Carnegie Mellon
Computer Science Department.

15-441 Spring 2008
Final

Name: _____

Andrew ID: _____

INSTRUCTIONS:

There are 16 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.

Generally speaking, you need to briefly explain your answer and you will not get full credit for yes/no answers or just numbers. Exceptions are multiple choice questions or questions where you fill in a table or header.

A	B	C	D	E	F	G	H	I	Total
20	15	15	8	15	15	10	15	7	120

A Protecting Your Web Server

Your friend Alyssa Carmellon recently started a fantastically successful new anti-social networking site where people who dislike people can meet other people, and then dislike them. Unfortunately, with popularity—and the particular clientele the site attracts—has come a serious problem: her site is under constant denial-of-service attack. She's offered you a generous consulting fee if you can help her solve this problem.

1. Alyssa mentions that the attacks seem to be sending 500 small packets every second. The attack causes her web server machine to print out an error message:

   ```
   Kernel Error:  Maximum number of pending TCP connections exceeded.
   ```

   and the machine then ignores most attempts to connect to port 80 (the Web server). She notes that during this attack, the web server software itself appears totally bored—it's not handling *any* GET requests from clients!

   You tell Alyssa that this is most likely a: (circle exactly one)

   A. TCP Bandwidth exhaustion flood
   B. TCP SYN flood
   C. UDP bandwidth exhaustion flood
   D. Smurf attack
   E. Blormig exploit
   F. It's a fluke, ignore it, it'll go away

2. Alyssa seems pleased with your explanation, but still wants to know how to solve it. She tells you that all of the packets seem to come from a single IP address, *IP_evil*. You suggest that she configure her border router to act as a firewall to drop the right packets to protect the Web server from attack. List in the table below the rules you would enter into the firewall. You can enter a single number or IP address, a range, or * as a wildcard. The action should be either "pass" or "drop".

   (Note: Specify the *entire* firewall configuration below. Don't assume anything about the way it was configured or working before or by default.)

Source	Dest	Protocol	Src Port Range	Dst Port Range	Action

3. You receive a check from Alyssa, but two weeks later she returns and tells you that the attacker got smarter. Now, whenever she installs a firewall rule to block the traffic, the source IP address changes to a new one, sometimes within minutes! Her sysadmins can't react quickly enough to keep the site online more than half the time.

   After thinking for a while, you ask her more about the traffic pattern. You observe that no single source should be sending very many of these packets at all, since a source should only need one or two concurrent connections with your Web server. The problem, then, is to ensure that over some period of time, that source doesn't send too many packets.

   (a) What mechanism do you suggest that Alyssa use to solve this new problem?

   A. BGP
   B. A firewall
   C. Token bucket
   D. CSMA
   E. A well-trained monkey

(b) You recall that this mechanism can be configured with both a long-term rate and a small "burst" number of packets. Assume that you will need to support a wide range of legitimate clients, from ancient browsers using the earliest HTTP 0.9 (recall too that many of these old browsers also opened 4 connections at a time) to modern browsers using HTTP 1.1 with its modern features for improved efficiency. Suggest to Alyssa:

How many connections to allow in a burst, and briefly (1 sentence) explain why:

How many connections to allow as a longer-term rate (per second or per minute):

(Your answers don't have to be precise, we're just looking for a reasonable answer and justification that makes sense.)

(c) With the parameters you chose, how many connections could a client open in:

1 second?

10 seconds?

1 minute?

B Waterpipe network

4. 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. 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?

   (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) This prisoners started planning a jail break. During this time, they were all talking to each other frequently. Unfortunately, they found that using CSMA/CD over the Water Pipe Network resulted in a significant packet loss rate. Can you identify the problem responsible for the packet losses and propose a solution?

C Layering/Tunnelling

5. The figure below shows a network topology connecting two LANs. LAN 1 is uses a NAT to connect to the Internet and includes a client host H1. LAN 2 includes a web server H2. Packets between the two LANs are routed through the tunnel between two routers as shown in the figure. The various interfaces of the routers and hosts have the IP and MAC addresses shown in the figure.

