1. **Processor pipeline (20 points)**
Consider an out-of-order, dynamically scheduled, superscalar processor with an L1 data cache that supports 1-cycle hits. Assume that 20% of all instructions are loads and that the L1 data cache hit ratio is 90%. The miss-penalty for an L1 miss is 30 cycles. With an ideal memory (all accesses complete in 1-cycle) the IPC of the processor is 2.5.

   (a) (13 points) How many outstanding misses must the processor support to achieve ideal performance? (Assume that instruction level parallelism is not a bottleneck; i.e., that there are enough independent instructions to tolerate miss latency. Also assume that the memory is capable of serving an arbitrary number of memory requests in parallel. Finally, assume that bursty behavior can be ignored.)

   (b) (7 points) What fraction of the ideal performance is achievable if the processor allows exactly one outstanding miss?

2. **Memory Hierarchy and Virtual Memory (20 points)**
(a) (13 points) Consider a system with X bits of page offset in the virtual address. Further assume that Y-bits of physical address are used for the index and block offset for L1 cache access.

   (i) (5 points) Explain how X and Y must be related to ensure that it is trivial to achieve virtually-indexed, physically-tagged cache operation without the synonym problem.

   (ii) (8 points) If page coloring is used to solve the synonym problem, how many colors are needed?

(b) (7 points) What is the key advantage of inverted page tables over page tables that hold an entry per virtual page?

3. **Multicore (25 points)**
(a) (15 points) A multi-threaded shared-memory application is known to be have a specific access pattern called "migratory accesses". In this access pattern, data is read/written by multiple threads, but the reads/writes occur in clusters originating at only one thread at a time. Modify the MSI coherence protocol shown in the figure to capitalize on this behavior. **The modification turns out to be a simplification. Your answer must include a brief justification.**

   Legend: PrRd = Processor Read
   PrWr = Processor Write
   BusRd = Bus read
   BusRdX = Bus read exclusive
   BusWB = Bus writeback

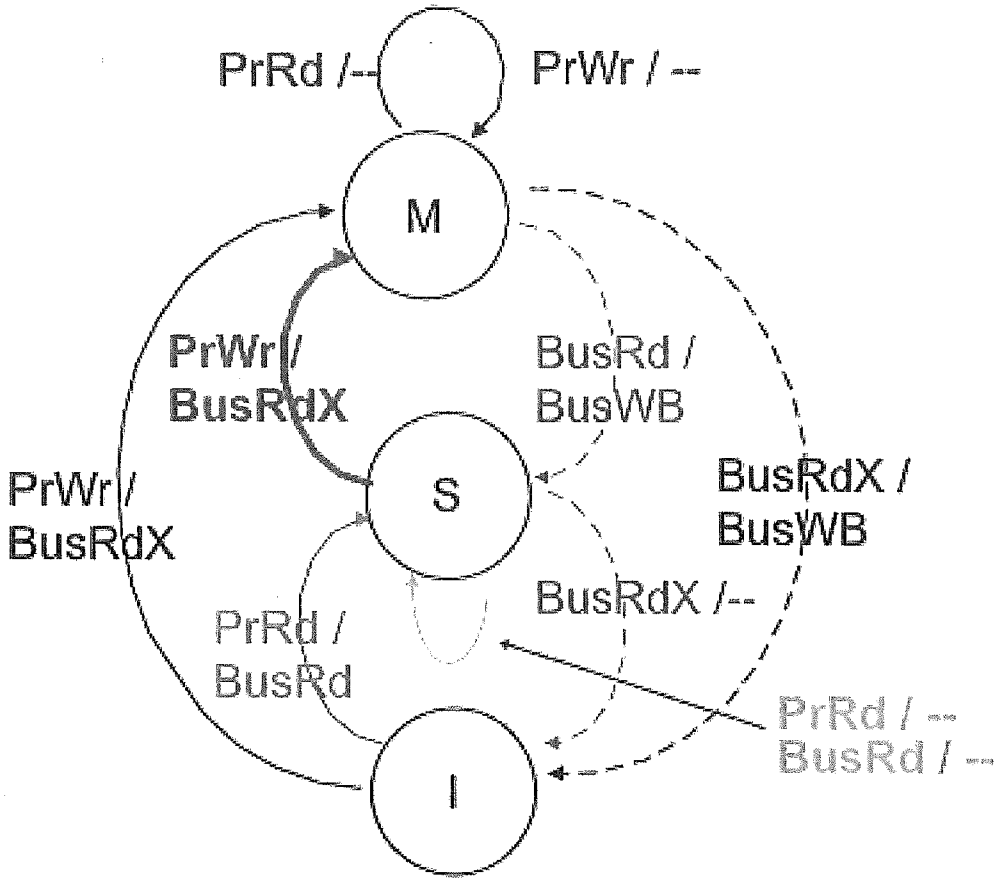

(b) (10 points) Consider the two-threaded, shared memory application shown below. The variables X and Y are both zero (0) before this code segment. The execution results in the printing of 0 and 0. One designer argues that the results are sequentially consistent because there is no consistency violation in printing out the initial values. Is the designer correct? Provide a brief justification for your answer.

Thread 1	Thread 2
X=1 Print Y	Y=1 Print X

4. Fundamentals: (35 points)
The logic technology used in a CPU offers several modes of operation that differ in their energy-performance demands. In the highest-performance operating mode, the CPU achieves P_{peak} performance at E_{peak} energy for a given task. The CPU can achieve **lower performance at a much lower energy cost** by switching to lower performance/energy operating modes.

The software model for the above CPU is a deadline-driven task model wherein the CPU receives a set of tasks with deadlines. The goal is to ensure all tasks

satisfy their deadlines while minimizing energy. The execution times of the tasks vary significantly, but

(1) the execution time is solely determined by the outcomes of a few early branches in the task and **the number of these branches is statically unknown but is no more than eight** (later branches do not affect execution time),

(2) the execution time is highly predictable after the early branches, and

(3) all tasks can complete by the deadline in the highest-performance operating mode.

(a) (10 points) Briefly describe where the energy minimization opportunity is. You do not have to describe how you would capture the opportunity.

(b) (10 points) What is your overall strategy to capture the opportunity? **Hints:** (1) You may assume simple tree-like control flow graphs; and not more complex directed-acyclic graphs (DAGs) or cyclic graphs. (2) Always deciding on the basis of the last branch may lose significant opportunity.

(c) (15 points) Describe the hardware that can exploit the above opportunity. Your solution must include both (a) how you use the information in the structures your design may include, and (b) how you update the information in these structures.