Carnegie Mellon
Computer Science Department.
15-441 Spring 2008
Midterm

Name: ______________________________

Andrew ID: ______________________________

INSTRUCTIONS:

There are 22 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.

A	B	C	D	E	F	Total
/ 20	/ 30	/ 10	/ 16	/ 20	/ 4	/ 100

A Routing

1. For the network depicted below, circle the answers that list valid paths that packets may take under valley free routing between a pair of clients.

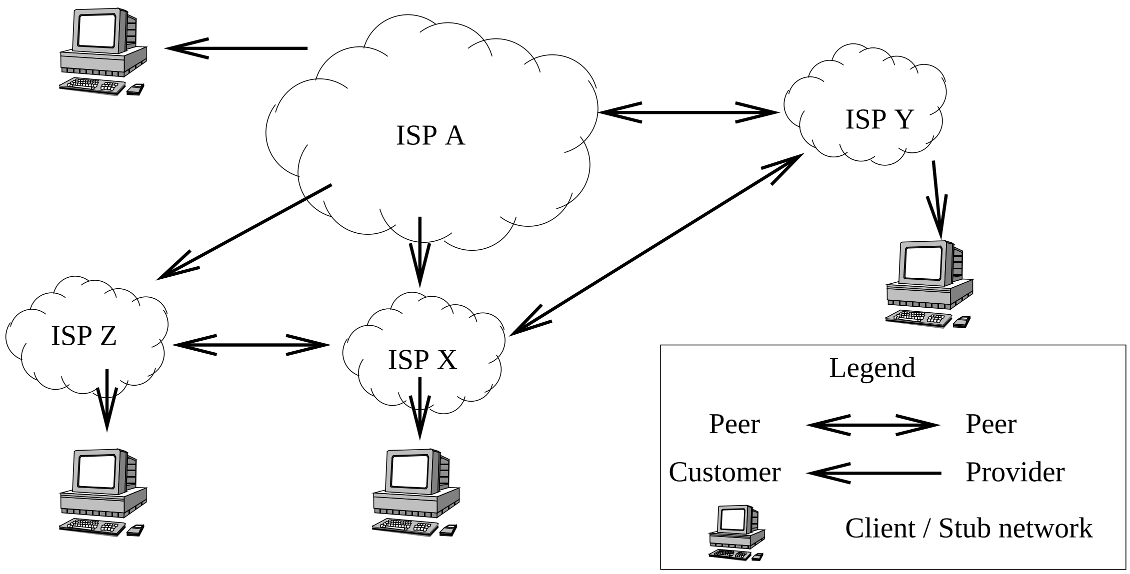

A. Client ⟶ ISP Z ⟶ ISP X ⟶ ISP A ⟶ Client

B. Client ⟶ ISP X ⟶ ISP Y ⟶ ISP A ⟶ ISP Z ⟶ Client

C. Client ⟶ ISP Y ⟶ ISP X ⟶ Client

D. Client ⟶ ISP Z ⟶ ISP A ⟶ ISP Y ⟶ Client

E. Client ⟶ ISP Z ⟶ ISP X ⟶ ISP Y ⟶ Client

Solution: Total: **max 5 pts**

(C)(**3 pts**) and (D)(**3 pts**) are correct paths.

3 pts deducted for each incorrect answer

2. 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 the protocol(s) for which the claim applies:

LS, DV, PV - Requires a map of the complete topology

LS, DV, PV - Sends its routing table to its neighbors

LS, DV, PV - Requires flooding

LS, DV, PV - Suffers the count to infinity problem

LS, DV, PV - BGP is this type of routing protocol

Solution: (a) LS (b) DV,PV (c) LS (d) DV (e) PV

B Multicast

3. The major use of multicast today is not on the wider Internet, but in the datacenter. It is used to distribute information or files to only those nodes interested in the information. One application in particular is used in a testbed called Emulab, which allows users to grab from 1 to 200 machines and run programs on them. Emulab's "frisbee" program distributes entire disk images (hundreds of MB each) to each testbed machine when a user starts an experiment, so that the machines start from a fresh, known operating system installation.

   (a) Frisbee runs *very* quickly, transmitting at over 100Mbit/ sec. It could be quite bad if this multicast traffic leaked out all over the Internet. What mechanism could its designers use to prevent this high-rate multicast traffic from leaking out of their network onto the Internet, even if someone on the Internet intentionally subscribed to it?

   > **Solution:** The designers should use TTL-scoped multicast. (We also accepted "use firewall rules to block the multicast address at the border routers", but this solution is not as robust or convenient.)