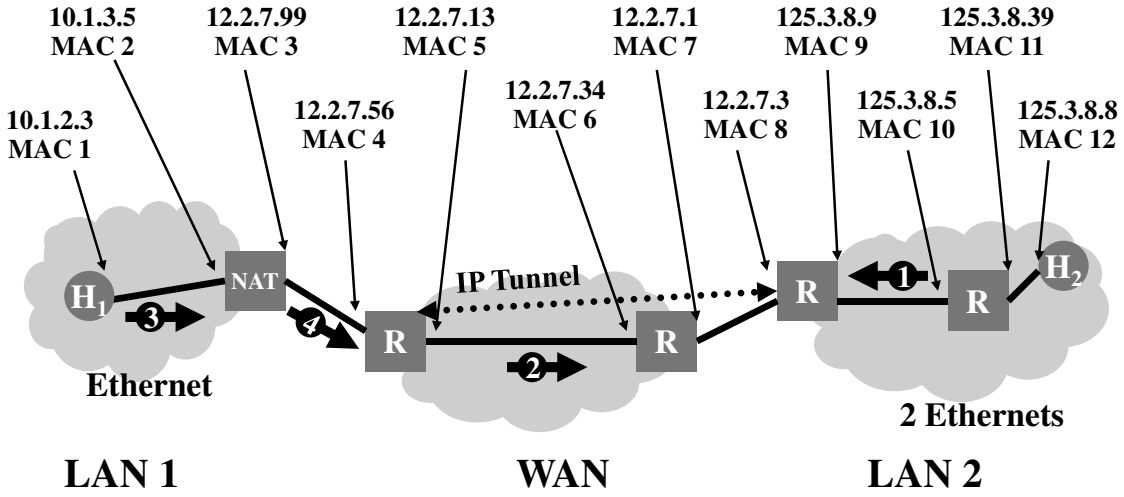

Host H1 has established an HTTP session with web server H2 and data packets are flowing between the two machines. You have to fill in the header type and the source and destination addresss for the network and datalink layer headers for packets 2, 3, and 4 that are shown in the figure. We have filled in the headers for packet 1 (traveling from H2 to H1) as an example. Note: you may not need all the headers (rows) in the figures.

Hdr Type	Src	Dest
Ethernet	MAC 10	MAC 9
IP	125.3.8.8	12.2.7.99
Data		

Header for Packet 2:

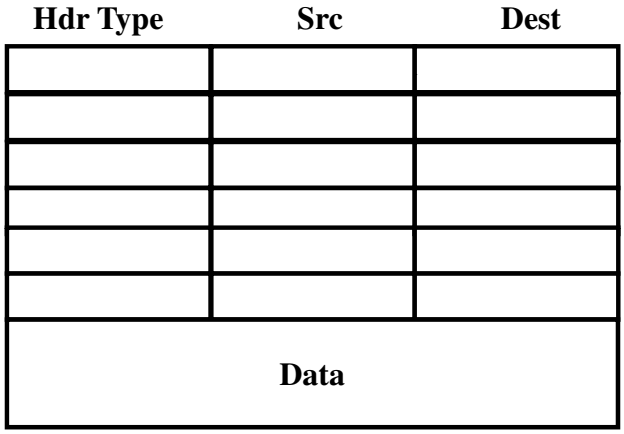

Header for Packet 3:

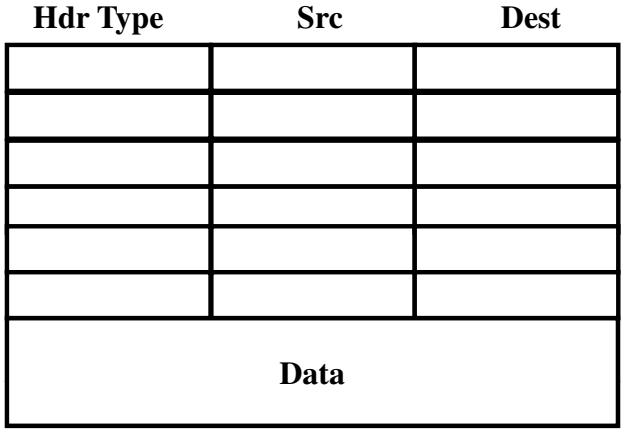

Header for Packet 4:

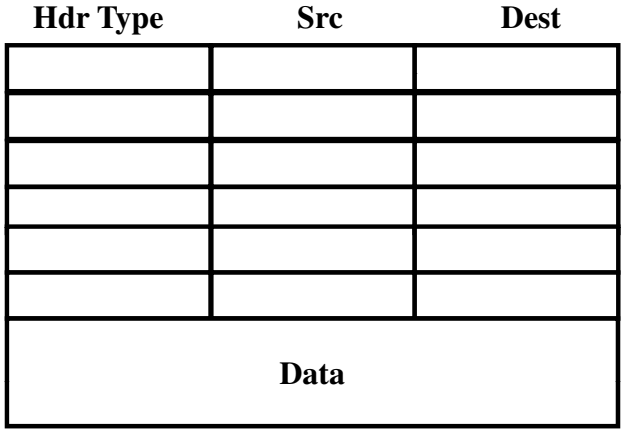

D Multiple Choice

6. Which of the following statements about wireless networks is/are true? (circle ALL that apply)
   (a) All wireless networks must use access points.
   (b) The sender can detect a collision without feedback from receiver.
   (c) Collisions can still happen when RTS/CTS mechanism is used.
   (d) Packet sniffing is easier in wireless networks than in wired networks.
   (e) Wireless networks generally have higher loss rates than that in wired networks.

