

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

EE122
MIDTERM EXAMINATION
Monday, 17 October 2011
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, packed with notes, formulas, and diagrams (but no font smaller than 10pt). 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.

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.

There are two blank sheets at the end of the test. Feel free to rip those off and use them as a place to work out some of the problems. Please show your work as needed on the test itself, but these pages will give you a place to get your thoughts in order.

Please write your SID on each page!

_____ (Name—Please Print!)

_____ (Signature)

SID: _____

Discussion Section (Day/Time): _____

Question	Possible Points	Points Obtained
1	20	
2	15	
3	18	
4	15	
5	14	
6	16	
7	20	
8	15	
Total	133	
Extra	20	
Fun	0	0

Question 3 was marked as only having 15 points on a few exams. Please check to make sure you were marked with this problem having 18 points. If not, contact us for a regrade.

1. True or False [20 points]

For each of the following, please indicate T or F

- F** i) In the Internet, all packets from A to B follow the same path through the network.
- T** ii) In circuit switched networks, all bits in the same transmission from A to B follow the same path.
- F** iii) The Internet designers picked Best Effort service because it satisfied the needs of all applications, even those with strict delay requirements.
- T** iv) Statistical multiplexing means that when you have more flows the overall burstiness is reduced (as in the law of large numbers).
- T** v) Time Division Multiplexing assigns each flow a slot in a frame, and one flow cannot borrow another flow's slot if it isn't being used.
- T** vi) IP fragments are reassembled at the destination
- F** vii) Performance is often more important than flexibility but typically only in hindsight.
- F** viii) DHCP is a transport level protocol
- F** ix) Link-state routing scales better than distance-vector.
- T** x) ACK packets can also carry data.
- T** xi) There are other transport protocols used in the Internet besides UDP and TCP
- F** xii) There are other internetworking protocols used in the Internet besides IPv4 and IPv6.
- F** xiii) The distributions for most Internet measurements (such as DNS access patterns and flow lengths) are exponential. One need only worry about the low end of the distribution, because the probabilities at the higher end die out so quickly.
- T** xiv) CIDR is the addressing scheme currently used in the Internet.
- F** xv) A RST message causes TCP to change its sequence numbers to zero and start again.
- T** xvi) ISPs in North America get their Internet addresses from ARIN.
- F** xvii) NAT is used to allocate addresses to hosts on a dynamic basis, while DHCP allows several hosts to use a single IP address.
- T** xviii) The Sequence Number in a TCP packet refers to the first byte in the packet payload.
- F** xix) The original Internet addressing scheme only supported 256 host addresses.
- T F** xx) I feel much better now that I've finished the first question on the test.

2. Short Questions: Principles/Foundations [15 points]

A. [3] We are faced with three basic design decisions

- a) How to break the overall system into modules
- b) Where these modules are implemented:
- c) Where the associated state is stored:

We have three general principles to guide us. Which of the above design decisions do the following help clarify: (choose either a, b, or c for principle below)

- i. Fate Sharing: **c**
- ii. End-to-End Principle: **b**
- iii. Layering: **a**

B. [2] Your instructor is fond of certain sayings, which he repeats *ad nauseum*. If you care about your sanity, make him stop by showing that you get it already. Consider the following:

*Architect for **a**, engineer for **d***

Fill in each of the blanks with one of the following words (for your answer, put the letter rather than writing out the word itself). You should only use two words, one for each blank.

- a) Flexibility
- b) Robustness
- c) Cost
- d) Performance

Note: *There are many combinations that are true (or at least reasonably close to true), but the question here is which does the instructor repeatedly pound into your head.*

C. [4] Consider the five layers of the Internet.

- a) Application
- b) Transport
- c) (inter)network
- d) Link-Layer
- e) Physical

i) Which is the highest layer implemented in hosts? **a**

ii) Which is the highest layer implemented in routers? **c**

iii) Which is the highest layer implemented in switches? **d**

iv) Application designers want to rely on a reliable byte-stream abstraction; network designers want to build networks that think only about packets (not byte-streams) and best-effort service (not reliable delivery). What layer bridges the gap between what application designers want and what networks provide? **b**

D. [3] Which of the following three choices is a *necessary and sufficient* condition for successful routing (assuming routing is deterministic): **b**

- a) Forward all packets along the shortest path, and never send a packet out along the port it came in on.
- b) Avoid loops (a packet returning to a node it has already visited) and dead-ends (where a packet is dropped before it reaches the destination).
- c) Always forward packets to a node closer to the destination.

