Midterm II

CS 414 Operating Systems, Spring 2007
April 26th, 2007
Prof. Hakim Weatherspoon

Name: _____NetId/Email:_____

Read all of the following information before starting the exam:
Write down your name and NetId/email NOW.

This is a **closed book and notes** examination. You have 90 minutes to answer as many questions as possible. The number in parentheses at the beginning of each question indicates the number of points given to the question; there are 100 points in all. You should read **all** of the questions before starting the exam, as some of the questions are substantially more time consuming.

Write all of your answers directly on this paper. ***Make your answers as concise as possible.*** If a question is unclear, please simply answer the question and state your assumptions clearly. If you believe a question is open to interpretation, then please ask us about it!

Good Luck!!

Problem	Possible	Score
1	24	
2	20	
3	28	
4	16	
5	12	
Total	**100**	

1. (26 points) True/False.
Circle either True or False. In the following, it is important that you EXPLAIN your answer in *TWO SENTENCES OR LESS* (Answers longer than this may get ***no credit***!). Also, answers without an explanation GET ***NO CREDIT***.

a. (3 points) "Marshaling" is the process by which Byzantine Generals are forced into making good decisions.

TRUE / FALSE
EXPLAIN:

b. (3 points) A "broadcast network" is one which uses radio-frequency transmission to send data from one party to another.

TRUE / FALSE
EXPLAIN:

c. (3 points) Randomness is essential to achieving good performance from the Ehternet communication algorithm (CSMA/CD).

TRUE / FALSE
EXPLAIN:

d. (3 points) A Remote Procedure Call (RPC) can be used to call a procedure in another process on the same machine.

TRUE / FALSE
EXPLAIN:

e. (3 points) Using the TCP/IP protocol over an unreliable network, the receiver can receive the same IP packet multiple times.

TRUE / FALSE
EXPLAIN:

f. (3 points) With the NFS distributed file system, it is possible for one client to write a value into a file that is not seen by another client when reading that file.

TRUE / FALSE
EXPLAIN:

g. (3 points) The fastest way to send a large document securely to a third party that you have not interacted with yet is to encrypt it with their public key. Assume you know everyone's public key.

TRUE / FALSE
EXPLAIN:

h. (3 points) Doubling the block size in the UNIX 4.2 BSD file system will exactly double the maximum file size.

TRUE / FALSE
EXPLAIN:

EXTRA CREDIT

(3 points) Consider two processes P and Q that are communicating using mailboxes. From P → Q they use mailbox A, and from Q → P they use mailbox B. Assume both the processes are asynchronous, mailboxes are currently empty, and the communication links between P and Q are unreliable. If P now wants to determine that Q has crashed, what is the sequence of instructions P should execute?
You may only use the following blocking instructions: send, reply, re-send, receive, deleteMailbox, createMailbox, and setAlarm. Note that an alarm will interrupt a blocking instruction.

2. (20 total points) Disks.

a. (12 points) Disk requests come into the disk driver for cylinders: 10, 22, 20, 2, 40, 6, 38, in that order. The disk has 60 total cylinders and the disk head is currently positioned over cylinder 20. A seek takes 6 milliseconds per cylinder moved. What is the sequence of reads and total seek time using each of the following algorithms?

i) (4 points) First-come, first-served:

ii) (4 points) Shortest Seek Time First:

iii) (4 points) LOOK (initialing moving upwards):

b. (7 points) RAID

i) (4 points) Give a brief (2-3 sentences) description of RAID 5

ii) (2 points) How many disk failures can RAID 5 tolerate without losing data?

iii) (2 points) How would you reconstruct a failed disk?

3. (28 points total) File Systems.

a. (8 points) Consider a file system with 2048 byte blocks and 32-bit disk and file block pointers. Each file has 12 direct pointers, a singly-indirect pointer, a doubly-indirect pointer, and a triply-indirect pointer.

i) (4 points) How large of a disk can this file system support?

ii) (4 points) What is the maximum file size?

b. (4 points) Briefly (2-3 sentences) state the differences between a hard link and a soft link.

c. (6 points) Rather than writing updated files to disk immediately when they are closed, many UNIX systems use a delayed *write-behind policy* in which dirty disk blocks are flushed to disk once every 30 seconds. List two advantages and one disadvantage of such a scheme.

Advantage 1:

Advantage 2:

Disadvantage:

d. (6 points) List the set of disk blocks that must be read into memory in order to read the file /home/cs414/test.doc in its entirety from a UNIX BSD 4.2 file system (10 direct pointers, a singly-indirect pointer, a doubly-indirect pointer, and a triply-indirect pointer). Assume the file is 15,234 bytes long and that disk blocks are 1024 bytes long. Assume that the directories in question all fit into a single disk block each. *Note that this is not always true in reality*.

e. (4 points) On a single UNIX machine, if some program B reads a block of a file after it has been updated by another program A, the copy of the file block B reads will include A's updates. In NFS this behavior is not guaranteed. Assuming that there are no failures, why doesn't NFS necessarily provide such update semantics when A and B are run on different machines? What semantics does it provide instead?

f. (4 points) The Andrew File System (AFS) solves the above problem (e) using state information it maintains at the server. What state is kept? How is the state used to solve the problem?

4. (16 points total) Network Performance.

a. (8 points) Consider a TCP network connection with a current window size for unacknowledged bytes of 1,000 bytes, over a cross-country link with a one-way latency of 50 milliseconds, and a link bandwidth of 100 Mbit/second. You may assume that no packets are lost for this particular problem, and that the size for an acknowledgement is essentially 0 bytes long.

How long does it take TCP to transmit 100,000 bytes across the link? That is, how much time elapses from when the first byte is sent by the sender to when the sender *knows* that the receiver has received the last byte?

b. (4 points) Assume that the receiver can process incoming data at greater than 100 Mbit/s, what is the optimal window size that the receiver should advertise?

c. (4 points) If the link is shared by N pairs of senders and receivers, does your answer for part (b) change? If so, how? If not why?

5. (12 points total) Security.

a. (4 points) Assume Alice and Bob have never met. Explain how Alice can use a public key infrastructure to prove her identity to Bob. You can assume that a PKI is essentially a list of everyone's public key signed by a Certificate Authority (CA). *Hint: make sure to prevent replay attacks*.

b. (4 points) Explain how to utilize a PKI to establish a private session key between both parties for fast symmetric encryption.

c. (4 points) What are two desirable properties for secure hash functions (ignoring the property where half the bits change for small changes in input). Why are these properties important for signatures?