

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

Spring 2016

Anthony D. Joseph

Midterm Exam #2
April 20, 2016
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	12	
2	26	
3	24	
4	22	
5	16	
TOTAL	100	

NAME: _____

1. (12 points total) True/False and Why? **CIRCLE YOUR ANSWER AND WRITE AN EXPLANATION (no credit will be given if no explanation is provided).**
- a. In addition to greatly reducing the overall failure probability, moving from a single hard drive to using RAID 1 (mirroring) can also increase the throughput of small random reads.

TRUE
Why?

FALSE

- b. The TCP transport protocol provides a reliable, in-order bytestream abstraction.

TRUE
Why?

FALSE

- c. The Byzantine failure model allows arbitrary behavior by a faulty system component.

TRUE
Why?

FALSE

- d. Remote Procedure Calls provide identical semantics to local calls.

TRUE
Why?

FALSE

NAME: _____

2. (26 points total) Memory and I/O.

a. (12 points) For each statement, **CIRCLE** the address translation scheme(s) for which the statement is true.

i) (4 points) A process cannot directly read or write another process' memory. Assume there is no `mmap()` operation.

No translation

Base and Bound

Segmentation

Paging

ii) (4 points) External fragmentation may occur.

No translation

Base and Bound

Segmentation

Paging

iii) (4 points) Part of the process' stack could be swapped out to disk while the process continues to execute using the rest of the stack. *Assume the stack is located in a typical location for the type of address translation in use.*

No translation

Base and Bound

Segmentation

Paging

NAME: _____

b. (10 points) Consider a system with the following specifications:

- 46-bit virtual address space
- Page size of 8 KBytes
- Page table entry size of 4 Bytes
- Every page table is required to fit into a single page

How many levels of page tables would be required to map the entire virtual address space?

In the space below, document the format of a virtual address under this translation scheme. Briefly explain your rationale.

c. (4 points) Briefly explain, in four sentences or less, what is DMA *and* why DMA enables a system to more efficiently handle I/O devices.

NAME: _____

3. (24 points total) Demand Paging.
- a. (18 points) Consider a demand paging system with 3 pages of physical memory. When a page fault occurs, we use a *page replacement algorithm* to make space for the new page by evicting a page from memory. Using the FIFO, Clock Algorithm, and LRU page replacement algorithms, *mark which pages will be in physical memory after each page request*. The first 3 rows for each algorithm are pre-filled for you.

Time	FIFO				Clock				LRU			
	A	B	C	D	A	B	C	D	A	B	C	D
---	X				X				X			
A	X	X			X	X			X	X		
B	X	X	X		X	X	X		X	X	X	
C												
D												
B												
A												
B												
A												
D												
C												
Number of page faults												

NAME: _____

b. (4 points) When a page fault occurs and a page is evicted, what entries in which data structures must the OS invalidate?

c. (2 points) In the N^{th} chance page replacement algorithm, explain the purpose of the parameter N . Consider the effects of a small or large N .

NAME: _____

4. (22 points total) Disks and File Systems.
- a. (6 points) Name two performance-improving features that are provided by modern disk controllers. For each of the features you name, briefly explain the benefit consequences of a disk controller that does not provide that feature.
- i) Feature #1 – Benefits and consequences if it is not provided:
- ii) Feature #2 – Benefits and consequences if it is not provided:
- b. (5 points) If a disk subsystem receives an average of 10 requests per second and each request has an average response time of 500ms, how many requests are in the subsystem on average? *Explain your answer.*
- c. (4 points) The FAT file system eliminates external fragmentation by allocating disk space in block-sized units. Considering that fact, briefly explain why defragmentation is still sometimes necessary.

NAME: _____

- d. (3 points) In the UNIX 4.2 BSD FFS, inodes have direct pointers, a singly indirect pointer, a doubly indirect pointer, and a triply indirect pointer. The maximum file size supported by this inode type is approximately the same as the maximum file size supported by an inode with only a triply indirect pointer. Briefly explain one disadvantage of an inode design that only uses a triply indirect pointer instead of a combination of pointers (*other than the slightly reduced maximum file size*).
- e. (4 points) Briefly, *in one to three sentences*, explain the purpose of the journal in a journaling file system?

NAME: _____

5. (16 points total) Pintos coding question: Verifying user pointers.

Your project group is tasked with implementing a new syscall for Pintos. This syscall takes a linked list as one of its arguments, and your responsibility is to create a function that checks the validity of the linked list, so the syscall handler can safely dereference it. The value of head may be NULL, if the list is empty. This is function signature of your syscall:

```
void foo(struct node *head);
```

We have also defined the structure of the node struct, as well as `exit_thread_if_invalid()`, a function that will verify the validity of user-specified pointers.

```
struct node
{
    int number;
    char buffer[16];
    struct node *next;
    // If this is the last node, next = NULL
};
```

```
/* Checks that P to P+SIZE-1 is a valid user buffer.
   Kills the current thread if it is invalid. */
void exit_thread_if_invalid(void *p, size_t size);
```

```
uint32_t SYS_F00 = 60; // The syscall number for foo()
```

Fill in the `syscall_handler()` function *on the next page* so that it *safely* handles the `SYS_F00` syscall.

NAME: _____

```
static void
syscall_handler (struct intr_frame *f)
{
    uint32_t* args = ((uint32_t*) f->esp);

    // Check if this is SYS_F00
    if (_____ ) {

        // This will actually handle SYS_F00,
        // assuming arguments are valid
        handle_foo(args, &f->eax);
    } else {
        // Code for ALL other syscalls will go here.
    }
}
```