

5 Comparative Architectures (TMJ)

- (a) A 4KB, blocking, private L1 cache with 16B lines sees the following sequence of accesses from its core.

```
0x00001000 Load
0x00001010 Store
0x00002000 Load
0x00001010 Load
0x00003000 Load
0x00001010 Store
0x00001010 Store
0x00002000 Load
0x00001000 Load
0x00002000 Load
```

Assuming a write-allocate cache that is empty at first and implements the least-recently-used (LRU) replacement algorithm, what is the hit rate if the cache is

- (i) direct-mapped;
- (ii) fully-associative;
- (iii) 2-way set-associative?

[6 marks]

- (b) If the core supports out-of-order execution, how might a non-blocking cache bring performance benefits? [4 marks]

- (c) How might the core's load/store queue be used to reduce the number of memory accesses seen by the cache? [4 marks]

- (d) Assume that this core and cache are part of a chip-multiprocessor, with the cache connected to a shared L2 via a bus that maintains coherence through a snooping MESI protocol. What sequence of steps would be taken if another core wanted to load from 0x00001010 after the given sequence had finished?

[6 marks]