

Part 1: Operating System

1. (10%) Explain the advantage of virtual-machine concept.
2. Consider a paging system with the page table stored in memory.
 - a. (2%) If a memory reference takes 100 nanoseconds, how long does a paged memory reference take?
 - b. (3%) If we add associative registers, and 90 percent of all page-table references are found in the associative registers, what is the effective memory reference time? (Assume that finding a page-table entry in the associative registers takes zero time, if the entry is there.)
 - c. (5%) Draw a figure to show the translation of a logical address (p, d) to a physical address (f, d) in a paging system with the associative memory.
3. (8%) The multilevel queue scheduling algorithm has been created for situations in which processes are easily classified into different groups. What advantage is there in having different time-quantum sizes on different levels of a multilevel queuing system?
4. (6%) The banker's algorithm is a deadlock-avoidance algorithm. For a system with n processes and m types of resources, what is the order of complexity of the banker's algorithm?
5. (6%) How would use of a RAM disk affect your selection of a disk-scheduling algorithm? What factors would you need to consider?
6. (10%) A variety of disk-organization techniques, collectively called redundant arrays of inexpensive disks (RAID), are commonly used to address the performance and reliability issues. For the following RAID levels and description, make the right bidding between them.

(a) RAID 0	(1) mirrored disks
(b) RAID 1	(2) bit-interleaved Parity
(c) RAID 2	(3) non-redundant striping
(d) RAID 3	(4) block-interleaved parity
(e) RAID 4	(5) block-interleaved distributed parity
(f) RAID 5	(6) memory-style error-correcting codes

Part 2: Computer Organization

7. (10%) Assume that there are two machines:

- M1: A multicycle datapath machine with 500 MHz clock. A Load instruction takes 5 cycles, a Store or R-type instruction takes 4 cycles, and a Branch instruction takes 3 cycles.
- M2: A machine like M1, except that all the cycles that access memory are broken into two clock cycles and the machine runs at a clock rate of 750 MHz.

Using the gcc instruction mixes in that Loads occupied %21, Stores occupied %12, R-type occupied 46%, and Branches occupied 21%. Find out which machine is faster? How much faster is it?

8. (20%) Please briefly explain the following terms: (a) translation-lookaside buffer (b) page fault (c) data forwarding (d) branch prediction

9. (10%) Suppose we have a processor with a base CPI of 1.0, assuming all reference hit in the primary cache, and a clock rate of 250 MHz. Assume a main memory access time of 200 ns, including all the miss handling. Suppose the miss rate per instruction at the primary cache is 5%. How much faster will the machine be if we add a secondary cache that has a 20 ns access time for either a hit or a miss and is large enough to reduce the miss rate to main memory to 2%.

10. (10%) One simple way to model time for logic is to assume each AND or OR gate takes the same time for a signal to pass through it. Time is estimated by simply counting the number of gates along the longest path through a piece of logic. Compare the number of gate delays for the critical paths of two 16-bit adders, one using ripple carry and one using two-level carry lookahead.