7. A network C advertises the network number 192.3.124/22 (and no other numbers). What network numbers (all 24 bits) could AS C own? (circle ALL that apply)
   A. 192.3.123
   B. 192.3.124
   C. 192.3.125
   D. 192.3.128
   E. 192.3.10
   F. 192.4.124

8. Consider an ideal, noiseless channel with 1500 Hz bandwidth. What does Shannon's law predict the capacity of this link will be?
   (a) Zero
   (b) 3000 bps
   (c) 5 dBm
   (d) infinite
   (e) 8934 bps

E Tag switching

9. The CMU Credit Union is thinking about creating a world-wide network of Cash and Homework Dispensers for CMU students. They've decided to use a network of ATM PVCs (permanent virtual circuits) to connect the dispensers back to their master bank computers, and want you to review their network configuration to make sure that it's set up properly.[1]

   The ATM network is shown in the box in the figure below.

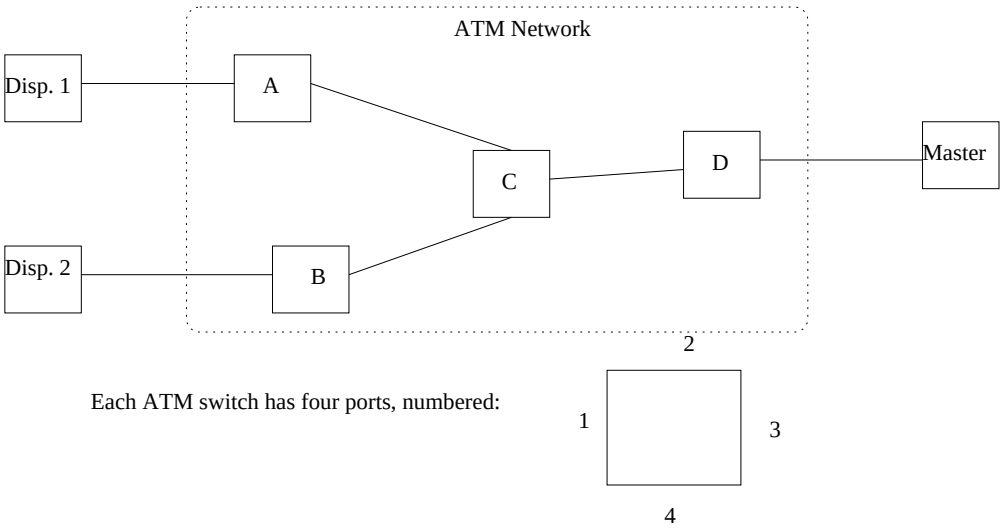

   (a) Fill in the missing parts of the configuration table, and use the smallest possible VCI number. To save some writing, we'll only worry about traffic flowing TO the master - we'll assume that we configure the outbound links using the same VCIs and swap the input/output parts.

Switch	Input Port	Input VCI	Output Port	Output VCI
A	1	1		
B	1	1		
C	2			
C	4			
D	1			
D	1			

   (b) Write the sequence of (Switch, Input Port, Input Label) tuples and final label for the two virtual circuits, i.e. Dispenser 1 to master and Dispenser 2 to master. We've given you the start node and starting label. The intermediate tuples should look like (A, 1, 999) [e.g., switch A, input port 1, label 999].

[1] Yes. We're using ATM to connect ATMs. Sorry.

Start node A, label 1. Switch tuples:

Final label:

Start node B, label 1. Switch tuples:

Final label:

(c) The bank comes back to you later and tells you that they're encountering problems when uploading new software images to their dispenser machines. Things are great for handling small customer transactions ("withdraw $50 from account 555"), but these large file transfers fail miserably. You realize that the ATM network provider has never needed to support data before, since their primary business has been voice up until now. Can you explain the problem and how to fix it?

F Routing and Bridging and Bears, oh my!

