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

Spring 2024 John Kubiatowicz

Midterm III

April 25th, 2024

CS162: Operating Systems and Systems Programming

Your Name:	
Your SID:	
TA Name:	
Discussion Section Time:	

General Information:
This is a **closed book** exam. You are allowed 3 pages of notes (both sides). You have 2 hours to complete as much of the exam as possible. Make sure to read all of the questions first, 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.* On programming questions, we will be looking for performance as well as correctness, so think through your answers carefully. If there is something about the questions that you believe is open to interpretation, please ask us about it!

Problem	**Possible**	**Score**
1	18	
2	20	
3	23	
4	17	
5	22	
Total	100	

[This page left for π]

3.14159265358979323846264338327950288419716939937510582097494459230781640628620899

Problem 1: True/False [18 pts]

Please *EXPLAIN* your answer in TWO SENTENCES OR LESS (Answers longer than this may not get credit!). Also, answers without an explanation *GET NO CREDIT*.

Problem 1a[2pts]: It is possible to make Remote Procedure Calls (RPCs) with integer arguments between clients that use big-endian integers and servers that use little-endian integers.

□ True □ False

Explain:

Problem 1b[2pts]: In the Pintos kernel, if you call intr_disable() and then dereference a NULL pointer, your kernel will be stuck in an infinite loop, because the page fault trap cannot reach the CPU when interrupts are disabled.

□ True □ False

Explain:

Problem 1c[2pts]: The Clock algorithm is used to find a physical page that has not been accessed in a while, thereby approximating a MIN replacement algorithm.

□ True □ False

Explain:

Problem 1d[2pts]: The FAT file system links together blocks in a file by including a pointer in each block on disk to the next block on disk. In a FAT32 file system, this means that each block stores 4 bytes less of user data than the size of the block on disk.

□ True □ False

Explain:

Problem 1e[2pts]: An M/M/1 queue has a deterministic distribution for both arrival times and service times.

□ True □ False

Explain:

Problem 1f[2pts]: A storage system protected by RAID6 can identify up to two faulty drives by looking at the data returned from reads.

□ True □ False

Explain:

Problem 1g[2pts]: Journaling file systems use a log to make multi-operation updates atomic in the face of OS crashes.

□ True □ False

Explain:

Problem 1h[2pts]: For modern disk drives, the transfer rate for sectors on outer cylinders is higher than for sectors on inner cylinders.

□ True □ False

Explain:

Problem 1i[2pts]: SSDs utilize FLASH memory which wears out as it is written.

□ True □ False

Explain:

Problem 2: Multiple Choice [20pts]

Problem 2a[2pts]: Which of the following best describes the source of the Meltdown security flaw (made public in 2018)? (*Choose one):*

A: ◯ A system call implemented by multiple operating systems neglected to properly check arguments and could thus be fooled into returning the contents of protected memory to users.

B: ◯ User code could execute speculative loads to addresses marked as "kernel-only" in a user's page table and use the result to effect the state of the cache before the load results could be squashed by the processor pipeline.

C: ◯ Unmapped physical pages (i.e. that were not in an page table) could be accessed by user-level code because of a timing bug in the x86 processor between TLB lookup and cache access.

D: ◯ Linux version 2.4 had a stack-overflow bug which would allow it to be tricked into mapping kernel pages into the user page table.

E: ◯ Speculative execution of arithmetic operations on out-of-order pipelines could be made to reveal the contents of kernel memory by tricking the kernel into changing the timing of operations based on the contents of its private data.

Problem 2b[2pts]: Little's Law has the following properties (*Choose all that apply):*

A: ☐ It applies only to memoryless arrival distributions.

B: ☐ It says that the average number of jobs in the system is equal to the average arrival rate of jobs multiplied by the average length of time that a job stays in the system.

C: ☐ It explains why the average length of time spent in a queue grows without bound as the system utilization approaches 100%.

D: ☐ It can be applied to systems that are not in equilibrium.

E: ☐ None of the above.

Problem 2c[2pts]: Memory-mapped I/O is *(Choose one):*

A: ◯ A security mechanism for I/O devices that that prevents user-mode applications from directly accessing these devices, forcing device access to go through the system call interface.

B: ◯ A software protocol that constructs a coherent distributed shared memory between multiple physical nodes, allowing communication between nodes to occur as reads and writes to the shared memory (address) space.

C: ◯ A technique for communication between processes using the mmap() system call. As a result, processes can interact via reads and writes to shared memory addresses

D: ◯ A hardware mechanism for assigning physical memory addresses to devices such that processor read and write operations to these addresses become commands to devices.

E: ◯ None of the above.

Problem 2d[2pts]: Which of the following are true regarding buffer caches? *(Choose all that apply):*

A: ☐ In file systems without a journal, delayed writes can allow temporary files to be created, written, read, and deleted without ever impacting the disk.

B: ☐ Since FIFO is an approximation to MIN, we can easily handle buffer cache replacement by linking all buffer cache pages into a circular list.

C: ☐ The buffer cache can improve file system performance in situations in which multiple processes simultaneously write to random parts of the disk.

D: ☐ The buffer cache can participate in the prevention of compulsory misses during sequential reads from large files.

E: ☐ The buffer cache is a write-through cache.

Problem 2e[2pts]: Which of the following are true about the Pintos file system? *(Choose all that apply):*

A: ☐ File names for a given directory are stored as strings inside the directory's inode.

B: ☐ A directory can only hold up to a maximum number of directory entries.

C: ☐ In order to resolve a relative path passed into the open syscall, the root inode sector has to be read from the disk.

D: ☐ A directory entry always corresponds to an inumber on the disk.

E: ☐ The directory path "/../../" is invalid.

Problem 2f[2pts]: Which of the following statements about I/O performance are true? *(Choose all that apply):*

A: ☐ The elevator algorithm can improve hard disk drive performance by handling I/O requests in order of physical location rather than in order of arrival.

B: ☐ For many I/O devices, the effective bandwidth of a request can be improved by increasing the size of the request (in bytes).

C: ☐ If the requests entering a queue are combined from many different (uncorrelated) sources, then the arrival distribution can be roughly modeled by an exponential (memoryless) distribution.

D: ☐ SSDs have a seek time that is higher than the time to perform a write.

E: ☐ None of the above.

Problem 2g[2pts]: In UNIX, which of the following I/O devices can be accessed using file operations like `open()` and `close()`? *(Choose all that apply):*

A: ☐ Network devices (WiFi, Bluetooth, wired Ethernet).

B: ☐ A mouse.

C: ☐ The internal hard drive.

D: ☐ A serial device connecting to a terminal.

E: ☐ A printer.

Problem 2h[2pts]: Which of the following are true about modern hard disks? (*Mark all that apply):*

A: ☐ They have an independent disk head on every surface of every platter. As a result, the disk can simultaneously read or write from different tracks on different platters.

B: ☐ Some of them gain bit density by overlapping tracks.

C: ☐ Their internal controllers contain memory for caching, allowing them to read a whole track at a time.

D: ☐ They have a lower bit density on the outside tracks from the inside tracks (because the surface of the outside tracks move under the disk head faster than that of the inside tracks).

E: ☐ Their internal controllers can queue requests and perform variants of the elevator algorithm without consulting the operating system.

Problem 2i[2pts]: The Berkeley FFS has the following properties (*Mark all that apply)*:

A: ☐ It changed the inode format from the BSD 4.1 file system to better reflect the overwhelming presence of small files in a typical UNIX filesystem.

B: ☐ It reserved 10% of the blocks to ensure that new files could get sequential groups of blocks for better read performance.

C: ☐ It performed skip-sector allocation of blocks to prevent processor delays during reading from missing blocks and forcing a complete rotation for each block read.

D: ☐ It placed the inodes and blocks for files within a directory close to the blocks and inodes for the directory to improve performance.

E: ☐ It introduced a B-tree format for directories in order get better lookup performance for directories with large numbers of files.

Problem 2j[2pts]: Which of the following are true about the Clock Algorithm? *(Choose all that apply):*

A: ☐ To examine the state of a particular *physical* page and thus determine if it should be replaced, the Clock Algorithm must examine all PTEs (in all of the page tables) that point at the physical page.

B: ☐ The Clock Algorithm arranges *physical* pages in a ring and keeps a pointer (clock hand) to a page within the ring. During a page fault, it uses the first physical page pointed at by the clock hand to hold incoming data, after which it advances the clock hand to the next page.

C: ☐ The Clock Algorithm arranges *physical* pages in a ring and keeps a pointer (clock hand) to a page within the ring. During a page fault, it chooses a page to hold incoming data by looking at pages one at a time, starting with the current clock hand, until it finds one that hasn't been used recently.

D: ☐ The Clock Algorithm provides an approximation to LRU.

E: ☐ The Clock Algorithm arranges *physical* pages into two groups – an active group that is mapped as "valid" and managed in a FIFO list and an inactive group that is mapped as "invalid" and managed as an LRU list.

[This Page Intentionally Left Blank]

Problem 3: Potpourri [23pts]

Problem 3a[2pts]: In Pintos, what is the significance of the PHYS_BASE constant? *No more than two sentence answer.*

Problem 3b[3pts]: Why is it important for the page fault exception to be precise? Make sure that you define "precise exception" in your answer. *No more than 2 sentences for the definition and one sentence for why the page fault exception needs to be precise.*

Problem 3c[8pts]: For the following problem, assume a hypothetical machine with 4 pages of physical memory and 7 pages of virtual memory. Given the access pattern:

`A B C D E A A E C F F G A C G D C F`

Indicate in the following table which pages are mapped to which physical pages for each of the following policies. Assume that a blank box matches the element to the left. We have given the FIFO policy as an example. Complete LRU and CLOCK. For CLOCK, assume that we check the page under the clock hand, then increment; also, assume that pages brought in have Use=0 initially.

Access→		**A**	B	C	D	E	A	A	E	C	F	F	G	A	C	G	D	C	F
FIFO	1	A				E									C				
	2		B				A										D		
	3			C							F								
	4				D								G						
LRU	1																		
	2																		
	3																		
	4																		
CLOCK	1																		
	2																		
	3																		
	4																		

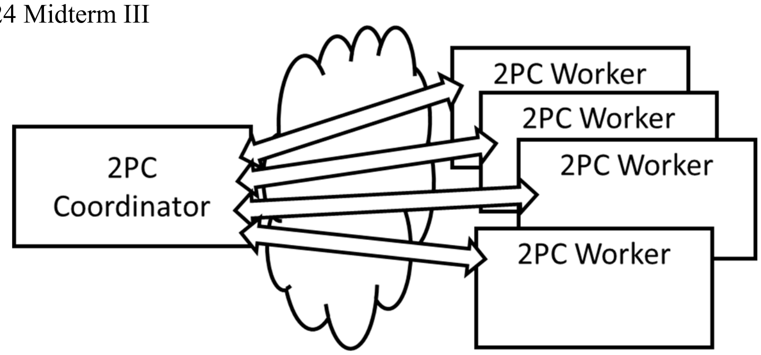

Two-Phase Commit examples: In the following scenarios, consider the following responses to a failure condition. Answer if it renders the system SAFE or UNSAFE and explain why:

Problem 3d[2pts]: The Coordinator sends a *VOTE-REQ* to the Workers and fails immediately. After recovering, it sends a *GLOBAL-ABORT* to all Workers.

□ SAFE □ UNSAFE

Explain:

Problem 3e[2pts]: Suppose a Worker sends a vote to *VOTE-COMMIT* to the Coordinator and does not hear back a response. After some timeout period, this means that the Coordinator lost its vote, so the Worker aborts.

□ SAFE □ UNSAFE

Explain:

Problem 3f[2pts]: Suppose the Coordinator had a 30 second timeout frame. A Worker logs their *VOTE-COMMIT* and before it can send its vote, it crashes and recovers within 2 seconds. This Worker sends its vote to the Coordinator and crashes and recovers again. Then this Worker sends its vote to the Coordinator for the second time.

□ SAFE □ UNSAFE

Explain:

Problem 3g[2pts]: Explain why RAID5 is no longer considered sufficient protection for storage systems with large hard disk drives. *No more than two sentence answer.*

Problem 3h[2pts]: List two reasons that SSDs must have a Flash Translation Layer (FTL). *One sentence per reason.*

[This Page Intentionally Left Blank]

Problem 4: Pintos/File Systems [17pts]

Consider the following inode_disk and indirect_block struct similar to what is in Pintos. Suppose sectors are 512 bytes.

```
struct inode_disk {
    block_sector_t direct[12];
    block_sector_t indirect;
    block_sector_t triply_indirect; // Note that this is a triply
                                    // and that there is no doubly!
    size_t size;
    void* unused[113];
};

struct indirect_block { // doubly and triply indirect blocks look like this
    block_sector_t block_nums[128];
}
```

Problem 4a[2pts]: What is the maximum file size supported by this design? Leave your answer unsimplified in the box:

Problem 4b[4pts]: How many sectors of each type are required to represent a maximum sized file in this design? Leave your answer unsimplified in the box:

Data blocks:

Indirect blocks:

Doubly indirect blocks:

Triply indirect blocks:

Problem 4c[6pts]: Suppose that we have a file that is represented using the inode structures of (4a) and (4b) and that the current size of the file is 12 * 512 bytes long. We now want to write 512 characters to the end of this file. You may use the following functions:

```
void block_read(void *buffer, block_sector_t block_num);
void block_write(void *buffer, block_sector_t block_num);
block_sector_t block_allocate();
```

The block sector corresponding to the file's inode_disk struct is 3:

```
block_sector_t FILE_INODE_BLOCK = 3;
```

Fill in the blanks in the following function such that after running this function, 512 characters are written to this file and persisted. Assume that the file is exactly 12 blocks long already. You may not need all of the blanks, but may only have 1 semicolon per line.

```
void Append512Chars () {
    struct inode_disk inode;

    ____________________________________________________________

    ____________________________________________________________

    struct indirect_block indirect_block;
    indirect_block.block_nums[0] = block_allocate();
    char buffer[512];
    memset(buffer, 'a', 512);

    ____________________________________________________________

    ____________________________________________________________

    ____________________________________________________________

    ____________________________________________________________
}
```

Problem 4d[2pts]: Explain what the mmap() system call does and how it can be used to enable fast communication between two processes. *No more than two sentence answer.*

Problem 4e[3pts]: 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: *Only one sentence for each category:*

Advantage 1:

Advantage 2:

Disadvantage:

[This Page Intentionally Left Blank]

Problem 5: Network Disk System [22pts]

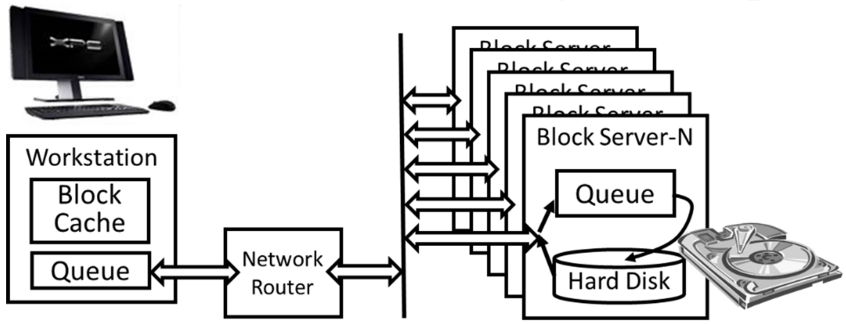

In this problem, we are going to build a network storage system as shown above. This system is going to serve storage blocks over the network to the workstation. This is a so-called "iSCSI" system, in which the client sends requests over the network to remote disks using the iSCSI protocol. The iSCSI protocol takes the normal SCSI commands that would be sent to a disk and "wraps" them with TCP/IP and iSCSI overhead to transport over the network.

The workstation will be responsible for building the file system out of blocks that it accesses. Consequently, the workstation will contain the block cache for this file system. Rather than requesting blocks from a local set of disks, it will send requests over the network to a set of block servers, each of which has its own queue and disk.

Properties of the network are as follows:

- Network bandwidth: **1Gb/s.** Assume that transmission errors do not occur.
- Maximum packet size in network (MTU): **9KB** (so-called "jumbo frames")
- Overhead for packets: **40B** (this is the TCP/IP packet overhead) + **60B** (iSCSI protocol). So, assume that every iSCSI packet is **100B** + data size. *Request packets have no data.*
- The router is fully pipelined and can forward packets at network speed with a **2.2μs** delay.

Properties of the disks on the block servers are as follows:

- There are **10** servers. Each disk is **16TiB** in size with a **4096B** sector size.
- Disks rotate at **10,000 RPM**, have a data transfer rate of **65,536 KB/s**, and have a **5.8ms** average seek time. They also have a SCSI interface with a **200μs** controller time. Assume that a group of consecutive sectors can be fetched with a single request.
- Our use of the disks to build a file system has a disk service distribution with **C=1.5**.
- Each disk can handle only one request at a time, but each disk in the system can be handling a different request.

EACH OF THE FOLLOWING ANSWERS SHOULD BE SIMPLIFIED TO SINGLE NUMBERS. HOWEVER, YOU MUST SHOW YOUR WORK TO GET CREDIT.

Problem 5a[3pts]: Suppose that the server takes **1μs** to receive packets after all the bits have arrived (this is the interrupt service time). What is the latency for a request to make it from the client to one of the servers? *Hint: for network piece, since router is pipelined, compute transmission time for request (iSCSI packet without data) along one network hop, then simply add router delay.*

Problem 5b[3pts]: Suppose that the network controller on clients takes **1μs** to receive packets (this is the interrupt service time). Also, suppose that the disk controller on servers is capable of DMA directly into the network. How long will it take to send a sector worth of data from server to client after the disk controller retrieves the data from disk? *Hint: don't forget the iSCSI packet overhead.*

Problem 5c[4pts]: What is the average *service time* to retrieve a *single* sector from a random location on a *single* disk, assuming no queuing time (i.e. the unloaded request time)? *Hint: there are four terms in this service time! Note: 4096 x 16 = 65,536 and 1/16=0.0625*

Problem 5d[2pts]: How many sequential sectors would we have to combine together into a block in order to achieve an *effective transfer rate* of at least 10% of the data transfer rate when reading a block of data?

Problem 5e[3pts]: Assuming that we arrange to keep all 10 of our servers busy with large requests, such as in **(5d)**, what fraction of our network bandwidth will be used for responses? You can assume that the protocol will never split a sector across network packets. Express your answer as a fraction. *Hint: consider using the effective transfer rate from* ***(5d)****. Also, compute number of complete sectors that will fit into a jumbo packet and account for the iSCSI packet overhead as a fractional increase to data bytes sent.*

Problem 5f[4pts]: Suppose that the distribution of arrivals is memoryless. Assume that we choose to use the large block size from **(5d)** and do not want add more than 75% latency because of queueing time. Also assume that the combination of file system layout, access behavior, and disk characterization leads to a coefficient of variance **C=1.5**. What is our target utilization on each node? How many requests will be queued on average in each node (leave as a fraction)? *Hint: Little's law might be helpful here, as would be* $u=\lambda T_{ser}$

Problem 5g[3pts]: Assume that we chose to use the large block size from **(5d)** and keep the latency increase due to queueing to be 75%, as mentioned in **(5f)** and that we assume data is stored using RAID6 (which means that 2 of the 10 disks are used for redundancy), what is the bandwidth of actual data we can get from our system? You can leave this result as a product of integers.

[This Page Intentionally Left Blank]

[Scratch Page: Do not put answers here!]

[Scratch Page: Do not put answers here!]