ethernet: **BEB**
12. Akamai is an example of one of these: **CDN**
13. The addressing structure currently used in IP: **CIDR**
14. This lets several computes share a single IP address: **NAT**
15. Your reaction upon seeing the length of this test: **WTF**

3. Miscellaneous Short Questions [5 points]

Narrow Waist: Consider the Internet hourglass, consisting of:
- (a) Protocols **a**bove the waist
- (b) Protocols **b**elow the waist
- (c) Protocols at the waist

For each of the following, identify whether the protocol is above, below, or at the waist by using (a), (b), or (c):

1. TCP:	**A**	
2. Ethernet:	**B**	
3. SONET:	**B**	
4. HTTP:	**A**	
5. IP:	**C**	

Ethernet delivery: The Ethernet adaptor for node A hears a frame on the wire and then passes it up to host A. Which of the following are legitimate reasons for the adaptor passing the packet to the host (list all that are legitimate):
- (a) The adaptor was in promiscuous mode.
- (b) The adaptor felt like it.
- (c) The destination MAC address was a multicast group to which A belonged.
- (d) The frame's CRC was damaged.
- (e)The destination MAC belongs in the same segment as host A.
- (f) The destination MAC address matched that of the adaptor
- (g) The destination MAC address was the broadcast address

A, C, F, G

4. Packet Recovery [5 points]

Consider hosts A and B connected by two links, L1 and L2, with a switch S in the middle. Both links drop packets: L1 has a loss rate of 0.0312%. L2 has a loss rate of 17.567%.

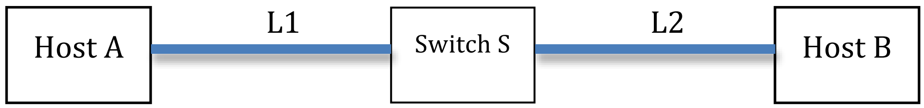

Consider two possible scenarios, and in each scenario don't worry about how the packet losses are detected.

(a) End-to-end recovery: Whenever a packet is lost (on L1 or L2), host A resends it to host B. Switch S does nothing but forward packets.

(b) In-network recovery: Whenever a packet is lost on L1, host A resends it to host B; whenever a packet is lost on L2, the switch S resends it to host B. Switch S buffers incoming packets from L1 and retransmits them when they are lost on L2.

Host A is trying to send a single packet to host B. This may require some retransmissions by host A and/or the switch S. A transmission on link L1 means that host A sent a packet (whether or not it was lost on L1 or later). A transmission on link L2 means that switch S transmitted a packet (whether or not it was lost on L2 after transmission), regardless of whether S was merely forwarding a packet from host A, or sending a retransmission to recover a loss on L2 as in scenario (b).

In answer to the questions below, write (a) if (on average) there are more transmissions in scenario a, write (b) if (on average) there are more transmissions in scenario b, and write (s) if (on average) there are the same number of transmissions in both scenarios.

i) In which scenario are there more transmissions on link L1? **A** ☐

ii) In which scenario are there more transmissions on link L2? **S** ☐

Hint #1: When sending over a single link L with failure rate f, the average number of transmissions before success is 1/(1-f).

Hint #2: If you think this is complicated, think again. You should not have to do any calculations involving the actual drop rates.

5. Routing Tables [10 points]

Consider the following routing table:

0000/4	⇒	port 1
0001/4	⇒	port 1
0010/4	⇒	port 1
0011/4	⇒	port 1
0100/4	⇒	port 2
0101/4	⇒	port 3
0110/4	⇒	port 2
0111/4	⇒	port 2
1000/4	⇒	port 1
1001/4	⇒	port 1
1010/4	⇒	port 1
1011/4	⇒	port 1
1100/4	⇒	port 4
1101/4	⇒	port 4
1110/4	⇒	port 1
1111/4	⇒	port 1

In the above table, each packet will match one and only one prefix. Rewrite the routing table using Longest Prefix Match, using the minimal number of prefixes. One of your entries should be for the default route (i.e., default ⇒ port X) and the rest should be written in the prefix notation used above (i.e., to avoid confusion, **do not** write it in the dotted-decimal form a.b.d.c/X; write it instead as "bit pattern"/X as above).

Default	**⇒**	**port 1**
01/2	**⇒**	**port 2**
0101/4	**⇒**	**port 3**
110/3	**⇒**	**port 4**

A variety of notations were accepted for the prefixes…..such as 01/2, 0100/2, etc.

6. Network Address Translation and FTP [5 Points]

When you connect to an FTP server, you are actually making two connections. First, the so-called control connection is established, over which FTP commands and their replies are transferred. Then, in order to transfer a file or a directory listing, the client sends a particular command over the control connection to establish the data connection. The data connection can be established two different ways, using active mode or passive mode.

