

# COE 501: Computer Architecture

## Problem Set 6: Cache Coherence

- 1) (10 pts) The bus-based dual-core multiprocessor shown in Figure 1 represents a symmetric shared memory architecture. Each processor has an L1 write-back private cache. Coherence is maintained using the MSI write-invalidate snooping protocol. For simplicity, each cache is directly-mapped with four blocks, and each block holds two words (8 bytes). For clarity, the tag contains the full address in hexadecimal, while the data is shown in decimal.

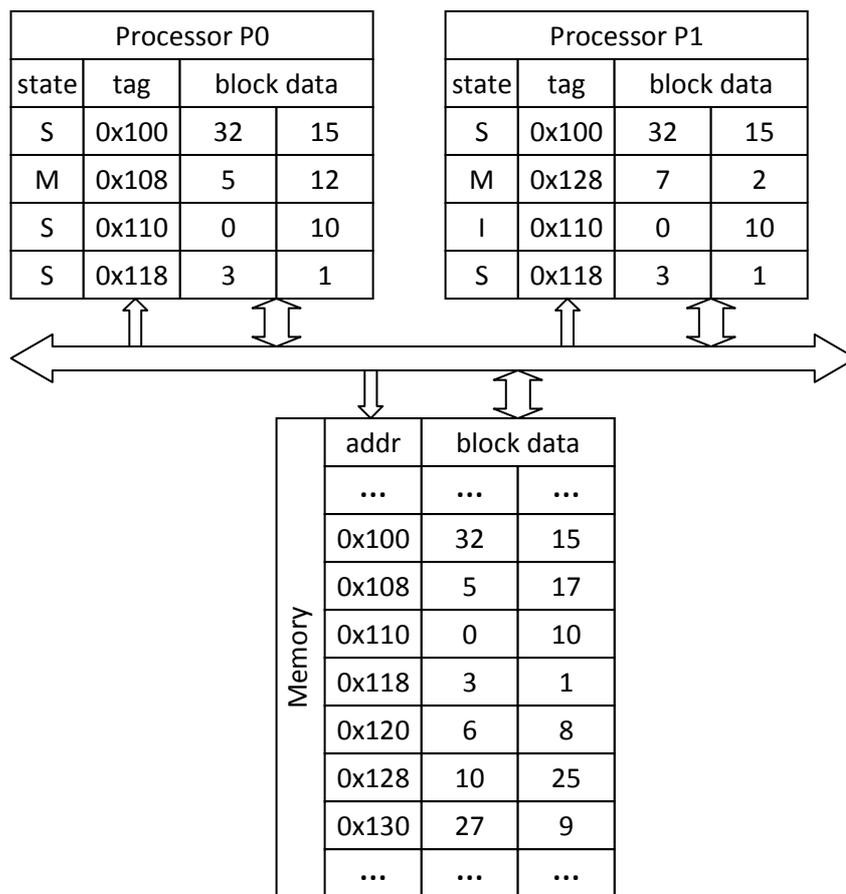


Figure 1: Bus-Based Dual Core multiprocessor

Each part of this exercise specifies a memory operation. Treat each operation as independently applied to the initial state given in Figure 1. What value is returned (by a read) and the bus transaction (if any), the resulting state, tag, and value of the caches and memory after the given operation? Show the content of the relevant cache blocks.

- a) P0 reads address 0x120
- b) P0 writes address 0x120 ← 38
- c) P0 reads address 0x128
- d) P1 reads address 0x114
- e) P1 writes address 0x11C ← 14

2) (10 pts) The performance of a snooping-cache coherent multiprocessor depends on many detailed implementation issues that determine how quickly a cache responds with data in the modified (M) state. In some implementations, a CPU read miss to a block in the M state in another processor's cache is faster than a read miss to a block in memory. This is because caches are smaller and faster than main memory. Consider the following latencies:

- CPU read and write hits generate no stall cycles
- CPU read and write misses generate  $N_{\text{memory}}$  and  $N_{\text{cache}}$  stall cycles if satisfied by memory and cache, respectively.
- CPU write hits that generate an invalidate incur  $N_{\text{invalidate}}$  stall cycles.
- A write-back of a block, due to replacement or another processor's request, incurs an additional  $N_{\text{writeback}}$  stall cycles.

