



Terminology

Coherence

- Defines what values can be returned by a read
- Coherent if:
 - If P writes to X then reads X, with no writes to X by other processors should return value written by P
 - If P writes to X and then another processor reads from X, if read/write sufficiently separated should return value written by original P
 - Writes to the same location are serialized; two writes to the same location by any two processors are seen in the same order by all processors
- Consistency
 - Determines when a written value will be returned by a read, we'll need to define a memory consistency model



Protocols for Coherency				
Write Invalidate: When one processor writes, invalidate all copies of this data that may be in other caches				
Processor activity	Bus activity	Contents of CPU A's cache	Contents of CPU B's cache	Contents of memory location X
				0
CPU A reads X	Cache miss for X	0		0
CPU B reads X	Cache miss for X	0	0	0
CPU A writes a 1 to X	Invalidation for X	1		0
	G 1 1 A 14	1	1	1

	Contents of	Contonts of	Contents of memory
copies that may be in other caches			
while Distances when one process		oudeust inte	and apaare an

Processor activity	Bus activity	Contents of CPU A's cache	Contents of CPU B's cache	Contents of memory location X
				0
CPU A reads X	Cache miss for X	0		0
CPU B reads X	Cache miss for X	0	0	0
CPU A writes a 1 to X	Write broadcast of X	1	1	1
CPU B reads X		1	1	1
				J



Implementing Invalidation

- Bus-based scheme
 - Processor to invalidate acquires the bus
 - Broadcasts the address to invalidate
 - All other processors continuously snoop on the bus watching the addresses
 - If an address is invalidated that matches an address in its cache, then the corresponding data is invalidated
 - Serialization of the bus forces serialization of access automatically

7

8

Implementing Invalidation

- Write-through cache
 - To locate a data item when a cache miss occurs, just go to memory (since memory will contain the most up-to-date value in a write-through cache)
- Write-back cache
 - What problem do we have reading in data on a cache miss when all processors use write-back caches?

Implementing Invalidation

- Write-Back Cache
 - May need to find the most recent value of a data item in some other processor's cache, not in memory
 - We can do this using the same snooping scheme for cache misses and writes
 - Each processor snoops every address placed on the bus when a read is requested from memory
 - If a processor has a dirty copy of the requested cache block (i.e. one we wrote to and is hence updated), provide that cache block to the requestor and abort the memory access

9

• Since write-back caches generate lower memory requirements, they are preferred in multiprocessors despite increased complexity









Performance of Snooping Coherence Protocols

- Use the four parallel programs described earlier as a benchmark
 - Split cache misses into two sets
 - Coherence Misses misses due to cache invalidation
 - Capacity Misses actually capacity, compulsory and conflict misses, but most of these are capacity. "Normal" cache misses from a uniprocessor





Distributed Shared Memory Architectures

- Snooping protocol not so efficient on most DSM machines
 - Snooping requires a broadcast mechanism, which is easy to do on a bus
 - Most DSM systems don't have a bus but a more complex system interconnect (mesh, hypercube, etc.) so broadcast becomes a much more expensive operation
- One solution:
 - Prevent coherency by marking shared data as uncacheable
 - Private data can still be cached
 - For shared data, we must always access through memory
 - Simple to implement, but can slow things down if programs are not written with this model in mind
 - Access to remote memory can be quite slow



- Another solution: software-based coherency
 - Possible but slow and conservative, every block that might be shared treated as if it is shared
- Most popular alternative: Directory Protocol
 - Directory keeps the state of every block that may be cached
 - Information in the directory includes which caches have copies of the block, whether it is dirty, shared, etc.
 - Centralized version of snooping this directory is always in the same location
 - Memory requirements for the directory are proportional to the number of memory blocks * number of processors





Directory Protocol Terminology

- Local node
 - Node where a request originates
- Home node
 - Node where the memory location and directory entry reside
 - Could be the local node as well
- Remote node
 - Node that has a copy of a cache block
 - Might be exclusive or shared
- Nodes will pass messages to one another; messages will move a directory between states in a transition diagram, just like with the snooping protocol

Message type	Source	Destination	Message contents	Function of this message
Read miss	Local cache	Home directory	Р, А	Processor P has a read miss at address A; request data and make P a read sharer.
Write miss	Local cache	Home directory	Р, А	Processor P has a write miss at address A; — request data and make P the exclusive owner
Invalidate	Home directory	Remote caches	А	Invalidate a shared copy of data at address A
Fetch	Home directory	Remote cache	А	Fetch the block at address A and send it to its home directory; change the state of A in the remote cache to shared.
Fetch/invalidate	Home directory	Remote cache	А	Fetch the block at address A and send it to its home directory; invalidate the block in the cache.
Data value reply	Home directory	Local cache	Data	Return a data value from the home memory.
Data write back	Remote cache	Home directory	A, Data	Write back a data value for address A.

- 3-5 : Home sends to remote cache when home needs to satisfy request
- 6 : Home sends requested data to local cache
- 7 : Block replaced, needs to be written back to home or fetch requested

11









Summary Coherence protocols may be needed for correct program behavior Most common protocol is write-back cache, write invalidation Can use snooping or directory based mechanism to implement coherence Coherence requests become more important in programs that are less optimized Optimized programs will access most data locally and have fewer requests Exactly how the cache miss rates affect CPU performance depends on the memory system, interconnect, latency,

bandwidth, etc.