(a) Passive Mode: the client sends the PASV command to the server, and the server responds with an address. The client then issues a command to transfer a file or to get a directory listing, and establishes a secondary connection to the address returned by the server.

(b) Active Mode: the client opens a socket on the local machine and tells its address to the server using the PORT command. Once the client issues a command to transfer a file or listing, the server will connect to the address provided by the client.

In both cases, the actual file or listing is then transferred over the data connection.

Consider a client behind a NAT box trying to use FTP using one of these two modes. Assume the NAT has not been modified to understand the semantics of FTP.

Does the client's FTP attempt succeed using Passive Mode? (yes/no) **Yes** ☐

Does the client's FTP attempt succeed using Active Mode? (yes/no) **No** ☐

7. Domain Name Service [10 points]

a) What are the DNS record types for the following:

IP address of a given hostname: **A**

Hostname of given domain's mail server: **MX**

Hostname with a given IP address: **PTR**

Hostname of authoritative DNS server for a given domain: **NS**

b) Consider a scenario with a Host (call it A), its local DNS server (call it B), and the rest of the DNS infrastructure. A asks B to resolve foo.com. Write down the steps taken to resolve foo.com *and* respond to A. In each step identify who B contacts. There is no need to use all the steps below (and if you think you need all eight steps, think again).

Step 1: B contacts ______ **Root DNS Server**________________

Step 2: B contacts ______ **TLD DNS Server for .com**__________

Step 3: B contacts ____ **Authoritative DNS server for foo.com**__

Step 4: B contacts _______ **Host A** _____________________

Step 5: B contacts ______________________________________

Step 6: B contacts ______________________________________

Step 7: B contacts ______________________________________

Step 8: B contacts ______________________________________

8. Link-Layer Framing [9 points]

Sentinel framing uses one bit pattern to designate the start of a frame, another bit pattern to designate the end of a frame, and a bit-stuffing rule to prevent these two special patterns from occurring in the payload of the packet. The bit-stuffing rule inserts a 0 after a particular pattern of bits occurs in the payload.

An adaptor uses sentinel framing, and is handed a payload of

100111101100110110

It sends the string of bits:

00111010011011011001100110001111

What bit pattern is used to signify the start of a frame?

001110

What bit pattern is used to signify the end of a frame?

001111

What is the particular pattern used in the bit-stuffing rule?

0 stuffed after pattern 0011

Hint: the bit-stuffing rule *must* ensure that the frame-start and frame-end patterns cannot occur in the payload of a frame, so keep this in mind when checking your proposed solution.

9. Self-Learning [7 points]

Consider the following set of four Ethernet segments connected by a central switch (not a hub). Each Ethernet segment has two hosts on it. Assume at the beginning that the switch has no state about where each of the hosts reside, and will use self-learning to build its routing table during the sequence of transmissions listed below. For each step in this sequence of transmissions, list which ports the Ethernet frame is forwarded to by the switch (and, if the frame is not forwarded to any ports, just write "drop").

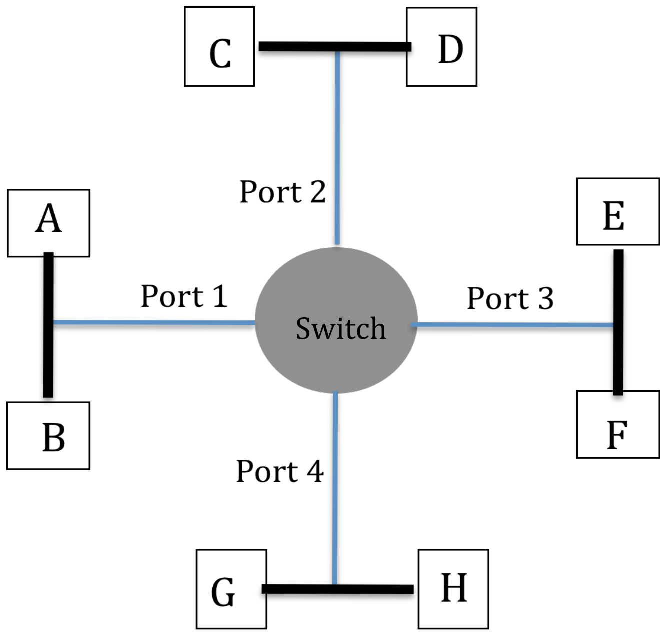

i) A sends to E 2,3,4

ii) A sends to B 2,3,4

iii) G sends to E 1,2,3

iv) E sends to A 1

v) E sends to B 1,2,4

vi) A sends to B 2,3,4

vii) H sends to G Drop

10. Spanning Tree [10 points]

The nodes in the graph below construct a spanning tree using the standard spanning tree protocol. The bridges (or switches), Bi, have MAC address i. After the spanning tree protocol is run, which links are in the spanning tree?