For the multiprocessor of Figure 1, consider the following parameters:

$N_{\text{memory}} = 80$  stall cycles,  $N_{\text{cache}} = 10$  stall cycles,  $N_{\text{invalidate}} = 4$  stall cycles,  $N_{\text{writeback}} = 5$  stall cycles.

Consider the following sequences of operations. Assuming the initial state of Figure 1, how many stall cycles are generated? Explain your answer and show the modified blocks.

a) P1 reads address 0x120  
P1 reads address 0x124  
p1 reads address 0x128

b) P0 reads address 0x100  
P0 writes address 0x108 ← 37  
P0 writes address 0x130 ← 29

c) P0 reads address 0x120  
P0 reads address 0x100  
P1 writes address 0x100 ← 24  
P0 reads address 0x104

3) (10 pts) Many snooping coherence protocols have additional states, state transitions, or bus transactions to reduce the overhead of maintaining cache coherence. An optimization to the MSI snooping protocol is to add the Exclusive (E) state, indicating that no other node has a copy of the block, but the block has not yet been modified. A cache block enters the Exclusive state when a read miss is satisfied by memory and no other node has a valid copy. CPU reads and writes to that block proceed with no further bus traffic, but CPU writes causes the state of the block to transition from Exclusive to Modified. There is a *Shared* wire on the bus that signals '1' when a read miss appears on the bus and at least one other processor has a copy of the block. Otherwise, the *Shared* wire signals '0', indicating that the block is exclusively read by one processor. Draw state transition diagrams for the MESI cache coherence protocol. Two diagrams are needed: first diagram should be based on requests from the processor and the second diagram should be based on requests from the bus.

4) (10 pts) The quad-core multiprocessor shown in Figure 2 uses directory-based cache coherence implemented in the shared L2 cache. Each processor has an L1 write-back private cache. Three states are defined in the L1 cache: M (Modified), S (Shared), and I (invalid). For simplicity, each L1 cache is directly-mapped with four blocks, and each block holds two words (8 bytes). For clarity, the tag contains the full address in hexadecimal, while the data is shown in decimal.

The L2 cache is directly-mapped, shared by all cores, and split into two banks. Blocks whose physical addresses are multiple of 16 (0x10 in hexadecimal) are mapped to L2 cache bank 0, while the other blocks are mapped to L2 cache bank 1. Four presence bits are associated with each L2 cache block, to indicate which L1 cache has a copy of the L2 cache block. The L2 cache is sufficiently large and contains a superset of the L1 blocks (inclusion policy). Four states are associated with each L2 cache block: M (Modified by a processor), O (Owned by L2 and can be shared in L1), S (Shared clean block), and I (Invalid). The L2 Owned state indicates that a block was modified in L1 and later written back to L2. The L2 cache is updated but not memory.