   (b) You recall from class that an alternative to native multicast is "application-layer" multicast, in which hosts transmit to each other. One logical way to set up application-layer multicast for Frisbee would be to use a distribution tree: the source sends to two nodes, 1 and 2. Each of these nodes sends to two more nodes, and so on, and so on.

   Assume that each node can send and receive at the same rate of R packets per second and can forward a packet as soon as it receives it. Assume that the switch is not a bottleneck (it can forward as many packets as the hosts can send).

   Assume that there are N=128 nodes.

   (a) How much more time, as a percentage (100%? 100000%?), will it take to send the disk images to the nodes using application-layer multicast than native multicast?

   > **Solution:** Given the assumption we stated above (a node forwards a packet as soon as it receives it), the answer is: It will take 200% of the time it would ordinarily take.
   > The sender now sends to two nodes. Therefore, it can only send at 50% of its normal sending rate to each of those nodes. Those nodes transmit two packets out for each packet they receive–but they only receive at 50% of their maximum traffic rate.
   > There is also a small additional delay from the time needed for the last packet to traverse the 7 layers of the forwarding hierarchy, but this delay (milliseconds) is tiny compared to the transmission time.
   > **However**: We were fairly flexible on grading this problem and also accepted the answer 700%, which would be correct iff a node could only send traffic after it received it. Well, kind of—technically, since it would have to send two copies of the disk image, it would really be 1400%, so we took that too.

   (b) If the disk image is 100,000 packets in size, how many total packet *transmissions* will there be using native multicast? Using application-layer multicast? Your answer doesn't have to be completely exact, but should be within a few percent.

   > **Solution:** Using native multicast, the source will send each packet exactly once. There will be 100,000 packet transmissions.
   > Using application-layer multicast, each packet is sent using unicast. Therefore, each packet must be transmitted in order to be received. There are 128 nodes that must receive the packet, and therefore, there will be $128 * 100,000$ total packet transmissions.

(c) Again, with a disk image of 100,000 packets, how many total packet *receptions* will there be using native? using application-layer?

> **Solution:** Using both schemes, every node must receive every packet. There will therefore be $128 * 100,000$ total packet receptions regardless of the scheme.
> (This illustrates the difference between native and application-layer multicast: In native, there is one transmission but there are N receptions.)

(d) Name one advantage of native multicast and one advantage of application-layer multicast.

> **Solution:** Native multicast requires fewer transmissions and is faster.
> Application-layer multicast works in more heterogenous environments, because it doesn't depend on any particular support in switches and routers. It can also take advantage of caching and retransmissions at the application layer, things that are fairly complex to implement with native multicast.

(Space for answers)

C True or False: Tunneling

4. Which of the following applications use tunneling? (Circle EVERY item that does.)

   A. Converting between IPv4 and IPv6

   B. Virtual Private Netorks

   C. Network Address Translation

   D. Mobile IP

   E. IP over ATM (VCs)

Solution: Following applications rely on tunneling: Private Virtual Networks and Mobile IP.

Converting between IPv4 and IPv6 sometimes relies on tunneling (e.g. to connect isolated IPv6 capable routers) but not always. (you get credit either way!)

D IP Layer

5. Sally Student runs a large network at her house and wants to subnet it to separate her work computers from the network that controls her toaster and fish tanks. She has purchased a NAT box, and divided her network up as follows:

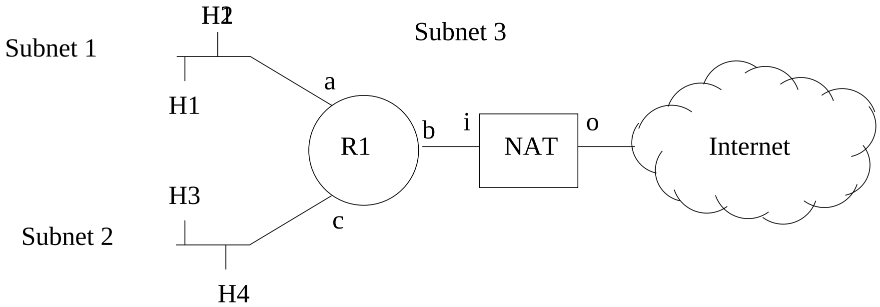

Her ISP has given her an IP address that she assigns to NAT-o (the outside or "o" interface on the NAT box). Sally knows that RFC1918 specifies three different address ranges that she could use for private addresses inside her home:

10.0.0.0	-	10.255.255.255	(10/8 prefix)
172.16.0.0	-	172.31.255.255	(172.16/12 prefix)
192.168.0.0	-	192.168.255.255	(192.168/16 prefix)

(a) Your job is to assign addresses to the subnets, routers, and NAT box inside her house. Use addresses from the 10.x netblock.

