

Midterm Exam

Amy Murphy
28 February 2001

Read before beginning: Please write clearly. Illegible answers cannot be graded. Be sure to identify all of your answers in your blue book (e.g., 1a, 2b, etc). All questions are worth 10 points, numbers in parenthesis indicate relative point values.

While I have tried to make the questions as unambiguous as possible, if you need to make any assumptions in order to continue, state these as clearly as possible as part of your answer. In the name of fairness, I would like to avoid answering questions during the exam.

CSC256: Students enrolled in 256 must choose **5 of the 7** questions to answer. For extra credit, you may answer the other 2 questions, indicating clearly on the front of the bluebook which questions are to be counted as extra credit. If you do not specify, I will assume the first 5 in your bluebook are for *normal* credit and the remaining 2 (if any) are *extra*. Grades will be assigned 0-50.

CSC456: Students enrolled in 456 must choose **6 of the 7** questions to answer. For extra credit, you may answer the 7th question, but you must clearly identify which question is to be counted as *extra*. If you do not indicate any, I will assume the 7th answer is *extra*. Grades will be assigned 0-60.

1. **Short answer:** your responses should be at most 3-4 lines of text:

- (a) (4) Name any three main components of most major operating systems (such as Unix or Windows-XX) and briefly describe their role.
 - (b) (3) Explain how the mixture of I/O bound processes with CPU bound processes maximizes system utilization. Why is this more important in batch systems than it is on most computers sitting around in our department?
 - (c) (3) Define a counting semaphore, its initialization, and the two operations used to operate on it.
2. **Scheduling.** Given the five processes below with their indicated number of run time units, answer the questions that follow. Assume processes arrived in numerical order at time 0.

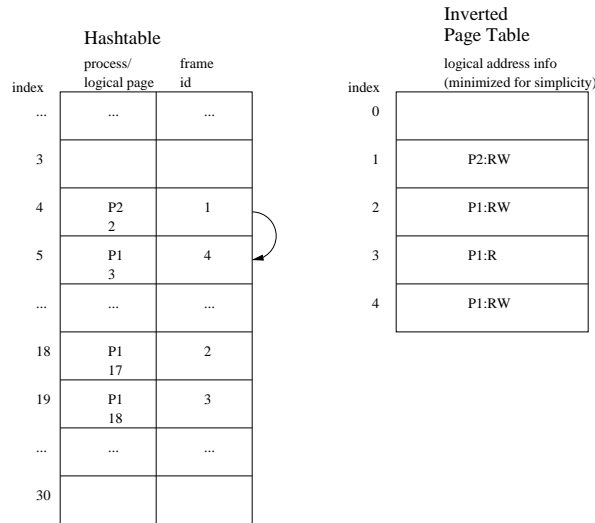
Process ID	CPU requirements
P1	7
P2	1
P3	2
P4	5
P5	3

- (a) (5) Show the scheduling order for these processes under first-come-first-served, shortest-job first, and round-robin scheduling (quantum = 1).
- (b) (4) For each process in each schedule above, indicate the wait time and turnaround time.
- (c) (1) Briefly explain why these two values are meaningful criteria for evaluating scheduling algorithms.

3. **Process Creation.** Modern unix operating systems use copy-on-write in order to implement the `fork()` operation for creating a new heavy-weight process.

- (a) (3) Define the difference between a heavy-weight process and a thread (or lightweight process).
- (b) (4) In class, we discussed copy-on-write for memory pages shared among multiple processes. We cannot apply this same concept blindly to process creation, but instead are forced to copy some parts immediately while other parts can be delayed. Knowing the components of general processes, which parts must be copied immediately, and which parts can be delayed and *copied-on-write*?
- (c) (3) Why is copy-on-write potentially better than copying the entire process immediately upon creation?

