

Operating Systems Comprehensive Exam

Spring 2011

Student ID # _____

2/17/2011

You must complete all of Section I

You must complete two of the problems in Section II

**If you need more space to answer a question, use the back of a page
and indicate which question is being answered.**

Section I: You must complete all problems in this section.

1. When an I/O request is being handled for a user's process, which term refers to the policy of returning control to the user process before the I/O is completed?

synchronous I/O asynchronous I/O delayed I/O none of these

2. Which of the following process scheduling algorithms can potentially result in starvation?

First-Come, First-Served Round Robin Shortest Job Next none of these

3. Given the following segment table:

Calculate the correct physical addresses for the following logical addresses [segment, offset], indicate any addressing errors that occur.

Segment	Base	Length
0	450	170
1	240	110
2	350	100
3	620	180

a) [3, 150] physical address? _____

b) [2, 105] physical address? _____

Matching: choose the best answer for each question from the list below:

4. The term _____ refers to a software-generated interrupt.
5. If a page's _____ is set, the page must be written to disk before it is replaced.
6. _____ provide a programming interface to the services provided by the OS.
7. _____ processes can affect or be affected by the execution of another process.
8. When a process uses a _____ send to deliver a message, it must wait until the receiving process gets the message.
9. The _____ module gives control of the CPU to the process selected by the short-term scheduler.

A. mutex B. job scheduler C. cooperating
D. limit register E. kernel mode F. privileged instructions
G. blocking H. atomic J. dispatcher
K. system calls L. invalid bit M. independent
N. dirty bit O. asynchronous P. trap

10. The working set model for paging is based upon the assumption of _____.

- locality fragmentation static linking none of these

11. If a system is in an unsafe state, it is guaranteed that a deadlock will occur. **True** **False**

Briefly explain your choice, i.e., why is your answer correct?

12. a) Does the following pair of operations correctly implement a *Semaphore* (assume that S is initially 1)?

<pre>wait(S) { while(S < 0); S--; }</pre>	<pre>signal(S) { S++; }</pre>
--	-----------------------------------

Yes **No** If not, briefly explain the error:

b) Assuming that the above Semaphore is either correct, or has been corrected, briefly explain two potential problems with (or disadvantages of) using this Semaphore to control access to a critical section.

13. A file handle (or file descriptor) is a pointer into the _____.

open file table

file control block

directory

file allocation table

14. List the three requirements that must be satisfied to implement a valid critical section:

15. a) Briefly explain the difference between internal and external fragmentation

b) give one example of a file allocation approach that can suffer from internal fragmentation and one example of a file allocation approach that can suffer from external fragmentation

16. Briefly explain the difference between deadlock avoidance and deadlock prevention:

Section II: You must complete two of the following three problems (A, B and C). If you complete more than two problems, clearly indicate which two problems you want graded. Otherwise, only the first two attempted problems will be graded.

- A. Consider the following list of disk access requests, in arrival order. For each disk scheduling algorithm, calculate the number of tracks that the read/write head crosses without stopping. The disk drive has 200 tracks, from 0 to 199, and the read/write head has a starting location of track **45**.

You must **show your work** to receive partial credit, without it an incorrect answer gets zero points.

65, 128, 150, 87, 79, 126, 51, 33

(a) First-Come, First-Served _____

(b) Shortest Seek Time First _____

Consider a paging system with the page table stored in memory.

- (a) If a memory reference takes 10 microseconds, what is the effective memory access time (with no TLB)?

_____ microseconds

- (b) We add a TLB where the time needed to find a page reference in the TLB is 4 microseconds. If the TLB contains 75% of the page references (assume that all TLB misses are still in memory), what is the effective memory access time?

_____ microseconds

B. Suppose that the following processes arrive for execution at the sequence P1, P2, then P3. Each process will run for the amount of time listed. In answering the questions below, use non-preemptive scheduling, and base all the decisions on the information you have at the time the decision must be made. Assume that all three processes are available at the start.

Process	Burst Time
P1	5
P2	2
P3	4

Justify your answer by drawing the corresponding Gantt chart for each case.

(a) What is the average waiting time for these processes using the FCFS scheduling algorithm?

(b) What is the average waiting time for these processes using the SJF scheduling algorithm?

(c) How many context switches will take place in total when using a Round-robin scheduling with quantum = 2?

C. a) Describe Belady's anomaly for page replacement.

b) Given the following list of page references, in execution order:

1, 2, 3, 4, 1, 3, 4, 5, 2, 1, 4, 3

Trace the LRU page replacement algorithm with only 3 frames available for pages (All pages are initially empty.) You must **show your work** to receive full credit.

Trace the LRU page replacement algorithm with only 4 frames available for pages (All pages are initially empty.) You must **show your work** to receive full credit.