

# CPS 210 Midterm Exam

Hallowe'en 1995

**Instructions.** Indicate whether each of the numbered statements below is true or false. For some of the statements a simple boolean answer is not sufficient. In these cases, indicate whether the statement is “mostly true”, “mostly false”, or “can’t say”, and explain your answer. For example, if the answer depends on other factors that are not reflected in the statement, then identify them and indicate their effect on the truth of the statement.

Please write your short answers directly on this paper.  
Write your explanations (clearly labeled) on separate sheets.  
This is a closed book examination.  
Collaboration is not permitted. BOO!

1. Variable partitioning reduces fragmentation.
2. Page table formats are architecturally defined.
3. Increasing the page size tends to improve TLB hit rates.
4. Increasing the page size tends to reduce the number of page faults.
5. Mutexes and condition variables can be implemented using only counting semaphores.
6. File caches tend to have higher hit rates than the virtual memory page cache because their replacement policies use better approximations to LRU.
7. Large virtual address spaces make tree-structured page tables (e.g., “linear” page tables as are used on the VAX) less attractive because they increase the space overhead of the tables.
8. Unix heap managers (**malloc** and **free**) operate in runtime library routines without making system calls to the kernel.
9. Suppose Calvin logs into an idle workstation and runs a program that makes repeated sequential passes over a file that is one page larger than the amount of memory allocated to the file cache. Calvin’s program will require fewer disk accesses under LRU file cache replacement than it will with either FIFO or MRU replacement.
10. If the file used by Calvin’s program is one page *smaller* than the file cache, then the program will require fewer disk accesses under LRU than either FIFO or MRU.
11. In Unix systems with virtual memory, a page reference beyond the top of the stack typically grows the stack with a zero-fill page fault.
12. Correct use of condition variables always requires that **wait** calls be enclosed in a **while** loop that tests the condition after each wakeup.
13. Virtual memory is not important for systems with very large physical memories.

14. User-level threads are attractive primarily because they reduce the amount of time spent trapping into kernel mode for thread operations, relative to kernel-level implementations.
15. “Microkernel” operating system structure (e.g., Mach) is attractive primarily because it reduces the amount of time spent trapping into kernel mode for (say) Unix system calls, since most Unix calls are executed by a user-mode server.
16. Imposing a fixed ordering on all resources (e.g., mutexes) requested by threads is a *necessary* condition to prevent deadlock.
17. A fixed resource ordering is *sufficient* to prevent deadlock.
18. A fixed resource ordering is sufficient to prevent starvation.
19. Performance of *multiprocessor* kernels can be improved by using spinlocks (rather than blocking mutexes) for short critical sections, without concern for the problems faced by user-level code that uses spinlocks.
20. Performance of *uniprocessor* kernels can be improved by implementing spinlocks without using interlocked atomic instructions (e.g., **test-and-set**), since those instructions always read/write memory rather than the cache.
21. For processor scheduling, a Shortest Job First policy always delivers the best average response time, but it can only be used when jobs have equivalent priority and their run times are known in advance.
22. Scheduler activations illustrate Lampson’s principle that judicious use of upcalls can yield system interfaces that are simple and general.
23. After a given reference string is processed, the set of pages in resident in physical memory is always a superset of the set of pages that would be resident in a smaller memory -- if but only if a “stack algorithm” is used for page replacement.
24. Page replacement using a two-handed global clock algorithm effectively ensures that physical page frames are allocated fairly among competing processes.
25. During periods of high memory demand, the TLB may contain translations for pages that are no longer resident in physical memory.

*Extra Credit:* The Mean Time Between Failures (MTBF) of a juggling professor is shorter using styrofoam balls than either erasers or magic markers, but only if the temperature in the room is below 68 Fahrenheit.