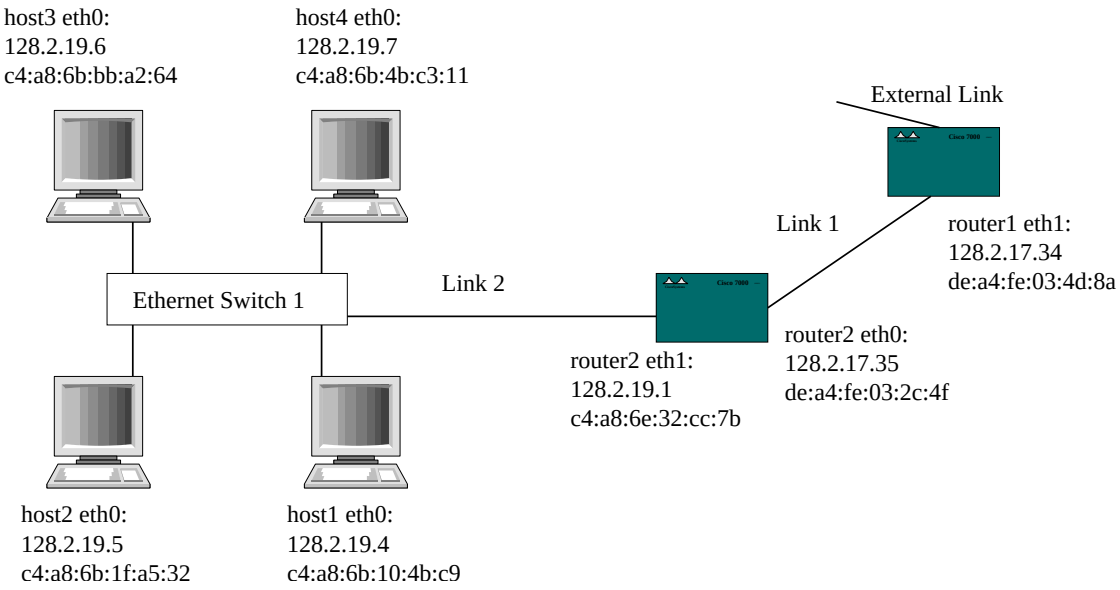

In the partial network topology shown above, a well-formed IP packet with a destination IP of 128.2.19.5 and TTL of 8 arrives at router1 via the external link. Link1 uses the subnet 128.2.17.34/31 and Link2 has the subnet 128.2.19.0/25 .

The hosts and routers on the ethernet subnet are all connected to their own port on an ethernet switch (a learning bridge).

10. What is the subnet mask of eth0 on host1?

11. What are the forwarding entries used by each router to forward the packet to 128.2.19.5?

Router	Destination	Mask	Next-Hop	Interface
router1				
router2	128.2.19.0	255.255.255.128 (/25)	directly connected	eth1

12. If this packet (from the outside to 128.2.19.5) is the first packet to be forwarded on the network, ARP requests will be sent out on both Link1 and Link2. Fill in the following ARP header fields for the requests.

Link	MAC Source Address	MAC Destination Address	IP address queried for
Link1			
Link2			

13. Circle all of the header fields in the following list that router2 will change between when it receives the packet on eth0 and when it transmits the packet on eth1.

Ethernet Header:

A. Source address
B. Destination address
C. EtherType
D. None

IP Header:

A. Source address
B. Destination address
C. Protocol
D. IP TTL
E. IP Header Length
F. IP Header Checksum
G. IP ID
H. None

14. Circle all of the header fields in the following list that will change from when the switch receives the packet on the port connected to router2 and when it transmits the packet on an outgoing port.

Ethernet Header:

A. Source address
B. Destination address
C. EtherType
D. None

IP Header:

A. Source address
B. Destination address
C. Protocol
D. IP TTL
E. IP Header Length
F. IP Header Checksum
G. IP ID

H. None

15. Assume that no packets have been transmitted on the network until the first packet to host2 arrives. In this question, assume that a host "sees" a packet if it is transmitted on its local medium by any device other than itself, but does not "see" packets that it generates.

    Select the correct answer below to indicate how many packets host2 sees and how many packets all other hosts on the network see. Include both ARP packets and data packets.