Subnet	**Number**	**Netmask**
Subnet 1		
Subnet 2		
Subnet 3		

Interface	**IP Address**
H1	
H2	
H3	
H4	
R1a	
R1b	
R1c	
NAT-i	

Solution: There are many solutions, here is one:

Subnet	**Number**	**Netmask**
Subnet 1	10.0.1.0	255.255.255.0
Subnet 2	10.0.2.0	255.255.255.0
Subnet 3	10.0.3.0	255.255.255.0

Interface	**IP Address**
H1	10.0.1.1
H2	10.0.1.2
H3	10.0.2.1
H4	10.0.2.2
R1a	10.0.1.3
R1b	10.0.3.1
R1c	10.0.2.3
NAT-i	10.0.3.2

Note that broadcast addresses and network addresses can NOT be assigned to interfaces.

(b) Give one reason that wide-spread deployment of IPv6 would let Sally get rid of her NAT device.

Solution: IPv6 would allow Sally to place a virtually unlimited number of devices on her home network without special IP address allocation from her provider.

(c) Give one reason that Sally might want to continue using her NAPT even if she could use IPv6.

Solution:
A NAPT provides a measure of security by not permitting unsolicited inbound connections.

(d) Assuming that the NAT box has no special support for any protocols, and merely translates TCP and IP ports and addresses, give an example of an application that would not work through this NAT, and very briefly explain why.

Solution: FTP would not work through this NAPT because it requires that the server open a connection back to the client. (Passive mode FTP would work—it has the client open the connection instead). Any protocol that embeds IP or TCP-layer information in the application stream is likely to be broken by a basic NAT.

(Space for answers)

6. The figure below shows how two private, physically disconnected IP networks are connected by an IP tunnel. Neither network advertises routes so they are not reachable from the general Internet. For each router interface, we show the IP address plus the datalink layer address (e.g., an IEEE 48-bit address for Ethernet or virtual circuit identifier for ATM).

   In the image below, let H1 be the host with MAC address 1 and H2 be the host with MAC address 12.

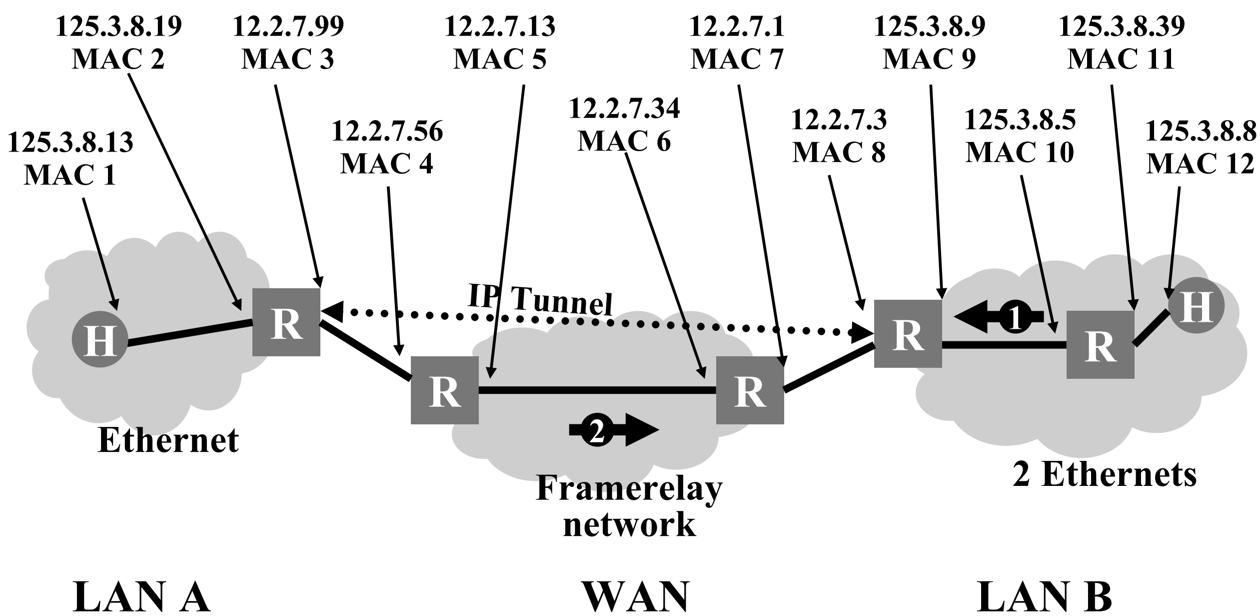

A packet sent from H1 to R1 (the left most router) has the following contents:

