


Lecture 18: Coherence Protocols

- Topics: coherence protocols for symmetric and distributed shared-memory multiprocessors (Sections 4.2-4.4)

SMP/UMA/Centralized Memory Multiprocessor



SMPs

- Centralized main memory and many caches → many copies of the same data
- A system is cache coherent if a read returns the most recently written value for that word

Time	Event	Value of X in	Cache-A	Cache-B	Memory
0			-	-	1
1	CPU-A reads X		1	-	1
2	CPU-B reads X		1	1	1
3	CPU-A stores 0 in X		0	1	0

Cache Coherence

A memory system is coherent if:


- P writes to X; no other processor writes to X; P reads X and receives the value previously written by P
- P1 writes to X; no other processor writes to X; sufficient time elapses; P2 reads X and receives value written by P1
- Two writes to the same location by two processors are seen in the same order by all processors – write serialization
- The memory *consistency* model defines “time elapsed” before the effect of a processor is seen by others

Cache Coherence Protocols

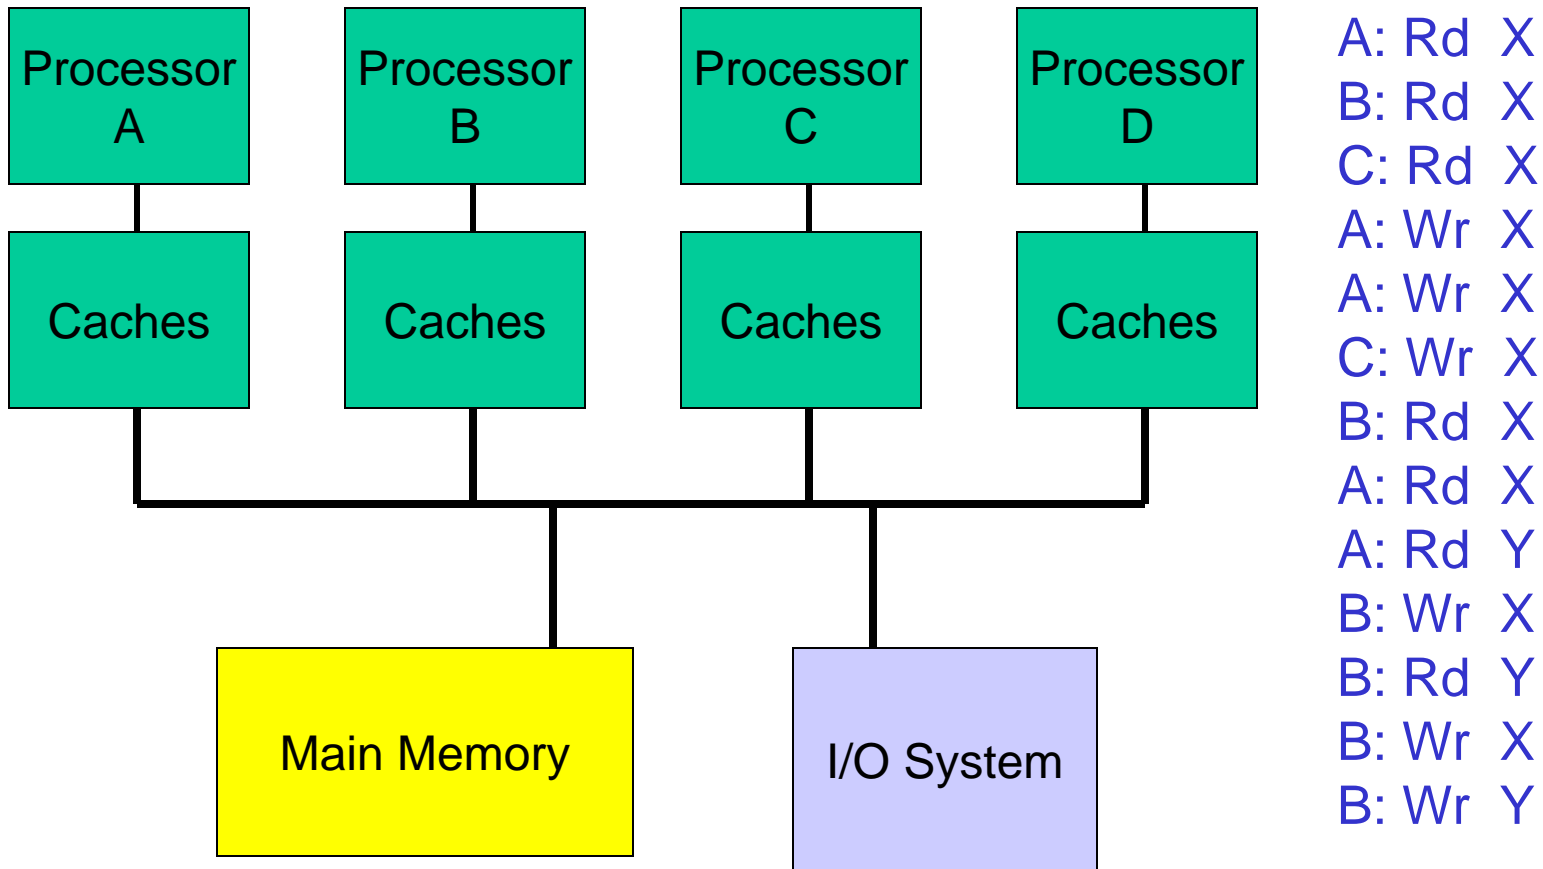
- Directory-based: A single location (directory) keeps track of the sharing status of a block of memory
- Snooping: Every cache block is accompanied by the sharing status of that block – all cache controllers monitor the shared bus so they can update the sharing status of the block, if necessary
 - Write-invalidate: a processor gains exclusive access of a block before writing by invalidating all other copies
 - Write-update: when a processor writes, it updates other shared copies of that block