    A. host2 = 1 All others = 0

    B. host2 = 1 All others = 1

    C. host2 = 2 All others = 0

    D. host2 = 2 All others = 1

    E. host2 = 2 All others = 2

G DNS

For answering the following question, consider the following name servers and the entries they contain:

Local name server

Name	Value	Type	Class
edu	a3.nstld.com	NS	IN
a3.nstld.com	192.5.6.32	NS	IN

a3 name server

Name	Value	Type	Class
princeton.edu	dns.princeton.edu	NS	IN
dns.princeton.edu	128.112.129.15	A	IN

princeton name server

Name	Value	Type	Class
mail.cs.princeton.edu	128.112.136.72	A	IN
platypus.cs.princeton.edu	128. 112.155.166	A	IN

(a) Suppose client1 queries the Local name server for mail.cl.princeton.edu. The Local name server will recursively resolve the request and cache the results with a TTL of 10 minutes. Right down the requests and responses the nameserver will get and the response it eventually sends back to client1.

(b) 42 seconds later client2 queries the Local name server for platypus.cs.princeton.edu. List out the requests and the responses from the point of view of the Local nameserver and the response it eventually sends back to client2.

H BGP - Does it really work?

16. An Autonomous System (AS) claims that it "owns" an IP CIDR block such as 128.2.0.0/16 by advertising a path to 128.2.0.0/16 that has only the single number in its AS Path. For example, CMU's routers claim to own the block 128.2.0.0/16 by advertising a route to this block with AS Path 9. There is nothing, however, to prevent two different Autonomous Systems from both claiming to own the same IP address, although one of the routers must be mistaken.

    (a) Explain why even if this happens, the BGP protocol prevents cycles from forming in BGP routing tables.

    (b) Suppose that a router in AS 9 (Carnegie Mellon) mistakenly claims that it owns CIDR block 18.0.0.0/8 (M.I.T.) by advertising a route to 18.0.0.0/8 with AS Path 9. What are the consequences of this? Recall that one of the principles of BGP is that a router should only advertise the path that it considers the best path to a destination.

17. In the following topology, network S is multi-homed to four providers, i.e., P_1, P_2, P_3 and P_4, where P_1, P_2 and $P3$ are customers of a common provider P while P_4 is a customer of $P^{'}$. Now suppose D is the destination of interest.

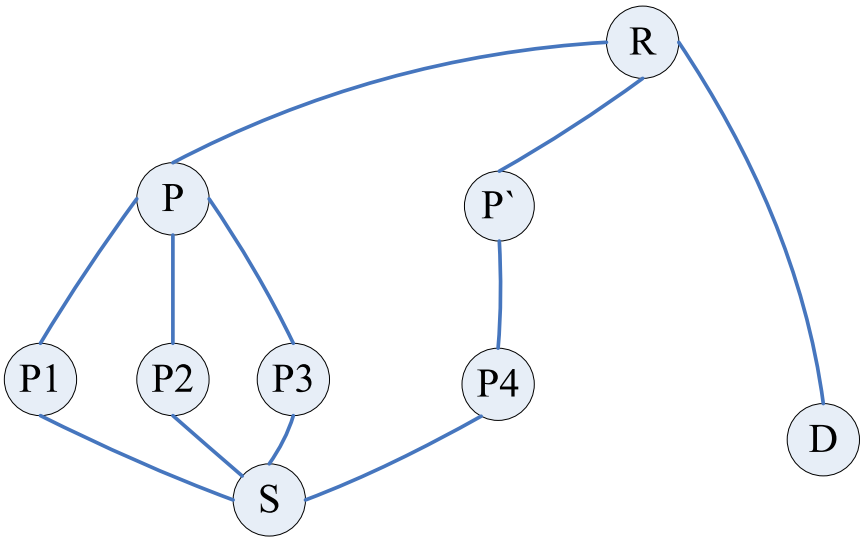

(a) What are the possible routing paths to D that S can have in its BGP routing table?

(b) Suppose S maintains all possible routing paths in its BGP routing table, and S prefers the paths going through the provider P (say, due to lower cost). Now suppose P is broken, what can P_1, P_2 and P_3 announce (withdraw) to S via BGP updates?

(c) Suppose only the announcement from P_1 arrives at S, while the announcements from P_2 and P_3 are lost. How can S find a working path to D? (show the path exploration process of S)

(d) If you think the above path exploration process is undesirable, how can you suggest to improve it? (you can modify the BGP routing protocol in any reasonable way)

I TCP

18. At time t, a TCP connection has a congestion window of 4000 bytes. The maximum segment size used by the connection is 1000 bytes. What is the congestion window after it sends out 4 packets and receives acks for all of them...

    (a) If the connection is in slow-start?

    (b) If the connection is in congestion avoidance (linear mode)?

The End – Phew!