

CS485/ECE440: Final Exam

December 11, 2008

Instructions

- This is a take-home exam. The answers are **due December 16, 2008 at 9:30 a.m.**
- That deadline is **firm!**
- Submit via e-mail to `riesen@cs.unm.edu` (Mail it before 9:30 a.m. on the 16th.) Your subject line must say: “Final Submission” (and **nothing else**).
- Make sure to show how you arrived at an answer. Your steps towards a solution are more important than the final result.
- Your answers to the questions below must be in a PDF or plain text file attached as file named `Final.pdf` or `Final.txt` according to its format. (Not `Final.pdf`, and no Microsoft Word or image files, such as jpeg, please.)
- Keep your answers (and programs) succinct.
- Obey the University rules on plagiarism. In particular, do use libraries and the web to find information you need to answer the questions, but do not copy whole answers or programs. Reference your sources. The work you turn in must be *your* work.

Problem 1: Encoding and Compression

Have a look at the table below. It contains binary encodings for most of the lowercase letters, the period, the carriage return `<cr>` (newline), and space (000). As you can see, it is a variable-length encoding; i.e. not every symbol is encoded with the same number of bits. In contrast, ASCII is a fixed-size 8-bit encoding.

Information beyond our textbook may help you answer the the last parts of this question. Search for *prefix codes*.

Sym	Code	Sym	Code	Sym	Code	Sym	Code
a	1011	h	11111110	p	010111	y	111111110
b	111110	i	0100	q	0110110	.	010110
c	01110	j	111111111	r	01010	<cr>	111111110
d	11110	l	1110	s	0011		000
e	100	m	01111	t	0010		
f	1111110	n	1010	u	110		
g	0110111	o	01100	v	011010		

Answer these questions:

- a). (5 points) Why do you think the space symbol (000) received a shorter encoding than the symbol j (111111111)?
- b). (10 points) What does the following bit string encode? (Broken into lines, scanning starts at the top left to right, line by line:


```
1111111010110101110101111111111
1100001001010111100000110011111
10000011101110101100110011010110
```
- c). (10 points) On the class web page you can find a file name `loremipsum.txt`. If you encoded it using the values in the table above, how many bits would you need? Is that less or more than the file size; i.e. its ASCII encoding?
- d). (5 points) Use a compression program of your choice; e.g., (`gzip`), `compress`, `pkzip`, etc. and compress `loremipsum.txt`. Is the result longer or shorter than using the encoding above?
- e). (10 points) Repeat the above experiments for the file `short.txt` on our web page. Is the compressed version shorter or longer than the encoded version? Compare your results for the two files and briefly explain these results.

Problem 2: Network Addressing

Given the network 192.12.88.64/26, answer these questions:

- a). (3 points) How many hosts could be support on this network?
- b). (2 points) What would be the broadcast for this network?

Problem 3: TCP Reset Flag

Figure 5.4 on page 388 of our textbook shows the TCP header format. The 6-bit flag field is explained on page 389. One of these flags indicates a reset of the connection. When a host receives such a reset TCP segment, it closes that connection. The behavior is sometimes used by firewalls to limit access to certain sites. For example, a company that wants to prevent its employees from surfing to certain web sites using company resources can monitor TCP traffic that goes through its firewall. When a connection attempt is made to a “forbidden” host outside the company firewall, the firewall sends a reset segment to the host that attempted the connection, which prevents connection establishment.

You will need to do a little research beyond the textbook to answer this question.

- a). (4 points) What, inside a TCP header tells the firewall that this segment is an attempt at a new connection?
- b). (8 points) Assuming you have control over how TCP is implemented or configured on your company computer, how can you get past the firewall anyway?
- c). (8 points) Such a firewall usually employs additional counter measures than sending reset packets. Mention another method you can use to get past the access control of the firewall.

Problem 4: Routing Protocols

Which nodes are involved when link-state and distance-vector protocols exchange information? (I.e., which nodes talk to each other?)

- a). (3 points) Link-state:
- b). (2 points) Distance-vector:

Problem 5: Packet Format

On an Ethernet, the following bytes (in hexadecimal) of a packet were intercepted: 00 60 CF 20 2B 8B 00 60 CF 20 2B 7F 08 00 45 00 00 30 C4 B2 40 00 80 06 2A 08 83 F7 02 10 83 F7 02 0F 0B 0B 00 50 50 20 0D 16 00 00 00 00 70 02 40 00 CE 80 00 00 02 04 05 B4 01 01 04 02. Answer the following questions. You will need some information about the format of TCP/IP header options that is not in our textbook.

- a). (3 points) What is the source hardware address of this packet?
- b). (3 points) What is the IP destination address?
- c). (3 points) What is the IP source address?
- d). (3 points) What is the TCP source port of this packet?
- e). (3 points) What is the TCP destination port of this packet?
- f). (3 points) What is the TCP window size?
- g). (3 points) Which TCP flags are set?
- h). (3 points) What service is the client that sent the packet trying to access?
- i). (3 points) What are the last four bytes for?
- j). (3 points) Does either the TCP or IP header have options? If so, which one, and what does the option mean?