Design Issues

- Invalidate
- Find data
- Writeback / writethrough
- Cache block states
- Contention for tags
- Enforcing write serialization



SMP Example



SMP Example

	A	B	C
A: Rd X			
B: Rd X			
C: Rd X			
A: Wr X			
A: Wr X			
C: Wr X			
B: Rd X			
A: Rd X			
A: Rd Y			
B: Wr X			
B: Rd Y			
B: Wr X			
B: Wr Y			

SMP Example

	A	B	C
A: Rd X	S		
B: Rd X	S	S	
C: Rd X	S	S	S
A: Wr X	E	I	I
A: Wr X	E	I	I
C: Wr X	I	I	E
B: Rd X	I	S	S
A: Rd X	S	S	S
A: Rd Y	S (Y)	S (X)	S (X)
B: Wr X	S (Y)	E (X)	I
B: Rd Y	S (Y)	S (Y)	I
B: Wr X	S (Y)	E (X)	I
B: Wr Y	I	E (Y)	I

Example Protocol

Request	Source	Block state	Action
Read hit	Proc	Shared/excl	Read data in cache
Read miss	Proc	Invalid	Place read miss on bus
Read miss	Proc	Shared	Conflict miss: place read miss on bus
Read miss	Proc	Exclusive	Conflict miss: write back block, place read miss on bus
Write hit	Proc	Exclusive	Write data in cache
Write hit	Proc	Shared	Place write miss on bus
Write miss	Proc	Invalid	Place write miss on bus
Write miss	Proc	Shared	Conflict miss: place write miss on bus
Write miss	Proc	Exclusive	Conflict miss: write back, place write miss on bus
Read miss	Bus	Shared	No action; allow memory to respond
Read miss	Bus	Exclusive	Place block on bus; change to shared
Write miss	Bus	Shared	Invalidate block
Write miss	Bus	Exclusive	Write back block; change to invalid

Coherence Protocols

- Two conditions for cache coherence:
 - write propagation
 - write serialization

- Cache coherence protocols:
 - snooping
 - directory-based

 - write-update
 - write-invalidate

Performance Improvements

- What determines performance on a multiprocessor:
 - What fraction of the program is parallelizable?
 - How does memory hierarchy performance change?
- New form of cache miss: coherence miss – such a miss would not have happened if another processor did not write to the same cache line
- False coherence miss: the second processor writes to a different word in the same cache line – this miss would not have happened if the line size equaled one word

How do Cache Misses Scale?

	Compulsory	Capacity	Conflict	Coherence	
				True	False
Increasing cache capacity					
Increasing processor count					
Increasing block size					
Increasing associativity					

Simplifying Assumptions

- All transactions on a read or write are atomic – on a write miss, the miss is sent on the bus, a block is fetched from memory/remote cache, and the block is marked exclusive
- Potential problem if the actions are non-atomic: P1 sends a write miss on the bus, P2 sends a write miss on the bus: since the block is still invalid in P1, P2 does not realize that it should write after receiving the block from P1 – instead, it receives the block from memory
- Most problems are fixable by keeping track of more state: for example, don't acquire the bus unless all outstanding transactions for the block have completed


Coherence in Distributed Memory Multiprocs

- Distributed memory systems are typically larger → bus-based snooping may not work well
- Option 1: software-based mechanisms – message-passing systems or software-controlled cache coherence
- Option 2: hardware-based mechanisms – directory-based cache coherence

Directory-Based Cache Coherence

- The physical memory is distributed among all processors
- The directory is also distributed along with the corresponding memory
- The physical address is enough to determine the location of memory
- The (many) processing nodes are connected with a scalable interconnect (not a bus) – hence, messages are no longer broadcast, but routed from sender to receiver – since the processing nodes can no longer snoop, the directory keeps track of sharing state

Distributed Memory Multiprocessors



Cache Block States

- What are the different states a block of memory can have within the directory?
- Note that we need information for each cache so that invalidate messages can be sent
- The block state is also stored in the cache for efficiency
- The directory now serves as the arbitrator: if multiple write attempts happen simultaneously, the directory determines the ordering

Title

- Bullet