Hdr Type	Src	Dest
Ethernet	DL 1	DL 2
IP	125.3.8.13	125.3.8.19
Data		

(a) Fill the table below with the header types and source/destination addresses for each layer 2 and layer 3 header in packet 1 marked on the figure. If a header has no source and/or destination address, please briefly explain what other information in the header is used to forward the packet. The packet is in transit from H2 to H1. (Note that the table has more spaces for headers than you will need.)

Hdr Type	Src	Dest
Data		

Solution: There is an IP header (with the addresses of the source and destination) and a M AC header with the MAC address of the source and destination on this network at location 1

Hdr	Src	Dest
MAC	10	9
IP	125.3.8.8	125.3.8.13

...data...

(b) Same question for packet 2, which is in transit from H1 to H2.

Hdr Type	Src	Dest
Data		

Solution: Again we have an IP header with the addresses of the source and destination. Th e next header is the tunneling header with the IP addresses of the tunnel entry and exit. Sin ce the packet is traveling over a framerelay network, the “destination” address i s a virtual connection identifier. There is no source address.

Hdr	Src	Dest
		VC 1
IP	12.2.7.99	12.2.7.3
IP	125.3.8.13	125.3.8.8
...data...		

E Physical Layer

7. You are building a network with distance of about 1km. Please rank the following technologies in terms of capacity (use a 1 to indicate the highest capacity), assuming we use commonly used transmit/receive technologies:

   ______ Single mode fiber

   ______ Cat 5 twisted pair

   ______ Multi-mode fiber

   ______ Coax cable

Solution:

1. Single mode fiber
2. Multi-mode fiber
3. Coax cable
4. Cat 5 twisted pair

8. Consider an ideal, noiseless channel with 1000 Hz bandwidth. What does Shannon's law predict the capacity of this link will be? (Capacity in bits per second).

> **Solution:** Shannon's law states: C = B x log (1 + S/N)
>
> If the channel is noiseless (i.e. N = 0), the capacity is infinite. Note that this is not realistic.

F Link Layer

9. Calculate the latency (from first bit sent to last bit received) for the following:

   (a) A 10-Mbps link with a single store-and-forward switch in the path, and a packet size of 1024 bits. Assume that each section of the link introduces a propogation delay of 10 microseconds, and that the switch begins retransmitting immediately after it has finished receiving the packet.

   > **Solution:** To put the packet on the first link it will take:
   > $\frac{1024bits}{10-Mbps*1{,}000{,}000bits/Mb} = .1024ms$
   > Then it will take another .01ms for the switch to finish receiving the packet for a total of .1124 ms. It will take another .1124 ms for the switch to send, and then for the destination to receive all of the packet. The total is:
   > 2 * .1124 ms = .2248 ms

   (b) Same as (a) but with five switches

   > **Solution:** From (a) we know that each link section will take .1124 ms. The total time will be 6 * .1124 ms = .6744 ms

   (c) Same as (a) but assume the switch implements cut-through switching: it is able to begin retransmitting the packet after the first 128 bits have been received.

   > **Solution:** To begin we can calculate the time it will take to send the first 128 bits:
   > $\frac{128}{10-Mbps*1{,}000{,}000bits/Mb} = .0128ms$
   > Then, .01 ms later, the switch can start sending the packet. It will take .1024 ms (from (a)) to put the entire packet on the wire. Then it will take .01 ms for the destination to finish receiving the entire packet:
   > $.0128ms + .01ms + .1024ms + .01ms = .1352ms$

10. (a) Given the Ethernet LAN shown in the figure below, indicate which ports are not selected (and therefore disabled) by the spanning tree algorithm. The hosts in the figure are labeled A-J and the bridges in the figure are labeled B1-B7. Hub Bi has an ID of i which is used as the tie breaker.

