University of California, Berkeley
College of Engineering
Computer Science Division – EECS

Fall 2013 | Anthony D. Joseph and John Canny

Midterm Exam #2

December 4, 2013
CS162 Operating Systems

Your Name:	
SID AND 162 Login:	
TA Name:	
Discussion Section Time:	

General Information:
This is a **closed book and one 2-sided handwritten note** examination. You have 80 minutes to answer as many questions as possible. The number in parentheses at the beginning of each question indicates the number of points for that question. 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 there is something in a question that you believe is open to interpretation, then please ask us about it!

Good Luck!!

QUESTION	POINTS ASSIGNED	POINTS OBTAINED
1	8	
2	21	
3	23	
4	32	
5	16	
TOTAL	100	

NAME: ____

1. (8 points total) True/False and Why? **CIRCLE YOUR ANSWER.**

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

TRUE FALSE

Why?

ii) Doubling the block size in the UNIX 4.2 BSD file system will exactly double the maximum file size.

TRUE FALSE

Why?

iii) 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 immediately afterwards.

TRUE FALSE

Why?

iv) In a replicated Key-Value storage system, Iterative PUTs achieve lower throughput than recursive PUTs on a loaded system.

TRUE FALSE

Why?

NAME: ____

2. (21 points total) Security.

   a. (5 points) Are digital signatures on digital documents more or less secure than ink signatures on printed documents? Give a brief (4-5 sentence) justification for your answer.

   b. (8 points total) X.509 Certificates.

      i) (5 points) The recipient of an X.509 certificate can authenticate the sender of the certificate without contacting a third party (e.g., a certificate authority) at the time of authentication. How does this work?

NAME: ____

ii) (3 points) Discuss a vulnerability in the use of X.509 certificates as described in part (i).

c. (8 points) SYN Cookies.

i) (3 points) What type of attack do SYN cookies protect against? Be specific in your answer.

ii) (5 ponts) Explain how SYN cookies are used?

NAME: ____

3. (23 points total) Transactions.
   a. (16 points) Consider the following bank transfer process for moving funds from account A to account B:

Read balance from source account:	**Read Source Account**
Decrement source account balance	
Write new balance of source account:	**Write Source Account**
Read balance from destination account:	**Read Destination Account**
Increment destination account balance	
Write new balance of destination account:	**Write Destination Account**

Suppose John and Alice share the same accounts and they each use different ATMs at the same time. John transfers $100 from Checking to Savings; while Alice transfers $200 from Savings to Checking – *note that one transfer goes from Checking to Savings, while the other is reversed*. Suppose that John and Alice's transfers happen simultaneously, and there is more than $200 in the account.

i) List a **serializable schedule** for the transactions, which is not simply already a **serial schedule**

ii) For each of the following three *potential* situations, state whether such a situation can exists, and if it can, give an example. Otherwise, explain why it cannot exist.

(1) Situation #1: A **non-serializable schedule** for the transactions

NAME: ____

(2) Situation #2: A **serializable schedule** using Two-Phase Locking, which is not simply already a **serial schedule**

(3) Situation #3: A schedule where Two-Phase Locking causes deadlock

NAME: ____

b. (7 points total) Consider the following schedule of four transactions:

T1:		R(B)	R(A)	W(A)				W(D)
T2:					W(B)			
T3:						W(A)		
T4:	W(C)						W(A)	

i) (4 points) Draw the dependency graph for this schedule.

ii) (3 points) Is this schedule serializable? If so, give a possible serial schedule. Otherwise, explain why not.

NAME: ____

4. (32 points) Filesystems.

   a. (8 points) Consider a file system that has 2,048 byte blocks and 32-bit disk block pointers to those blocks. Each file header has 12 direct pointers, one singly-indirect pointer, one doubly-indirect pointer, and one triply-indirect pointer.

      i) (4 points) How large of a **disk** can this filesystem support? *You may leave your answer in symbolic form.*

      ii) (4 points) What is the maximum file size? *You may leave your answer in symbolic form.*

NAME: ____

b. (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 approach.*

Advantage 1:

Advantage 2:

Disadvantage:

c. (4 points) Briefly (in 2-3 sentences) explain the differences between a hard link and a soft link.

NAME: ____

d. (6 points) List the set of disk blocks that must be read into memory in order to read the file `/home/cs162/midterm2.txt` 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*.

NAME: ____

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 (as described in lecture) this behavior is not guaranteed. Assuming that there are no failures, why doesn't NFS necessarily provide such update semantics when programs 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 in part (e) by using state information it maintains at the server. What state is kept? How is the state used to solve the problem?

5. (16 points) Networking.

a. (4 points) Consider a TCP network connection with packet size 1000 bytes, and current receiver advertised window size of 100 packets, over a cross-country link with one-way latency (for a 0-byte packet) of 50 milliseconds in each direction, and a link bandwidth of 100 Mbit/second. You may assume that no packets are lost for this particular problem, and that the times to assemble, unpack and process packets at each end of the connection are negligible.

How long does it take TCP to transmit 1 million 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 with zero latency, what is the optimal window size that the receiver should advertise?

NAME: ____

c. (4 points) Polling a device for the completion of an operation is typically a bad idea because, as with busy-waiting, while polling, the CPU is not doing useful work. However, this is not always true. Describe a situation where polling might be a better choice than using interrupts.

d. (4 points) Explain the difference between an IP address and a MAC address, and why we cannot use only MAC addresses or IP addresses alone.

NAME: ____

This page intentionally left blank as scratch paper

Do not write answers on this page