

Quiz 2 Review

Ryan Lee

(Adapted from prior course offerings)

Quiz 2 logistics

- Time: 1pm EST on Friday, November 17
- Location: 32-141
- Please read handout before quiz

Topics


- Advanced memory operations
- Multithreading
- Cache coherence
 - Snooping-based vs. Directory-based
 - VI, MSI, MESI, MOSI, ...
 - Transient states
 - Synchronization primitives
- On-chip Networks
 - Topology
 - Routing
 - Flow control
 - Router micro-architecture
- Memory consistency model
 - Sequential consistency
 - Total Store Order (TSO)
 - Relaxed consistency

Advanced memory operations

- Write policy
 - Hits: write through vs. write back
 - Misses: write allocate vs. write no allocate
- Speculative loads/stores
 - Cause 1: control dependency: All instructions are speculative until commit
 - Just like other instructions
 - Solution: buffer the stores and commit them in order
 - Cause 2: (memory-location-based) data dependency
 - Simple solution: buffer stores; loads search addresses of all previous stores
 - Problem: addresses of previous stores may be unknown
 - Solution: speculate no data dependency
 - Use a data structure to keep track of this speculation: speculative load buffer

Store Buffer

- » Enables data forwarding
- » Handles OoO stores
- » Handles speculative stores



» On store execute:

- mark valid and speculative; save tag, data and instruction number.

» On store commit:

- clear speculative bit and eventually move data to cache

» On store abort:

- clear valid bit

» One entry per store

» Written by stores


» Searched by loads

» Writes to data cache

Load Buffer

- » On load execute:
 - mark entry valid, and instruction number and tag of data.
- » On load commit:
 - clear valid bit
- » On load abort:
 - clear valid bit

*Speculative
Load Buffer*



- » One entry per load
- » Written by loads
- » Searched by stores

- » Enables aggressive load scheduling
- » Detects ordering violations

Multithreading

- Fine-grain multithreading
- Coarse-grain multithreading
- Simultaneous multithreading
 - Scheduling policies
 - Round-robin: Equalize *throughput* between threads
 - ICOUNT: Equalize *instr. in flight* between threads

Cache coherence

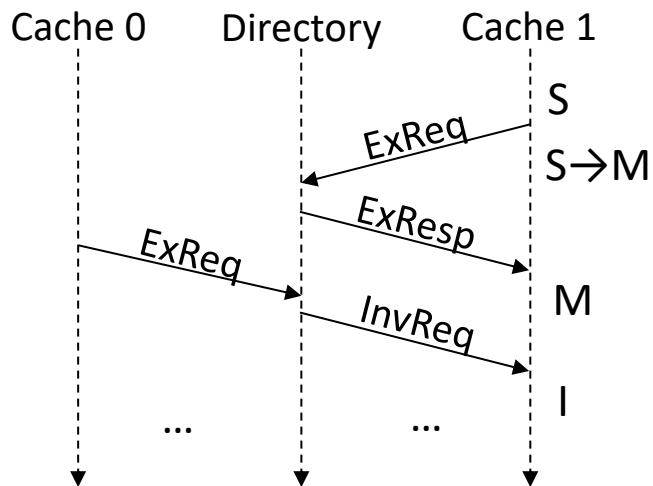
- Simplify building shared memory systems
- Definition:
 - Write propagation **Liveness: do something good**
 - Writes eventually become visible to all processors
 - Write serialization **Safety: don't do anything bad**
 - Writes to the same location are serialized (all processors see them in the same order)

Cache coherence

- Transient states: required by lack of atomicity
 - Two types
 - Split states: to implement one transaction
 - E.g., S transitions to $S \rightarrow M$ (instead of M), waiting for an ExResp
 - Race states: to handle overlaps of two transactions
 - Not all such overlaps require transient states
 - See the following examples


Cache coherence

- Split example
 - $S \rightarrow M$



Cache coherence

- Race example



If the arriving message is from a younger transaction:

- Either defers processing it
- Or handles it immediately and transitions to a race state (e.g., $S \rightarrow I \rightarrow M$)

On-chip networks

- Allow sharing communication resource
- Topology
 - Metrics: routing distance, diameter, average distance, bisection bandwidth, ...
- Routing
 - Properties: deterministic, adaptive, deadlock-free, ...

On-chip networks

- Flow control
 - Bufferless
 - Circuit switching, dropping, misrouting, ...
 - Buffered
 - Store-and-forward, virtual cut-through, wormhole, virtual channel
- Router architecture

Memory (consistency) model

- Concerns reads/writes to multiple memory locations
- Interacts with many parts and optimizations of the system
 - Probably more than what you would have imagined...
- Coherence is an useful (but not necessary) building block
 - Recall: Coherence guarantees writes are visible in some global order.

Sequential consistency

- Definition
 - “The result of any execution is the same as if the operations of all the processors were executed in some sequential order, and the operations of each individual processor appear in the order specified by the program”
 - Arbitrary order-preserving interleaving of memory references of sequential programs
- Implementation
 - In-order instruction execution + atomic loads and stores
- Advantage: easy to understand
- Disadvantage: limits performance
 - Uniprocessor optimizations often violate them!
 - E.g., committed store buffers, non-blocking caches, speculative execution, memory address speculation, ...

Total Store Order (TSO)

- Allows loads to go ahead of stores waiting in the store buffer
- Implementation
 - Sequential consistency implementation + per-core FIFO store buffer with store-load bypassing

Relaxed memory consistency

- Allows more reordering
 - Store-load
 - Store-store
 - Load-load
 - Load-store
- Re-ordering can be disabled by fences/barriers

Tips on consistency problems

- Keep definitions in mind
- Think systematically
 - E.g., For questions asking all allowed execution results: search invariants to minimize brute-force search
 - E.g., For questions asking to add minimal barriers/fences: find the precise reordering that violates the target model

Wish you all the best!