E. [3] Which statement below characterizes a simple invariant that (i) can be implemented in the source and (ii) ensures reliable delivery: **b**

- a) Resend packet if and only if the previous transmission was lost or corrupted.
- b) Always resend packet if the previous transmission was not acknowledged.
- c) Resend packet once for every packet drop

3. Miscellaneous Short Questions: [15 points]

A. [3] Consider the following fields in the IP packet header, with the additional answer of none of the above.

- a) 16-bit Identification Field
- b) 8-bit TTL
- c) 8-bit Protocol
- d) 16-bit Checksum
- e) 4-bit header length
- f) none-of-the-above

Checksum is only over packet header, so it does not catch payload corruption. Too busy being proud of devising such a clever trick question, I then got fooled myself when I filled out the answer key, so we marked it incorrectly. If you answered f and were marked wrong, ask for a regrade. If you answered d and were marked right (which is most of you), be thankful for my stupidity.

- i) Which is used for fragmentation? **a**
- ii) Which can be used to indicate the presence of options? **e**
- iii) Which is used to prevent packets from looping forever? **b**
- iv) Which is used to detect corruption of the packet payload? **f**
- v) Which is used to instruct the network stack how to handle incoming packets? **c**
- vi) Which of these fields (name all of them) were completely removed in IPv6? **a, d, e**
 Note: *the IPv6 next-header field is similar in spirit to the protocol field, so don't include that header in your answer.*

Speaking of stupid, some of you listed c in your answer even though we explicitly told you not to. But if you lost credit for this, come ask for a regrade.

B. [2] DHCP: When obtaining an IP address through DHCP, the following messages are sent to/from the requesting host.

- a) DHCP discover
- b) DHCP offer
- c) DHCP request
- d) DHCP ACK

Which of these are broadcast? **a, b, c, d**

C. [4] Consider the following protocols.

- a) DHCP
- b) DNS
- c) Ethernet
- d) FTP
- e) HTTP
- f) IP
- g) TCP
- h) UDP
- i) WTF

You will be asked to describe the set of headers of a packet. This description should be the ordering of the header bits, starting with the bits at the front of the packet. In other words, for a packet that has protocol a's bits before protocol b's bits, write: a b. You can use each protocol as often as you'd like in the answer.

i. Host sends an HTTP request (using Ethernet at link-layer). Please list the resulting headers: **c f g e**

If you missed this question, make sure you understand why.
Please come to office hours for an explanation.

ii. Host sends a DHCP request (using Ethernet at link-layer). Please list the resulting headers: **c f h a**

D. [1] How many bits are there in the host portion of the address 12.34.15.9 with a subnet mask of 255.255.192.0? **14**

E. [2] Consider the following routing protocols:

- a) Distance-vector
- b) Link-state

i) Which involves global state but local computation? **b**

ii) Which involves local state but global computation? **a**

SID: _____

F. [6] Consider a node A behind a NAT, sending packets to a node B. We will follow a simple packet exchange of A sending to B, having the packet altered by NAT, and then the response coming back, again altered by NAT, and then arriving at A.

You know the following:

Packet leaving A:

Destination address: 4.3.2.1
Source address: 192.168.2.7
Destination port: 80
Source port: 1067

Packet leaving B:

Destination address: 2.2.2.2
Source address: 4.3.2.1
Destination port: 2500
Source port: 80

i) For the packet leaving the NAT towards A, please fill in the blanks:

Destination address: **192.168.2.7**
Source address: **4.3.2.1**
Destination port: **1067**
Source port: **80**

ii) Fill in the following blanks:

IP address of the NAT: **2.2.2.2**
IP address of node A: **192.168.2.7**
IP address of node B: **4.3.2.1**

4. Putting It All Together [15 points]

The client on host A is accessing a web page. The following messages are seen traversing the link to A (in either direction). Indicate which order they are in, and whether they are going to, or coming from, A. We ignore ARP (which we did not cover in class yet), we assume that the necessary DNS entry is not in A's cache, and (for clarity) we assume A does not piggyback its HTTP request on a TCP ACK (i.e., that these are two separate packets).

- a) TCP SYN
- b) DNS Request
- c) HTTP Request
- d) TCP ACK
- e) TCP SYNACK
- f) TCP FIN
- g) TCP FINACK
- h) DNS Reply
- i) HTTP Response

For the answers below, enter the letter of the packet and circle either From A or To A. We assume that the client initiates closing the TCP connection, and that all responses are a single packet. You can reuse packets as needed.

Packet 1: **b** **From A**

Packet 2: **h** **To A**

Packet 3: **a** **From A**

Packet 4: **e** **To A**

Packet 5: **d** **From A**

Packet 6: **c** **From A**

Packet 7: **i** **To A**

Packet 8: **f** **From A**

Packet 9: **g** **To A**

Packet 10: : **d** **From A**

5. Learning Switch [14 points]

Consider a learning switch with five ports. Immediately after being turned on, it receives the following series of packets (these are the only packets the switch sees during this period) which are described using the following notation:

port 2: $A \rightarrow B$ denotes packet sent from A to B that arrives on port 2

The packets listed below arrive in sequence, one right after another. For each packet, write down which outgoing ports the packet is forwarded out of. The only acceptable answers are one of the following choices:

- a) The switch forwards the packet out port 1.
- b) The switch forwards the packet out port 2.
- c) The switch forwards the packet out port 3.
- d) The switch forwards the packet out all ports except the port the packet arrived on.
- e) The switch drops the packet without forwarding it.

