# Lecture 09: Hardware Memory Consistency

#### Vivek Kumar Computer Science and Engineering IIIT Delhi vivekk@iiitd.ac.in

CSE513: Parallel Runtimes for Modern Processors

## Last Lecture (Recap)

- Memory latency continues to limit the performance of multicore processors
  - Several optimizations inside processors for hiding the load/store latency
    - As a side effect of these optimizations, load/store inside a program could be reordered, and hence may not happen in the source code order as expected by programmer
- Memory consistency model defines a set of rules for valid set of reordering of two different memory accesses
  - Both compiler and processor can perform reordering
- Sequential consistency is the most primitive form of memory consistency that basically says memory access to any location always happens atomically, and the effect is visible to each and every core
  - Modern programming languages supports sequential consistency only for code block within a mutex lock/unlock operation (Data Race Free)







## **Today's Class**

- x86-TSO memory consistency model
- Store buffer



#### **Few Alternatives to Sequential Consistency**



#### Let's take a deep dive in x86-TSO model



• Rule-1&2: Read-Read & Write-Write reordering is NOT allowed (over same core)



Question: What is an invalid set of result for R1 & R2? Assume initially A=B=0

Answer: R1=1 & R2= 0

• **Rule-3**: Writes are NOT reordered with earlier Reads (over same core)



Question: What is an invalid set of result for R1 & R2? Assume initially A=B=0

Answer: R1=1 & R2= 1



• **Rule-4a**: Reads CANNOT be reordered with earlier Writes to **SAME** memory locations (over same core)



Question: What is an invalid set of result for R1 & R2? Assume initially A=B=0

Answer: R1=0 & R2= 0



• **Rule-4b**: Reads CAN be reordered with earlier Writes to **DIFFERENT** memory locations (over same core)



Question: What is an invalid set of result for R1 & R2? Assume initially A=B=0

Answer: None!



Core-1	Core-2	
<b>WR1</b> : A=1	<b>WR2</b> : B=1	
<b>RD1a</b> : R1=A	<b>RD2a</b> : R3=B	
<b>RD1b</b> : R2=B	<b>RD2b</b> : R4=A	

- Can we change the orderings below?
  - Core-1: WR1->RD1a->RD1b
  - Core-2: WR2->RD2a->RD2b
  - o Not possible due to the Rules 1 and 4
- It could happen that R2 = R4 = 0
  - o But why?



#### Store Buffer [edit]

A store buffer is used when writing to an invalid cache line. Since the write will proceed anyway, the CPU issues a read-invalid message (hence the cache line in question and all other CPUs' cache lines that store that memory address are invalidated) and then pushes the write into the store buffer, to be executed when the cache line finally arrives in the cache.

A direct consequence of the store buffer's existence is that when a CPU commits a write, that write is not immediately written in the cache. Therefore, whenever a CPU needs to read a cache line, it first has to scan its own store buffer for the existence of the same line, as there is a possibility that the same line was written by the same CPU before but hasn't yet been written in the cache (the preceding write is still waiting in the store buffer). Note that while a CPU can read its own previous writes in its store buffer, other CPUs *cannot see those writes* before they are flushed from the store buffer to the cache - a CPU cannot scan the store buffer of other CPUs.

From Wikipedia

MESI protocol. (2022, May 2). In Wikipedia. https://en.wikipedia.org/wiki/MESI\_protocol

CSE513: Parallel Runtimes for Modern Processors

#### **Coherence without Store Buffer**





CSE513: Parallel Runtimes for Modern Processors



CSE513: Parallel Runtimes for Modern Processors

- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



- Store writes in a local buffer and then proceed to next instruction immediately
   Absorbs writes faster than the next cache => prevents stalls
- The cache will pull writes out of the store buffer when it's ready
  - Aggregates writes to same cache line => reduces cache traffic



CSE513: Parallel Runtimes for Modern Processors Credits: James Bornholt, UW, CSE451





	Thread 1	Thread 2	
	(1) $A = 1$ (2) $rQ = B$	(3) $B = 1$	
	$(2) \mathbf{I} \mathbf{U} - \mathbf{D}$		
	Can <b>r0 = 0</b> a	nd r1 = 0?	
1			



Thread 1	Thread 2
<ul><li>(1)</li><li>(2) rØ = B</li></ul>	(3) $B = 1$ (4) $r1 = A$
Can r0 = 0 a	and <b>r1 = 0?</b>





Thread 1	Thread 2
<ul> <li>(1)</li> <li>(2) r0 = B</li> </ul>	(3) $B = 1$ (4) $r1 = A$
Can r0 = 0 a	and r1 = 0?