4. **Inverted Page Tables.** Assume we have a simple, demand paging environment, with no segmentation. Processes $P1$ and $P2$ both have logical memory addresses in the range $0 \dots 99$, inclusive. The page size is 5. The hash table portion of the structure uses a hash function which simply calculates the index by adding the numerical portion of the process id and the logical page id. (Note, this is a bad hash function because it requires a large hash table, but for our exercise, this is good enough.) The hash table and the inverted page table are below. Arrows in the hash table indicate chaining of entries when a hash conflict has arisen. Empty entries in the IPT indicate the physical frame is not allocated.



- (a) (6) Calculate the physical addresses for the following logical addresses where the number after the colon is the logical address (not the logical page). Assume the first page of a process has id 0: P1:17. P1:92. P1:111. Show you work for these calculations (in other words, what calculations/table lookups were necessary to get the actual address in memory from the logical address).
- (b) (4) When process P2 attempts to read logical address 97, what happens? Describe the changes in the data structures, and what the process perceives of these changes. Assume a mechanism exists to find the free frame in memory.

5. **Synchronization.** In parallel programs (a single process with multiple threads), a common way to design the processing is in stages. All threads work independently during each stage, but must synchronize at the end of each stage. When all threads reach this synchronization point (called a barrier), they are notified and begin execution on the next phase of the computation.

Two potential complications: In many cases, there is no master thread watching over the children threads, waiting for each of them to get to the barrier, and then telling them to re-start. In other words, the children must monitor themselves! Second, it is often not known in advance how many children threads there will be during the lifetime of the parallel program. In other words, a child can spawn another child (realizing that this new child will start in the same program stage as the child which created it)! In all cases, all children must synchronize at the barrier before the processing is allowed to continue to the next phase.

- (a) (8) Your task is to provide the pseudo code for a monitor class called **Barrier** which enables this style of barrier synchronization. Things to take into consideration are the creation of a new child thread (one more process that needs to synchronize), processing when a thread reaches the barrier, and the releasing of waiting threads when the *last* thread reach the barrier. Remember that some of the standard monitor methods available to you are `wait()` and `signal()` and `signalAll()`. (Hint: if you know about Java synchronized objects, this concept is very similar.)
- (b) (2) Did you put the call(s) to `wait()` at the end of your method(s)? If so, why? If not, what possible complication might arise after your call to `signalAll()`?
6. **Deadlocks.** Although deadlock is a real problem, it is often ignored during the design of the operating system due to the inefficiencies which deadlock prevention and/or deadlock detection introduce. One approach to reduce the overhead is to group resources into multiple classes, use a total ordering among these classes, and use a technique tailored to the class for resource allocation within that class. For example, four reasonable classes may be (1) swap space, (2) process resources such as files, tape drives, etc. (3) main memory, (4) internal resources such as communication channels among processes.
- (a) (4) Is the sequence C1:R2, C1:R1, C3:R5, C1:R4, C4:R1 allowed (assuming that in C_x, the 'x' refers to the class number and R_x is a resource in that class as described above, and that once the resource is held, it is held for the duration of the process execution)? Is the sequence C1:R2, C1:R1, C2:R3, C4:R7 allowed? What general statement can be made about the required ordering of resource acquisitions within a process? Is the above ordering of resource classes reasonable to expect of processes? Why or why not?
- (b) (3) Thinking about main memory as a resource is interesting, because much of the time, main memory can be preempted from the process it is allocated to. In most systems, what happens when a process' memory is preempted? Why is preemption not typically an option for other resource types (such as a file or tape drive) for handling deadlock recovery?
- (c) (3) If we want to use the banker's algorithm for handing the swap space, what information must be provided by every process before it can be allocated swap space? What can the operating system do with a running process that requests swap space when the banker's algorithm rejects the allocation as a potential deadlock situation?

7. Memory Management.

- (a) (5) Define spatial locality and temporal locality. If a program has a data space much larger than main memory, how will a lack of spatial locality affect its performance? In other words, will it slow down or speed up, and why. For this question, ignore caching, and consider only main memory management.
- (b) (5) Define both internal and external fragmentation. Draw a pictorial representation of a pure paging system (not paged segmentation, or segmented paging, but just plain paging) which demonstrates internal fragmentation. Draw a pictorial representation of a pure segmentation system which demonstrates external fragmentation.