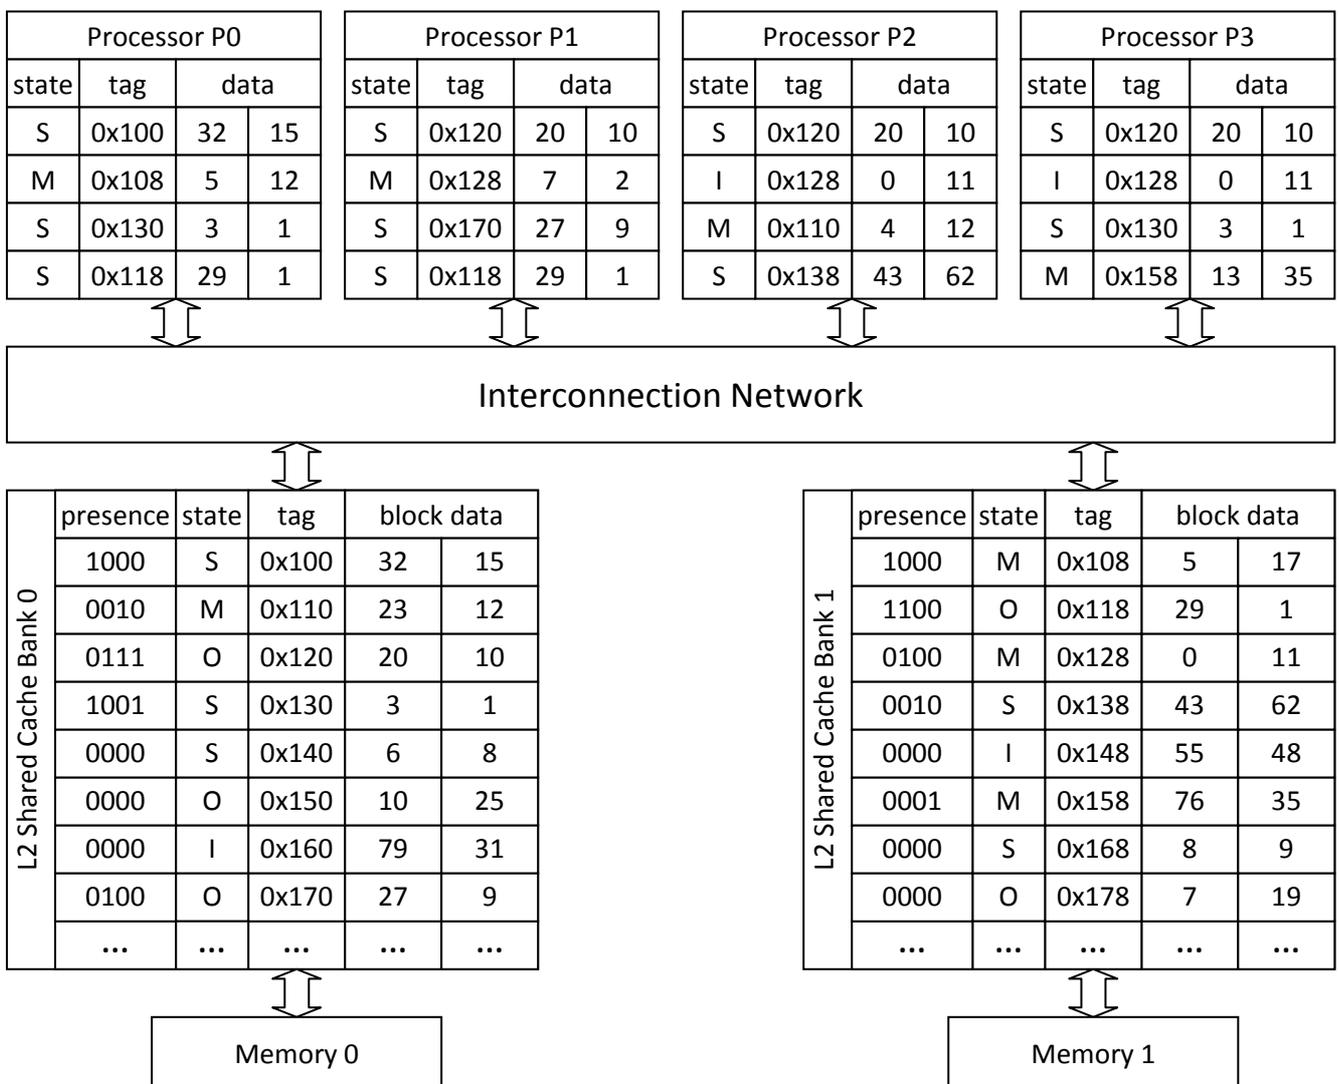


Figure 2: Directory-Based Quad-Core multiprocessor

Each part of this exercise specifies a memory operation. Treat each operation as applied to the initial state given in Figure 2. What value is returned (by a read) and the resulting state, tag, and value in L1 and L2 caches? Show the presence bits and content of the relevant cache blocks.

- a) P0 reads address 0x120
- b) P3 writes address 0x124 ← 38
- c) P2 reads address 0x128
- d) P1 reads address 0x158
- e) P1 writes address 0x11C ← 14

5) (10 pts) Directory protocols are more scalable than snooping protocols because they send explicit request and invalidate messages to those nodes that have copies of a block, while snooping protocols broadcasts all requests to all nodes. Consider again the directory-based quad-core multiprocessor shown in Figure 2. Show all the messages that are sent in the network for each of the following requests:

- a) P0 reads address 0x120
- b) P3 writes address 0x124 ← 38
- c) P2 reads address 0x128
- d) P1 writes address 0x11C ← 14