

UNIVERSITY OF CALIFORNIA
College of Engineering
Department of Electrical Engineering
and Computer Sciences
Computer Science Division

CS 162, FALL, 2004

Prof. Alan Jay Smith

Midterm 2, Monday, November 8, 2004

You have until the time announced for this exam. The exam is **closed book, closed notes, and no laptops, cellphones or programmable calculators**. All answers should be written on the exam paper; you may continue on the back of a page if necessary. Anything that we can't read or understand won't get credit. Any questions for which you give *no answer at all will receive 25% partial credit*. Please answer in standard English; illiterate or illegible answers to essay questions will lose credit. If we can't understand your answer, it is wrong. Partial credit will be given for "computation-type" problems only if your errors are obvious to the grader; we are not going to spend significant time debugging your solution. Please watch the front board for corrections and other information. This exam has 7 questions on 7 pages. You should put your name on every page; every missing (name) entry will cost you one point.

1. Is the throughput in an open system affected by the schedule algorithm, assuming that $\rho < 1$? How about for a closed system? Why or why not (in both cases)? Why do we care if ρ is less than 1?
(10)

2. The Unix, file system file descriptor contains 15 pointers, which points to the first 12 file blocks, and then to 3 subtrees, respectively 1, 2, and 3 levels deep. In Unix, all blocks are the same fixed size – e.g. 4Kbytes. Consider the Twinkies file system, in which a file descriptor also contains 15 pointers. In Twinkies, each pointer points to a file system block, but the block size is variable – blocks can be 1K, 4K, 16K, 64K, 256K and 1Mbytes. There are no indirect blocks, so the file may consist of no more than 15 blocks. What are the advantages and disadvantages of the Twinkies design vs. the Unix design? Please consider both writing and reading small, medium and large size files. Please consider and discuss everything relevant. (15)

3. a. Explain what a RAID-5 storage system is, and how it allows you to recover from the failure of a single disk. b. Discuss whether the disk failures can really be considered to be independent. (15)

4. What is thrashing? Why is it a problem? What is the working set paging algorithm and how is it intended to combat the thrashing problem? (10)

5. Disk Scheduling

Suppose you are given a disk with 5000 cylinders (numbered 0-4999). Suppose also that the drive is currently servicing cylinder 143, and previously serviced cylinder 125.

Requests to the following cylinders have been made to the disk in the following time sequence order and are all pending. (i.e. They have not been done yet.)

22 4789 3000 4790 1111 53 1115 2100 11 711 912 962

Consider the following disk arm scheduling algorithms: i) FCFS; ii) SSTF; iii) SCAN; iv) C-SCAN

a. For each of the above algorithms, what is the order in which the requests are serviced for the requests listed above, and what is the total distance moved by the disk head (in cylinders). (show work) (12)

b. In order to determine which is the “best” arm scheduling algorithm, (in general, not particularly for this problem), what would you consider? (List the factors, in their relative order of importance, and explain their relative importance.) (8)

6. Page Replacement (18)

Given the following reference string:

8 7 6 5 8 7 9 8 7 6 5 9

determine the number of page faults when there are 3 or 4 page frames, for each of LRU, OPT, and FIFO replacement. Show your work.

	3	4
LRU		
FIFO		
OPT		

7. We discussed 4 mechanisms to avoid having to allocate and keep in memory very large page tables. (E.g. Assume a 20-bit page number. Then there are 2^{20} page table entries. If each is 4-bytes, then we would need to allocated 4Mbytes of real memory for each page table.) What are these 4 mechanisms? Please explain each briefly. (16)