1. (15%, 5% each) Answer each of the following questions.
   a) Memory compaction is one of the solutions to the external fragmentation problem. Is memory compaction possible for each of the following cases?
      (i) Binding of instructions and data to memory addresses is done at compile time.
      (ii) Binding of instructions and data to memory addresses is done at load time
      (iii) Binding of instructions and data to memory addresses is done at execution time.
   b) Why is demand paging more commonly used than demand segmentation?
   c) Why sector sparing may invalidate any optimization by a disk-scheduling algorithm?

2. (10%) Supposed that the LRU page-replacement algorithm is implemented by keeping a stack of page numbers.
   a) If a page with number *k* is referenced without causing a page-fault, what shall we do? (3%)
   b) If a page with number *k* is referenced such that a page-fault occurs and page-replacement is necessary, what shall we do? (3%)
   c) What kind of data structure is best suited for implementing the stack. (4%)

3. (5%) Answer Truth or False:
   a) UNIX provides a file system with general graph directories.
   b) UNIX is a multi-programmed system.
   c) In UNIX, processes are represented by thread structure and user structure.
   d) Linux is a multi-user system and multiple processes are running according to a round-robin scheduler.
   e) The Linux kernel does not allow paging out of kernel memory.

4. (6%) For the processes shown in the box:
   What is the average waiting time using
   a) the Short-Job-First scheduling?
   b) the First-Come-First-Serve scheduling?
   c) the round-robin scheduling with a quantum = 10ms?

| Process | Burst(ms) |
|---|---|
| P1 | 10 |
| P2 | 29 |
| P3 | 3 |
| P4 | 7 |
| P5 | 12 |

5. (4%) Is the resource-allocation graph shown below in a deadlock state? Briefly explain why.

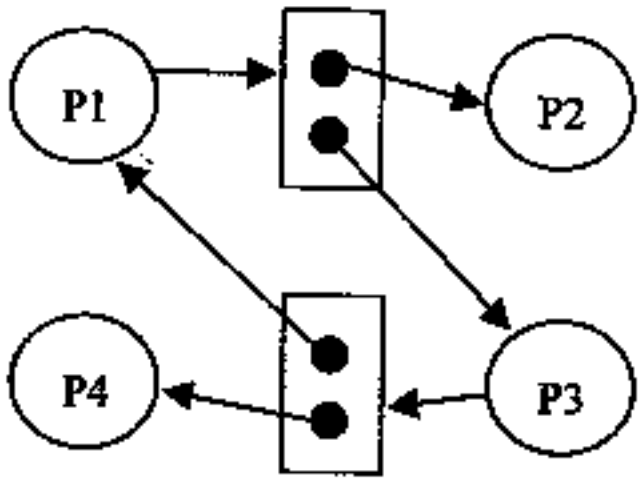


6. (10%) Answer the following two questions:

a) Does the following solution to the critical-section problem satisfy the mutual exclusion, progress and bounded waiting conditions? If not, which one is violated?

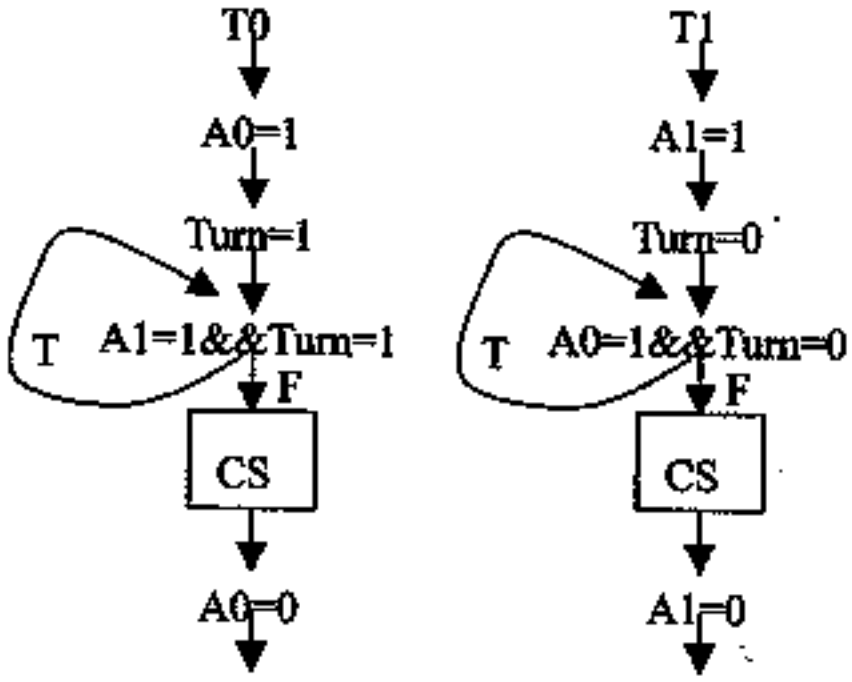


b) For a semaphore *Key*, there are two operations: P for test and V for increment. What will happen if the P() and V() operations on the semaphore *Key* are executed as below:

(i) Key.V();
Critical_section();
Key.P();

(ii) Key.P();
Critical_section();
Key.P();

7. (15%)
   a) What defines an "Instruction Set Architecture (ISA)?" Are cache configuration and number of pipeline stages part of ISA definition?
   b) What constitutes a technically good ISA?
   c) As a computer scientist, what should you care about ISA from each of following points of view? Limit your answer to one short paragraph each.
      (i) Application program development
      (ii) Compiler design
      (iii) Assembly language programming
      (iv) Hardware implementation

8. (10%) Given below is a direct-mapped cache with 1K 32-bit entries and 26-bit physical address.
   a) What are the values of I, J and K?
   b) Use this cache as building blocks; design a 2-way set-associative cache with 2K 32-bit entries and 26-bit physical address.
   c) How many 1Kx32 SRAM chips are needed in your design?

TAG DATA

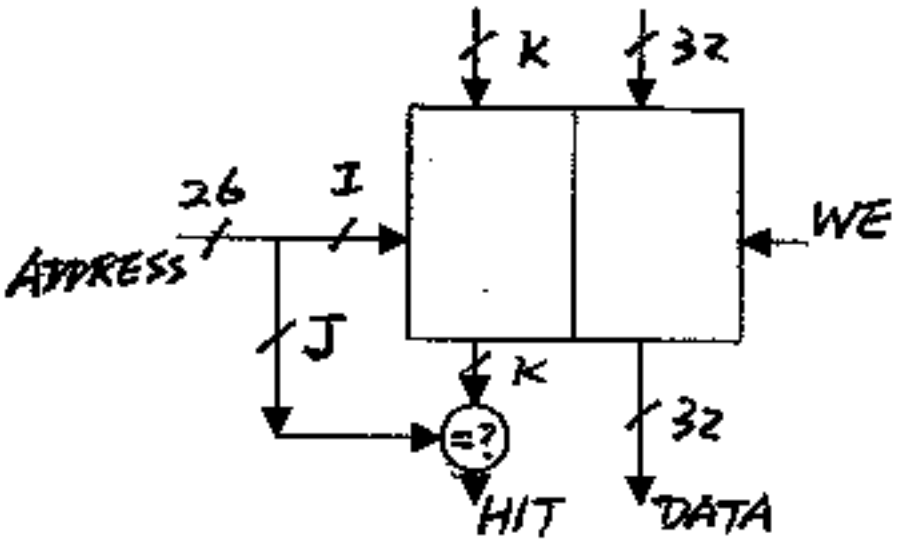


9. (6%) Explain the concepts related to the booth multiplier below.
   a) Describe the algorithm and hardware diagrams for the booth multiplier.
   b) Will the swap of the two operands input to the booth multiplier in executions affect the amount of the workloads (energy consumption) to the booth multiplier? Justify your answers.

10. (3%) Which one of the following techniques is not for reducing power consumption?
    (a) Code Morphing techniques (Transmeta)
    (b) To reduce the transition activities of the hardware components
    (c) To use the clock-gating technique
    (d) Speeding step (Intel Mobile Pentium III)
    (e) To increase the clock rate

11. (3%) Which one of the following descriptions related to memory hierarchies of a computer system is correct?
    (a) The direct-mapped cache scheme is always the best cache scheme for data placement.
    (b) A computer system with virtual memory always out-performs a system without one.
    (c) If a system is designed with a fast bus and fast memory systems, it needs supports with more levels of cache systems.
    (d) The trend of using more and more levels of cache is a result of the imbalance between the performance of CPUs and memories.
    (e) A von-Neumann style architecture greatly reduces the performance bottlenecks of the memory system.

12. (3%) Which one of the followings is not designed to improve the performance of super-scalar and pipeline architectures?
    (a) Delayed branch
    (b) Predicate execution
    (c) Branch prediction
    (d) Multiple issues of instructions per cycle.
    (e) Hazards in instruction scheduling.

13. (10%) Answer the questions related to Amdahl's law below.
    a) Give the key equations for Amdahl's law.
    b) Suppose a company in Taiwan, called PervasiveTech, is building a hardware Java processor. The hardware stack machine can speedup up the bytecode executions. Suppose the hardware execution is 50 times faster than a JVM interpreter. One of the important applications in Java is browsing. In the measurements of a browser application in Java, the researchers find only 60% of the time spent on bytecode executions for stack machines and the rest of them spent on object creations, synchronizations, and garbage collections. A naive hardware implementation of a Java processor with stack machines will only improve the performance of bytecode executions on stack machines, but not the part for object creations, synchronizations, and garbage collections. Under the assumptions above, please use Amdahl's law to explain the performance improvements for the Java processor made by PervasiveTech for a browser application.