L1, L2, L5, L6, L9, L10

Assume node B1 fails, and the bridges reform a spanning tree. Which links are now in the spanning tree?

L4, L7, L9, L10, L12

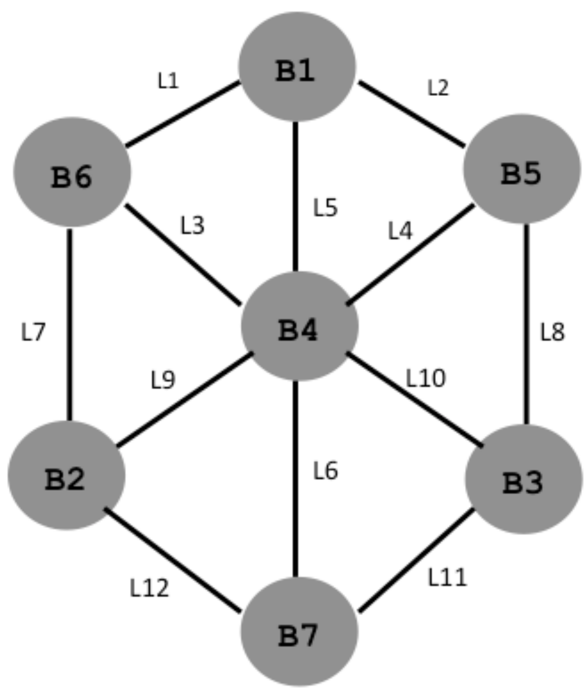

Hint: (i) the root is the bridge with the lowest ID, and (ii) when there is more than one shortest path to the root, the path whose first hop goes through the bridge with the lower ID is chosen. Also, each link is considered to have a length of one when computing path lengths.

11. Slotted Aloha [9 points]

Consider slotted Aloha where the probability to resend after a collision is p. Consider the case where there are N hosts sending, and they all collide on slot 1.

(i) What is the probability that there will be a successful transmission on the next slot (expressed in terms of N and p)?

$Np(1-p)^{(N-1)}$

(ii) For what value of p is this maximized?

$p = 1/N$

(iii) When p is set to this optimal value, what is the resulting answer for (i) in the limit of large N?

$1/e$

12. Ethernet [5 points]

Hosts A and B collide on an Ethernet (in what we'll call slot 0). Each only has a single packet to send, and no other nodes have data to send during the period in question. Which of the following three resend sequences (described below) are possible.

B

Below, "slot" refers to minimum packet transmission times, and each sequence starts with the same initial collision.

Sequence (a):

Slot 0: A and B both send (their first collision)
Slot 1: A and B both silent (idle)
Slot 2: A and B both silent (idle)
Slot 3: A sends, B silent (success)
Slot 4: A silent, B sends (success)

Sequence (b):

Slot 0: A and B both send (their first collision)
Slot 1: A and B both send (collision)
Slot 2: A and B both silent (idle)
Slot 3: A sends, B silent (success)
Slot 4: A and B both silent (idle)
Slot 5: A silent, B sends (success)

Sequence (c):

Slot 0: A and B both send (their first collision)
Slot 1: A sends, B silent (success)
Slot 2: A and B both silent (idle)
Slot 3: A and B both silent (idle)
Slot 4: A silent, B sends (success)

13. HTTP [10 points]

Consider a client accessing a base HTML page with four URL references to images to be displayed in the page. None of the images can be requested until the entire base page is retrieved. Assume that it takes time T1 to establish a TCP connection, time T2 to retrieve the base page, and time T3 to retrieve each of the images. Ignore issues of bandwidth constraints, congestion, server processing time (i.e., assume the web server can process requests as fast as they come in), and the time it takes to close a TCP connection, and assume all HTTP request packets are tiny (compared to the returning responses).

For each of the scenarios below, describe the total time (in terms of T1, T2 and T3) it takes to retrieve the base page and its embedded images.

Sequential requests with nonpersistent TCP connections:

(T1+T2) + 4(T1+T3)

Concurrent requests with nonpersistent TCP connections:

(T1+T2) + (T1+T3)

Sequential requests within a single persistent TCP connection:

(T1+T2) + 4T3

Pipelined requests within a single persistent TCP connection:

(T1+T2) + T3

14. Suggestion for Improving Course [3 points]

(there are no right or wrong answers, but the lack of answers will cost you points!)

i) Is the course going too slowly or too quickly? ☐

ii) Should the lectures be harder or easier? ☐

iii) Should the homeworks be harder or easier? ☐

iv) Do you like the additional material on the history and politics of networking, or should the lectures focus only on the technical details? ☐

v) Should the lectures contain more worked examples? ☐

vi) Is the project too easy or too hard? ☐

vii) Are you getting the support you need in section? ☐

viii) Any other suggestions or comments? Please write them below: