

Honor Code of the School of Engineering

"All students taking courses in the School of Engineering agree, individually and collectively, that they will not give or receive unpermitted aid in examinations or other course work that is to be used by the instructor as the basis of grading."

-From the Graduate/Undergraduate Bulletin

I have read, understood, and agree to abide by the Honor Code of the School of Engineering.

Name: _____

ID: _____

Signature: _____

Date: _____

1. [10 points]
2. [40 points]
3. [50 points]
4. [20 points]
5. [20 points]
6. [20 points]
7. [20 points]
8. [20 points]

Total Score:

Final Examination

CSEN 283 Operating Systems

***Department of Computer Science and Engineering
Santa Clara University***

Dr. Ming-Hwa Wang

Phone: (408) 805-4175

Course website:

Office Hours: Tuesday & Thursday 9-9:30pm and Saturday 11am-12pm

Fall Quarter 2024

Email address: m1wang@scu.edu

<http://www.cse.scu.edu/~m1wang/os/>

1. [10 points] Please give the whole names for the following acronyms: ECC, CC-NUMA, CMP, IIOP, RMS.
2. [40 points] True or false problems with wrong-answer penalties:
 - a. The operating system knows the current state of all processes in the distributed system.
 - b. RAID is a set of physical disk drives viewed by the operating system as a single logical drive.
 - c. Small file I/Os are often dominated by updates to file system metadata (e.g., i-node, file attributes, etc.).
 - d. The I/O access rights are managed at the physical organization layer.
 - e. A true global state cannot be determined in distributed systems because of the time lapse associated with message transfer.
 - f. Does it make sense to para-virtualize an operating system if the source code is available?
 - g. The superuser is exempt from the usual file access control constraints and has system wide access.
 - h. Double buffering is when a process transfers data to (or from) one buffer while the operating system empties (or fills) the other.

3. [50 points] Simple questions:
 - a. Consider a 4-TB disk that uses 4-KB blocks and the free-list method. How many block addresses can be stored in one block?
 - b. Given a disk-block size of 4 KB and block-pointer address value of 4 bytes, what is the largest file size (in bytes) that can be accessed using 10 direct addresses and one indirect block?
 - c. A RAID can fail if two or more of its drives crash within a short time interval. Suppose that the probability of one drive crashing in a given hour is p . What is the probability of a k -drive RAID failing in a given hour?
 - d. When transferring data from RAM to a network interface, pinning a page can be used, but suppose that system calls to pin and unpin pages each take 1 μ sec. Copying takes 5 bytes/nsec using DMA but 20 nsec per byte using programmed I/O. How big does a packet have to be before pinning the page and using DMA is worth it?
 - e. How much cylinder skew is needed for a 7200-RPM disk with a track-to-track seek time of 1 msec? The disk has 200 sectors of 512 bytes each on each track.
4. [20 points] For each of the following applications, indicate whether the system should use RAID level 0, RAID level 1, or RAID level 5 and explain why.
 - a. Storing account transactions for a financial institution
 - b. Storing high-resolution telescope images (that have been transferred from a tape drive) before processing on a supercomputer
 - c. Running a personal Web server
5. [20 points] Consider a file currently consisting of 100 blocks. Assume that the file control block (and the index block, in the case of indexed allocation) is already in memory. Calculate how many disk I/O operations are required for contiguous, linked, and indexed (single-level) allocation strategies, if for on block, the following conditions hold. In the contiguous-allocation case, assume that there is no room to grow at the beginning but there is room to grow at the end. Also assume that the block information to be added is stored in memory.
 - a. The block is added at the beginning.
 - b. The block is added at the middle.
 - c. The block is added at the end.
 - d. The block is removed at the beginning.
 - e. The block is removed at the middle.
 - f. The block is removed at the end.
6. [20 points] Suppose that a disk drive has 5000 cylinders, numbered 0 to 4999. The drive is currently serving a request at cylinder 143, and the previous request was at cylinder 125. The queue of pending requests, in FIFO order, is 86, 1470, 913, 1774, 948, 1509, 1022, 1750, 130. Starting from the current head position, what is the total distance (in cylinders) that the disk arm moves to satisfy all the pending requests, for each of the following disk-scheduling algorithms? a) FCFS, b) SSTF, c) SCAN, d) LOOK, e) C-SCAN, and f) C-LOOK.
7. [20 points] Consider a program that accesses a single I/O device and compare unbuffered I/O to the use of a buffer. Show that the use of buffer can reduce the running time by at most a factor of two. Generalize the result to the case in which a program refers to n devices instead of one. **Hint:** If the calculation time exactly equals the I/O time (which is the most favorable situation), both the processor and the peripheral device running simultaneously will take half as long as if they ran separately.
8. [20 points] Let a be the percentage of program code that can be executed simultaneously by n computers in a cluster, each computer using a different set of parameters or initial conditions. Assume that

the remaining code must be executed sequentially by a single processor. Each processor has an execution rate of x MIPS.

- a. Derive an expression for the effective MIPS rate when using the system for exclusive execution of this program, in terms of n , a , and x .
- b. If $n = 16$ and $x = 4$ MIPS, determine the value of a that will yield a system performance of 40 MIPS.