

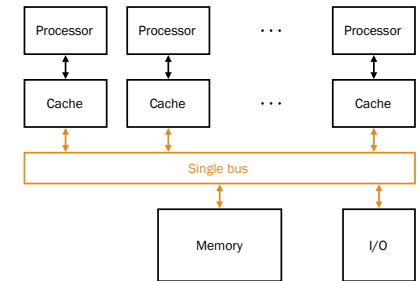
Directory-Based Cache Coherence

CSE 240B

Dean Tullsen

Review – Bus-based MP Architectures

- Centralized Memory
- Uniform Memory Access Time
- Bus as point of serialization and broadcast




CSE 240B

Dean Tullsen

Review – Network-based Multiprocessors

- “Scalable” network interconnect
- No broadcast medium
- Serialization at memory.



© 2003 Elsevier Science (USA). All rights reserved.

CSE 240B

Dean Tullsen

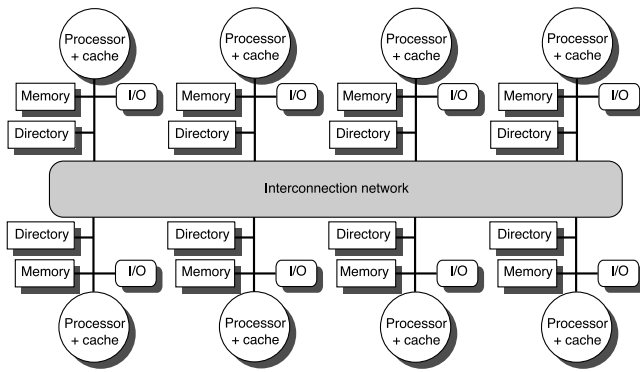
Large-scale MPs

- Separate memory per processor
- No broadcast mechanism
- Use a directory, per memory block, that tracks the state of every block in every cache (why?)
- Accounting per cache block, or per memory block?
 - Pro (per memory) – less complex – one entry per block
 - Minus – memory much bigger than caches
- Distribute directories with memories to prevent single directory as bottleneck

CSE 240B

Dean Tullsen

Distributed Directory MP



Example (simple!) Directory Protocol

- Similar to snoopy: each block could be:
 - Shared
 - Uncached
 - Exclusive
- In addition to state, must keep track of who is sharing (when shared), or who owns (when exclusive).
- Terminology
 - Local node (where the request [load] originates)
 - Home node (where the memory location permanently resides)
 - Remote node (a node that has a copy of the block)

Directory Messages

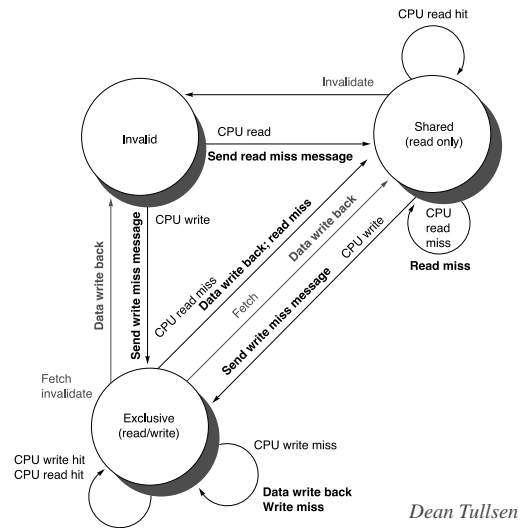
Type	Source	Destination
Read miss	Local	Home
Write miss	Local	Home
Invalidate	Home	Remote
Fetch	Home	Remote
Fetch/invalidate	Home	Remote
Data value	Home	Local
Data write back	Remote	Home

Example Directory Protocol

- Uncached, shared, exclusive
- If uncached
 - Memory is valid
 - Can only receive read miss, write miss messages
- Shared
 - Memory value is valid
 - Can receive read miss, write miss
- Exclusive
 - Memory value is invalid
 - Read miss, write miss, writeback

Cache Block Actions

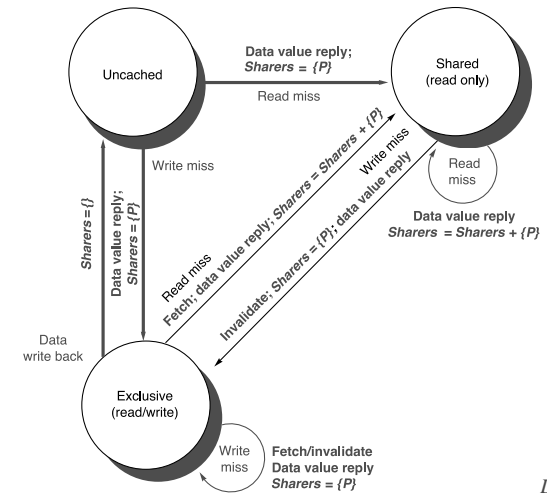
- Cache must also track state (ESI)



CSE 240B

Dean Tullsen

Directory Actions



CSE 240B

Dean Tullsen

Example

P1	P2	Directory	Memory
----	----	-----------	--------

P1: write 10 to A1

P1: read A1

P2: read A1

P2: write 20 to A1

P2: write 40 to A2
(same block)

CSE 240B

Dean Tullsen

Coherence Summary

- Snoopy vs. Directory
- Directory keeps track of state of line and of all sharers
- Snoopy cost is function of cache size
- Directory cost is function of memory size (* number of nodes?)
- Coherence misses can be significant part of miss rate

CSE 240B

Dean Tullsen