

國立中山大學資訊工程學系
107學年度第2學期博士班資格考試

科目：作業系統

INSTRUCTIONS: *If any question is unclear or you believe some assumptions need to be made, state your assumptions clearly at the beginning of your answer.*

1. (10%) Consider the two-dimensional array A:

```
double A[] [] = new double[200][200];
```

where each double occupies 8 bytes and A[0][0] is at location 200, in a paged system with pages of size 200 bytes. A small process is in page 0 (locations 0 to 199) for manipulating the matrix; thus, every instruction fetch will be from page 0. For three page frames, how many page faults are generated by the following array initialization loops, using LRU replacement and assuming (1) page frame 0 has the process in it, (2) the other two are initially empty, and (3) the array is stored in memory column-major?

- (a) for (int i = 0; i < 200; i++)
 for (int j = 0; j < 200; j++)
 A[i][j] = 0;
- (b) for (int j = 0; j < 200; j++)
 for (int i = 0; i < 200; i++)
 A[i][j] = 0;
2. A disk has 12000 cylinders, each with 32 tracks of 512 blocks. A seek takes 1 ms per cylinder moved. If no attempt is made to put the blocks of a file close to each other, two blocks that are logically consecutive (i.e., follow one another in the file) will require an average seek, which takes 6 ms. If, however, the operating system makes an attempt to cluster related blocks, the mean interblock distance can be reduced to 2 cylinders and the seek time reduced to 200 μ s. Assuming that the rotational latency is 3 ms and the transfer time is 10 μ s per block, answer the following questions.
- (a) (5%) How long does it take to read a 200 block randomly placed file?
- (b) (5%) How long does it take to read a 200 block clustered file?
3. Given a UNIX *i*-node with fifteen direct blocks and four levels of indirect blocks (i.e., a single, a double, a triple, and a quadruple) and assuming that the sizes of a pointer and a block are, respectively, 8 bytes and 8 Kbytes, answer the following questions.
- (a) (5%) What would be the size of the smallest file allowed in bytes?
- (b) (5%) What would be the size of the largest file allowed in bytes?
4. (10%) A computer has 32 GB of RAM allocated in units of 4 KB. How many KB are needed if a bit map is used to keep track of free memory?
5. (10%) A computer whose processes have 4092 pages in their address spaces keeps its page tables in memory. The overhead required for reading a word from the page table is 500 nsec. To reduce this overhead, the computer has a TLB, which holds 32 (page, frame) pairs and can do a lookup in 50 nsec. What hit rate is needed to reduce the mean overhead to 150 nsec or less, *assuming that TLB is looked up first before the page table is read*?
6. Assume a page reference string for a process with h frames (initially all empty). The page reference string has length m with n distinct page numbers occurring in it. For any page-replacement algorithms,
- (a) (5%) What is an upper bound on the number of page faults? **Justify your answer for credit.**
- (b) (5%) What is a lower bound on the number of page faults? **Justify your answer for credit.**

7. (10%) A machine has 48-bit virtual addresses and 36-bit physical addresses. Pages are 4 KB. How many entries are needed for the page table?

8. Consider the following page reference string:

1, 2, 3, 4, 2, 1, 5, 6, 2, 1, 2, 3, 7, 6, 3, 2, 1, 2, 3, 6.

How many page faults would occur for the following replacement algorithms, assuming three, four, or five frames? Remember all frames are initially empty, so your first unique pages will all cost one fault each.

(a) (5%) LRU replacement

(b) (5%) Optimal replacement

9. Suppose that a disk drive has 1000 cylinders, numbered from 0 to 999. The drive is currently serving a request at cylinder 200, and the previous request was at cylinder 125. The queue of pending requests, in FIFO order, is

50, 500, 250, 800, 350, 550, 400, 600, 100.

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) (5%) LOOK

(b) (5%) SCAN

(c) (5%) C-LOOK

(d) (5%) C-SCAN