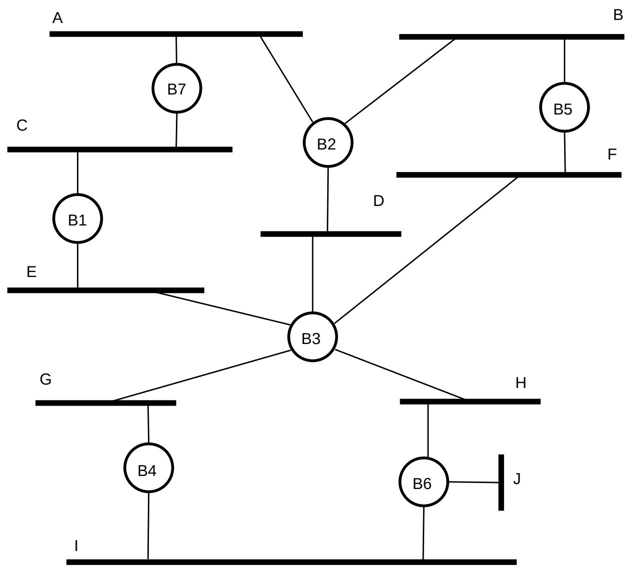

Solution: The following ports are not selected: B2-A B5-B B5-F B6-I

In the above picture mark the paths that would be used for sending data over the following routes (make sure to label the paths):

(a) Host A to Host H
(b) Host I to Host J
(c) Host D to Host F

Solution: Each solution lists the bridges that the data would go through

(a) B7, B1, B3
(b) B4, B3, B6
(c) B3

(d) Consider using the same network topology, but replace the bridges with routers using OSPF. Assume all link costs are equal and use the router's ID as a tie breaker. The hosts are nodes or subnets that are connected to multiple routers. On the picture below, mark (and label) the paths that would be taken to transmit data between the following hosts.

(a) Host A to Host H
(b) Host I to Host J
(c) Host D to Host F

Solution: Each solution lists the bridges that the data would go through

(a) B2, B3
(b) B6
(c) B3

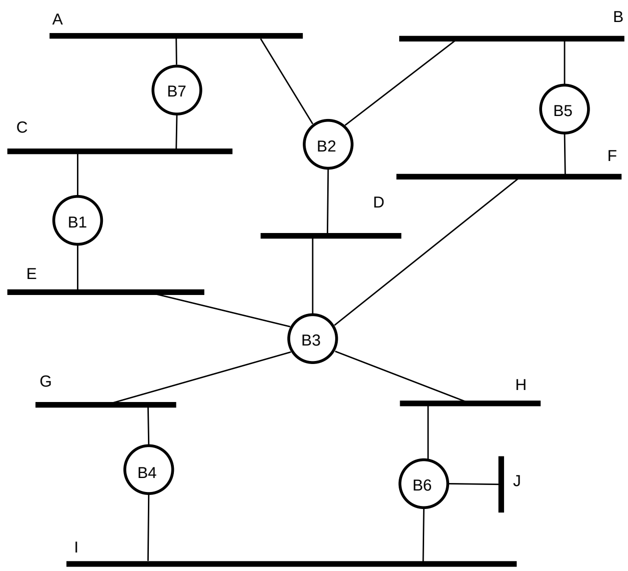

(d) All three transfers are occuring at the same time and each host is sending as much data as it can. The bandwidth and latency on all the links are the same. If two transfers use the same link, they will share the bandwidth equally. A bridge can carry any amount of traffic, only individual links are limited in their capacity. How much capacity will each of the 3 transfers get, in terms of the fraction of the capacity C of a single link:

When the network is using OSPF to create its topology:

Host A to host H:

Host I to host J:

Host D to host F:

Solution: Host A to host H: $C/2$ Host I to host J: C Host D to host F: $C/2$

When the network is using Ethernet and spanning tree to create its topology:

Host A to host H:

Host I to host J:

Host D to host F:

Solution: Host A to host H: $C/2$ Host I to host J: $C/2$ Host D to host F: C

11. Suppose Ethernet was the only existing LAN technology, so every host in the Internet was part of a local Ethernet and therefore had a globally unique Ethernet address. Would you now recommend to get rid of IP addresses by simply using Ethernet addresses instead of IP addresses? What about the other way around, why do we not simply assign network adapters IP addresses instead of Ethernet addresses so we don't have to deal with both, IP and Ethernet addresses?

Solution: We still want to keep IP addresses because they are assigned in a hierarchical fashion which allows hierarchical routing. Ethernet addresses don't have a structure that can be used for hierarchical routing.

We also don't want to replace Ethernet address by IP addresses, since that would mean hardcoding the IP address in a machine. Therefore, we couldn't do things like dynamically reassigning IP addresses (DHCP) and it would also create problems if a host moves to a new location. Also, Ethernet addresses are longer than IP addresses, so IP addresses are more likely to be used up.

(Space for answers)

The End – Phew!

G 4 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: