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

Spring 2023 John Kubiatowicz

Midterm III

April 27th, 2023

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.* If there is something about the questions that you believe is open to interpretation, please ask us about it!

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

[This page left for π]

3.14159265358979323846264338327950288419716939937510582097494459230781640628620899

Problem 1: True/False [20 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]: The clock algorithm provides an approximation to a FIFO allocation policy, since it links physical pages in a large circular queue.

□ True □ False

Explain:

Problem 1b[2pts]: An empty Amazon Kindle gains weight when you add electronic books to it.

□ True □ False

Explain:

Problem 1c[2pts]: When producing a queueing model of some system (such as a filesystem) and the distribution of request arrival times is unknown, the best estimate is to use a *deterministic* distribution.

□ True □ False

Explain:

Problem 1d[2pts]: For disks with constant *areal bit density* (bits stored/unit area of disk media), the disk head reads bits at a different rate on the outer tracks than on the inner tracks.

□ True □ False

Explain:

Problem 1e[2pts]: The Bitcoin blockchain could be used for distributed decision making.

□ True □ False

Explain:

Problem 1f[2pts]: Direct Memory Access (DMA) can reduce the load on the processor when accessing block devices such as disks.

□ True □ False

Explain:

Problem 1g[2pts]: Supposing that during Two Phase Commit for transaction T, the following occurs:

1. The Coordinator sends VOTE-REQ to all workers
2. Worker A sends VOTE-COMMIT to the Coordinator
3. Worker A dies
4. The Coordinator sends GLOBAL-COMMIT to all workers
5. Worker A recovers

Then Two Phase Commit for T fails because Worker A never received the GLOBAL-COMMIT message, so T was never committed to its log.

□ True □ False

Explain:

Problem 1h[2pts]: Network Address Translation (NAT) automatically switches the IP address of a remote server to spread load over multiple servers (thereby improving overall response time).

□ True □ False

Explain:

Problem 1i[2pts]: Since the File Allocation Table (FAT) does not actually store file contents, the FAT is not stored on disk.

□ True □ False

Explain:

Problem 1j[2pts]: Because RAID 5 is an instance of an erasure code, it requires individual disks to be able to identify when they have failed using a different code.

□ True □ False

Explain:

Problem 2: Multiple Choice [20pts]

Problem 2a[2pts]: Which of the following is *true* about Pintos' file system? **(*Choose one):***

A: ☐ Directory data (i.e., directory entries) are stored on the directory's inode on disk.

B: ☐ Two threads can always write to the same file concurrently.

C: ☐ The seek syscall should never require blocks in the file to be allocated or freed.

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

E: ☐ None of the above.

Problem 2b[2pts]: Which of the following is *false* about Pintos' file system? ***(Choose one):***

A: ☐ To determine the end of a directory when reading directory entries, PintOS relies on an offset with a file length.

B: ☐ Once the buffer cache has been implemented, inodes should no longer store the data for inode_disk.

C: ☐ To enforce synchronization, we add a lock to inode_disk.

D: ☐ Extending a file requires that you fill up the empty space with zeros.

E: ☐ None of the above are false (i.e. all of the above are true).

Problem 2c[2pts]: Which of the following is *true* about distributed decision making protocols ***(Choose one):***

A: ☐ The two-phase commit protocol is tolerant to Byzantine failures.

B: ☐ In the two-phase commit protocol, failure of the coordinator at any time before it writes "commit" into its log will cause the transaction to be aborted.

C: ☐ If a worker in the two-phase commit protocol fails *after* voting to commit, it will commit the transaction when it later recovers.

D: ☐ The PAXOS protocol is less robust than two-phase commit because it can block on a failed coordinator.

E: ☐ None of the above.

Problem 2d[2pts]: Which of the following are true of the Fast File System (FFS)? ***(Mark all that apply):***

A: ☐ FFS 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: ☐ Direct pointers in the inode allow for reduced access times for small files.

C: ☐ Free blocks are tracked using a bitmap, allowing for files to be more easily allocated in contiguous memory.

D: ☐ It placed all the inodes together on the inner tracks of the disk for better performance.

E: ☐ FFS keeps 10% of disk blocks free to help with contiguous block allocation of files (and to improve overall read performance).

Problem 2e[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: ☐ Their internal controllers can queue requests and perform variants of the elevator algorithm without consulting the operating system.

C: ☐ They have internal 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: ☐ None of the above.

Problem 2f[2pts]: Little's Theorem has the following properties ***(Mark all that apply):***

A: ☐ It applies only to memoryless arrival distributions.

B: ☐ It can be used to compute the average time spent in a queue, given the average length of the queue and average arrival rate.

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

D: ☐ It applies only to systems in equilibrium.

E: ☐ None of the above.

Problem 2g[2pts]: Which of the following statements about I/O performance are true? ***(Mark 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 2h[2pts]: Which of the following statements about I/O performance are *false*? ***(Choose one)*:**

A: ☐ Little's law applies to all stable systems, regardless of structure, bursts of requests, or variation in service.

B: ☐ It is impossible for utilization to be 1 in any stable system.

C: ☐ Utilization can be computed by λ/μ, where μ is the completion rate in jobs/second and λ is the arrival rate in jobs/second.

D: ☐ The equation $T_{ser} \times u / (1 - u)$ calculates the amount of time that a job spends waiting in the queue for a device with an exponential service time.

E: ☐ None of the above are false (i.e. all of the above are true).

Problem 2i[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 2j[2pts]: The Meltdown security flaw (made public in 2018) was able to gain access to protected information in the kernel because ***(Choose one)*:**

A: ☐ Data-specific variability in the processing of system-calls (related to how Intel processors handled synchronous exceptions) by one process (the victim) allowed another process (the attacker) to successfully guess values stored in the kernel stack of the victim process.

B: ☐ A POSIX-compatible 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.

C: ☐ There was a wide-spread bug in the x86 architecture that caused certain loads to ignore kernel/user distinctions in the page table under the right circumstances. This allowed user-code to directly load and use data that was supposed to be protected in kernel space.

D: ☐ Many processors allowed timing windows in which illegal accesses could be performed speculatively and made to impact cache state – even though the speculatively loaded data was later squashed in the pipeline and could not be directly used.

E: ☐ None of the above.

Problem 3: Short Answer Potpourri [23pts]

Problem 3a[5pts]: Suppose that a system uses an N^{th}-chance clock algorithm for demand paging. Also assume that physical memory has 8 pages and each PTE has a dirty bit, a use bit, and a counter. Assume that the clock algorithm evicts dirty pages when its counter reaches 2, and evicts clean pages when its counter reaches 1. Assume that the clock hand starts at PPN 7 (wrapping around to PPN 0) and that the algorithm increments *first* before examining control bits. Assume that the physical pages are allocated to virtual pages as follows:

PPN	Contains VPN	Dirty	Use	Counter
0	1	1	1	0
1	3	0	1	0
2	5	0	0	0
3	2	0	0	0
4	8	0	1	0
5	12	1	0	1
6	9	1	0	0
→ 7	7	0	1	0

For the following access pattern, provide the PPN of the page frame at the time that the access succeeds. Assume all accesses are reads. We have filled the first entry for you:

Accessed VPN 3: PPN =? **1**
Accessed VPN 2: PPN =? ________________
Accessed VPN 6: PPN =? ________________
Accessed VPN 8: PPN =? ________________
Accessed VPN 4: PPN =? ________________
Accessed VPN 10: PPN =? ________________

Problem 3b[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 D D E F A A C F G D A C G D C E`

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.

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

Problem 3c[3pts]: What is a *precise exception* and why is it the concept helpful for handling page faults during demand paging? (2 pts for definition, 1 pt for reason it is helpful).

Problem 3d[3pts]: What is the mmap() system call? How could it be used to set up shared memory between two processes so that they could share a linked list between them (here sharing means that the two of them could add items to the list and traverse the list assuming that suitable synchronization was employed to avoid inconsistencies):

Problem 3e[2pts]: Prior to the discovery of the Meltdown exploit, operating systems typically mapped some or all of kernel memory into the top of each user's page table (with PTEs that had the "User accessible" bit set to 0). What was the advantage of this practice?

Problem 3f[2pts]: Explain why two-phase commit is not subject to the General's Paradox.

Problem 4: File Systems [20pts]

Problem 4a[3pts]: Suppose that you are producing a Pintos-like file system for which the `inode_disk` contains the following data pointers:

```
#define BLOCKSIZE 1024
typedef uint32_t block_sector_t;

// Assume that this is one BLOCKSIZE in size
struct inode_disk {
      block_sector_t direct_pointers[12];
      block_sector_t indirect_pointer;
      block_sector_t double_indirect_pointer;
      block_sector_t triple_indirect_pointer;
      ...
}
struct indirect_block { // indirect blocks look like this
      block_sector_t block_nums[BLOCKSIZE/sizeof(block_sector_t)];
}
```

What is the maximum file size (in bytes) that can be supported by this filesystem? (You can leave an unsimplified expression for your answer).

Problem 4b[4pts]: Utilizing the inode structure from Problem 4a, what is the *minimum* and *maximum* number of disk accesses required to read 512 bytes from a file? Assume that the buffer cache has not been implemented, but that the contents of an inode and various types of indirect blocks can be reused during the processing of a single request. Explain your answer. *(Hint: assume that you are starting from anywhere within a very large file, so that all indirect pointers are in use).*

Problem 4c[4pts]: Suppose that you are building a file system on top of an SSD using the inode structure from Problem 4a. Assume that the SSD has the following parameters:

- Page size (i.e. minimum size of items transferred): 1024B
- Controller time: 10 μsec
- Page read time: 40 μsec
- Page write time: 90 μsec

Further, assume that reads or writes to the SSD are *not pipelined*, which means that every access incurs latency through the controller. Assuming there is no buffer cache, how much time is required to create a new file that is 256 KiB in size? Assume that this is the second file created in the root directory. Show your work and give the answer in milliseconds (ms):

Problem 4d[4pts]: At a particular point in time, the buffer cache has dirty data that needs to be flushed to disk. Suppose that the identities of these blocks can be listed in [track:sector] form as follows:

[10:5], [22:9], [11:6], [2:10], [20:5], [32:4], [32:5], [6:7]

Assume that the disk head is currently positioned over track 20. What is the sequence of writes under the following disk scheduling algorithms:

a) Shortest Seek Time First:

b) Scan (initially moving upwards):

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 (one sentence each):

Advantage 1:

Advantage 2:

Disadvantage:

Problem 4f[2pts]: What is a "Journaled" filesystem and how does it improve the durability of data on a disk relative to a system such as (4a). Give your answer in two sentences or less.

Problem 5: Parallel Disk Systems [17pts]

NOTE: IN THE FOLLOWING PROBLEM, ALL OF THE NUMBERS HAVE BEEN CHOSEN SO THAT ANSWERS SIMPLIFY EASILY WITHOUT THE NEED FOR A CALCULATOR.

Suppose that we build a disk array to handle a high rate of I/O by coupling many disks together. Properties of this system are as follows:

- Has a total of 16 disks, each of which is 2TiB in size
- Uses disks that rotate at 15,000 RPM, have a data transfer rate of 32768 KBytes/s (for each disk), and have an 3.5 ms average seek time, 4096 Byte sector size.
- Has a SCSI interface with a 500μs controller command time. Assume that a group of consecutive sectors can be fetched with a single request.
- Is limited only by the disks (assume that no other factors affect performance).

Each disk can handle only one request at a time, but each disk in the system can be handling a different request. The data is not striped (all I/O for each request has to go to one disk).

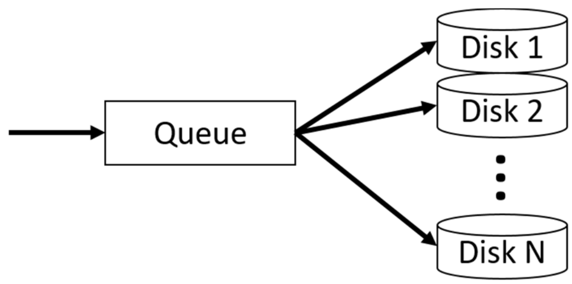

EACH OF THE FOLLOWING ANSWERS SHOULD BE SIMPLIFIED TO SINGLE NUMBERS. HOWEVER, YOU MUST SHOW YOUR WORK TO GET CREDIT.

Problem 5a[6pts]: 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!*

Problem 5b[3pts]: Assume that we plan to build a file system on top of this disk which achieves considerably more locality than spreading sectors randomly across the disk. How many sectors would it need to read in a row in order to achieve an effective bandwidth of 25% of the transfer rate from the disk? Express your answer in number of sequential sectors it would need to read at a time. *Hint: Express your latency from 5a in a form that is a fixed overhead + a component that is linearly related number of sequential sectors to fetch at a time. Solve in a way that the linear component is 25% of the total amount. Show your work!*

Problem 5c[2pts]: Let's use your answer from Problem 5b as our file-system block size. What is the average *service time* to retrieve a *block* from a random location on a *single* disk? Note that the answer to this problem is actually part of the work to solve Problem 5b.

Problem 5d[2pts]: Let's use your answer from Problem 5c. Now, assume that we build a file system on the disk array that can have block-sized requests outstanding for all disks. Allowing each disk request to be for a random block on disk, what is the maximum number of I/Os per second (IOPS) for the whole disk subsystem (an "I/O" here is a block request)? We will give partial credit, so make sure to show your work.

Problem 5e[4pts]: Treat the entire system as an M/M/m queue (that is, a system with one queue but m servers rather than one), where each disk is a server. For simplicity, assume that any disk can service any request. Assuming FIFO scheduling by the OS again, what is the maximum request rate (λ) that the system can handle while increasing the total time to service a block request so that the queue adds no more than 25% to the overall service time from Problem 5c. Show your work!

You might find the following equation for an M/M/m queue (where any server can handle any request from the queue) useful:

$$\text{Server Utilization}\,(\zeta) = \frac{\lambda}{\dfrac{1}{\text{Time}_{\text{server}}/m}} = \lambda \times \frac{\text{Time}_{\text{server}}}{m}$$

$$\text{Time}_{\text{queue}} = \text{Time}_{\text{server}} \times \left[\frac{\zeta}{m(1-\zeta)}\right]$$

[Scratch Page (feel free to remove)]

[Scratch Page (feel free to remove)]