List the answers by putting the appropriate letter in the blanks below:

- i) port 1: $A \rightarrow B$ **d**
- ii) port 2: $B \rightarrow C$ **d**
- iii) port 2: $B \rightarrow A$ **a**
- iv) port 3: $C \rightarrow B$ **b**
- v) port 3: $C \rightarrow D$ **d**
- vi) port 3: $A \rightarrow D$ **d**
- vii) port 2: $E \rightarrow B$ **e**

6. Routing [16 points]

A set of routing entries is valid (for a particular destination) if a packet is always forwarded to the destination, no matter where it starts. These entries need not route along the shortest path to be valid. Consider the following routing entries for destination A. This table lists, for each node, the next node it would send a packet to if that packet had a destination address A. You do not need to see the topology to answer this problem! *Make sure you don't confuse the first two subproblems (i,ii) with the graph used in (iii).*

i) Are these entries valid? **N** (Y or N). If no, why not? **Loop** (Loop? Deadend? Other?)

Node	Next Hop Taken To Reach A
B	H
C	G
D	C
E	B
F	A
G	D
H	F

ii) Are these entries valid? **Y** (Y or N). If no, why not? ____ (Loop? Deadend? Other?)

Node	Next Hop Taken To Reach A
B	E
C	G
D	A
E	G
F	B
G	A
H	E

SID: _____

iii) Consider the following graph, with links costs listed, and assume we are using shortest-path (or lowest-cost) routing, and that routing has equilibrated to a constant set of routing tables. The routing algorithm uses poisoned reverse, advertising an infinite weight for the poisoned paths.


Consider routing entries with A as the destination.

What distance to A does B advertise to C? ∞ or 6


What distance does C advertise to B? 5

What path does a packet take from B to A? **B-C-D-E-A** or **B-D-E-A**

If you picked the first route (in the third question), then the answer to the first question is ∞ , but if you picked the second route (in the third question) then the answer is 6.



7. Packet Transmission [20 points]



i) Consider two links: link 1 has bandwidth B and latency L , while link 2 has bandwidth $2B$ and latency $L/2$ (that is, link 2 is twice as fast as link 1, but has half the latency). **Here, and elsewhere on this test, latency refers to the propagation delay of a link; it includes no other delays (such as transmission time, processing time, time on the reverse path, etc.).**

A sends a data packet of size P to C at time zero (that is when the first bit hits the wire). Assume there is no processing time at the intermediate node D. When does the last bit arrive at C? Express your answer in terms of P , B , and L .

Last bit arrives at time = $1.5 \times (P/B + L)$

ii) Let RTT be the time it takes for the data to go from A to C plus the time it takes for an ACK to return from C to A. If A is using sliding window flow control, what is the throughput it achieves when:

It has a window size $W = 0.5 \times B \times \text{RTT}$ **Throughput = $B/2$**

It has a window size $W = 2 \times B \times \text{RTT}$ **Throughput = B**

Express your answer in terms of the various variables in the problem.

SID: _____

iii) Consider a simpler situation with a single link with bandwidth B and latency L , with nodes A and C on either end. Starting at time $t=0$, node A sends two data packets back to C , which then immediately ACKs them (requiring no processing delay). The data packets are 60 bytes, the ACK packets are 40 bytes. The first ACK packet arrives at $t=3.8\text{msec}$, and the second arrives at $t=4.28\text{msec}$.



What is the value of B (in Mbps): _____ (Note: **not** bytes/sec or bits/msec)
[Mbps is 10^6 bits/second]

What is the value of L (in msec): _____

Please show how you arrived at these answers.

Let T_1 be the time the first packet arrives (arrives means time the last bit leaves the wire). Let T_2 be the time the second packet arrives. P is the size of data packets, A the size of ACK packets.

$$T_1 = (P/B + L) + (A/B + L)$$

$$T_2 = (P/B) + T_1$$

$$\text{Therefore, } P/B = T_2 - T_1, \text{ so } B = (480 \text{ bits}) / (4.28 - 3.8 \text{ msec}) = 1 \text{ Mbps}$$

$$T_1 = (P+A)/B + 2L$$

$$\text{So } 2L = T_1 - (P+A)/B = 3.8 \text{ msec} - 800 \text{ bits} / 1 \text{ Mbps} = 3 \text{ msec}$$

$$\text{Thus, } L = 1.5 \text{ msec}$$

8. Addressing and Forwarding [15 points]

Consider the following routing table:

56.160.0.0/16 ⇒ port 1
 56.161.0.0/16 ⇒ port 1
 56.162.0.0/16 ⇒ port 5
 56.163.0.0/16 ⇒ port 1
 56.164.0.0/16 ⇒ port 1
 56.165.0.0/16 ⇒ port 6
 56.166.0.0/16 ⇒ port 6
 56.167.0.0/16 ⇒ port 1
 56.168.0.0/16 ⇒ port 2
 56.169.0.0/16 ⇒ port 2
 56.170.0.0/16 ⇒ port 2
 56.171.0.0/16 ⇒ port 2
 56.172.0.0/16 ⇒ port 1
 56.173.0.0/16 ⇒ port 4
 56.174.0.0/16 ⇒ port 3
 56.175.0.0/16 ⇒ port 3

In the above table, the destination address of each incoming packet will match one and only one prefix. Rewrite the routing table using Longest Prefix Match, using the minimal number of prefixes. Note that the table is listed in order, so any aggregation of prefixes will occur between adjacent prefixes. *Hint: start off by expanding these prefixes into bit patterns....*

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. Write the entries below (you will not need more than eight lines, but may need fewer):

Default ⇒ port **1**
56.162.0.0/16 ⇒ port **5**
56.173.0.0/16 ⇒ port **4**
56.165.0.0/16 ⇒ port **6**
56.166.0.0/16 ⇒ port **6**
56.174.0.0/15 ⇒ port **3**
56.168.0.0/14 ⇒ port **2**

Extra Credit (not on curve) [20 points]

EC1. Delays: Consider a stream of packets arriving at a switch destined for a single destination, and assume there is no other traffic at the switch. The switch has a buffer in which to store packets while they are waiting for the transmission line to become free (so the packet is buffered before it is transmitted, but during its transmission we don't count that as time in the buffer). Assume the stream has average arrival rate R (measured in packets per second) and average time in the buffer T .

i) What is the average number N of packets in the buffer? $T \times R$ (expressed in terms of R and T)

ii) We now change the definitions so that T includes the time while in transmission, and N includes the packet being transmitted. What is the average number N of packets in the buffer? $T \times R$ (expressed in terms of R and T)

EC2. RTT Estimation: Assume we have a TCP implementation that uses an exponential weighting factor of $\frac{1}{2}$ (rather than $\frac{7}{8}$) in the update formula for the RTT. Assume that the initial estimate is 20msec, and then the next four samples are: 60msec, 40msec, 20msec, 10msec. What is the RTT estimate after the fourth RTT sample? **20 msec**

$$E_0 = 20\text{msec}$$

$$E_1 = (20+60)/2 = 40\text{msec}$$

$$E_2 = (40+40)/2 = 40\text{msec}$$

$$E_3 = (40+20)/2 = 30\text{msec}$$

$$E_4 = (10+30)/2 = 20\text{msec}$$

SID: _____

EC3. Open-ended design question: Current routing designs (e.g., Link-State, Distance-Vector, Path-Vector) require time to reconverge after a failure. During this interval, loops and deadends can arise which result in packets being dropped. Devise a routing scheme that never has loops or deadends, so it does not drop packets, ever. Feel free to change IP (by including new fields) if that would help, and assume routers can recalculate routes instantaneously.

Wait for the lecture on advanced routing.....

For Fun [No points, just a beanbag]

Who said the following quotes?

- a) Michele Bachmann
 - b) David Clark
 - c) Bill Gates
 - d) Steve Jobs
 - e) Leslie Lamport
 - f) Barbara Liskov
 - g) Ron Paul
 - h) Ru Paul
 - i) Bernie Sanders
 - j) Eric Schmidt
 - k) Scott Shenker
-
- i) I don't know where they're going to get all this money because we're running out of rich people in this country. **a**
 - ii) The Internet is the first thing that humanity has built that humanity doesn't understand, the largest experiment in anarchy that we have ever had. **j**
 - iii) If you don't love yourself, how in the hell you gonna love somebody else? **h**
 - iv) Be a yardstick of quality. Some people aren't used to an environment where excellence is expected. **d**
 - v) Modularity based on abstraction is the way things get done. **f**
 - vi) Be nice to nerds. Chances are you'll end up working for one. **c**
 - vii) We reject kings, presidents, and voting. We believe in rough consensus and running code. **b**
 - viii) A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable. **e**
 - ix) Our country's founders cherished liberty, not democracy. **g**
 - x) There is a special place in hell reserved for the designer of Powerpoint's animation package. **k**
 - xi) I fear very much that for the first time in the modern history of our country the next generation will have a lower standard of living than their parents, and that would be a real tragedy. **i**