Thread 1	Thread 2
(1) (2) <b>rØ = B</b>	(3) (4) <b>r1 = A</b>
Can <b>r0 = 0</b> an	dr1 = 0?





Thread 1	Thread 2
(1) (2) $\mathbf{r0} = \mathbf{B}$	(3) (4) $r1 = A$
Can <b>r0 = 0</b> a	and r1 = 0?





Thread	d 1 Thread	2
(1)	(3)	
(2)	(4) $r1 = A$	
Can <b>r0 =</b>	0 and r1 = 0?	Executed
		r0 = B (= 0)





	Thread 2	Thread 1
	(3)	(1)
	(4)	(2)
Executed	nd r1 = 0?	Can <b>r0 = 0</b> ar
r0 = B (= 0)		
r1 = A (= 0)		





Thread 1	Thread 2	
(1)	(3)	
(2)	(4)	
Can <b>r0 = 0</b> a	nd r1 = 0?	– Executed
		r0 = B (= 0)
		r1 = A (= 0)
		A = 1





• **Rule-6**: Writes that are causally related, appear to all processors to occur in an order consistent with the causality relation

Core-1	Core-2	Core-3	
WR1: A=1			
	<b>RD2</b> : R1=A	<b>RD3a</b> : R2=B	
	<b>WR2</b> : B=1	<b>RD3b</b> : R3=A	

• If R1 = R2 = 1, can R3 = 0

- Implies WR1 → RD2
- RD2 → WR2 (Rule-3)
- RD3a  $\rightarrow$  RD3b (Rule-1)
- R2 = 1
  - Implies WR1→RD2→WR2→RD3a
  - Hence, R3 cannot be 0
    - R3 = 1

• **Rule-7**: Any two Writes must appear to execute in the same order to all cores other than those performing the writes

Core-1	Core-2	Core-3	Core-4
<b>WR1</b> : A=1	<b>WR2</b> : B=1	<b>RD3a</b> : R1=A	<b>RD4a</b> : R3=B
		<b>RD3b</b> : R2=B	<b>RD4b</b> : R4=A



Rule-8: All cores agree on a single execution order of locked instructions
 I.E., Rule-7 hold true even with LOCK statements

Core-1	Core-2	Core-3 RD3a: R1=A RD3b: R2=B	Core-4 RD4a: R3=B RD4b: R4=A	<ul> <li>Is it possible         <ul> <li>R1=1, R2=0, R3=1, R4</li> </ul> </li> <li>R1=1 &amp; R2=0         <ul> <li>Implies WR1 → WR2 w.r.t. Core-3</li> </ul> </li> </ul>
Same (Rule #	e example as in pr ‡7) except the LOC	evious slide CK instruction Same exar slide (Rule even with	<ul> <li>RD4b: R4=A</li> <li>Implies WR1 → WR2 w.r.t. C</li> <li>R3=1, R4=0 implies:</li> <li>Either RD3b → RD3a</li> <li>Violates Rule-1</li> <li>Or, WR2 → WR1 w.r.t. Core</li> <li>Recall value propagation of cache coherence</li> <li>If A=1 is propagated before B=1 then A=1 have been propagated</li> </ul>	
CSE513: I	Parallel Runtimes for M	odern Processors	©ν	/ivek Kumar 34

Rule-9: Reads and Writes are NOT reordered with Locked instructions



## **Reference Materials**

- Intel processor manual
  - o <u>https://cdrdv2.intel.com/v1/dl/getContent/671200</u>
    - Section 8.2
- Book
  - Sorin et al., A Primer on Memory Consistency and Cache Coherence
  - <u>https://course.ece.cmu.edu/~ece847c/S15/lib/exe/fetch.php?media</u>
     <u>=part2\_2\_sorin12.pdf</u>



## Next Lecture (L #10)

- Language memory consistency model
- Next lecture tomorrow at 2.30pm

