

## x86-TSO


Shared-mem concurrency.

Weak memory complexity.

Clarifies model for x86.

Not just proofs.

## Sequential consistency (SC)




Atomic actions: Mem read  $(a) \rightarrow \underline{\underline{v}}$   
Mem write  $(a, v) \rightarrow \checkmark$

- SC used in real HW.

★ - Clean abstraction for devs.

# Example



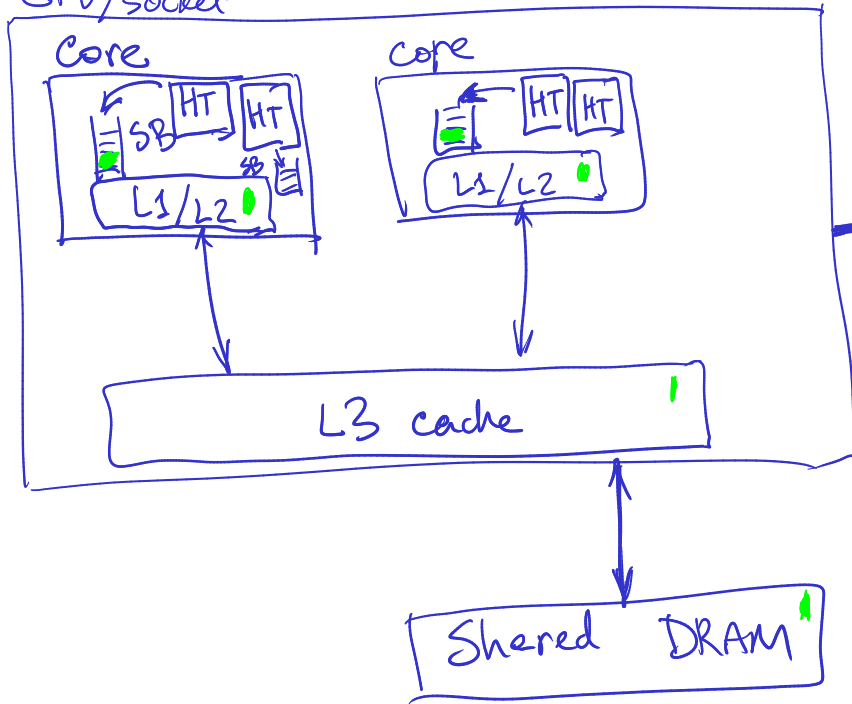
main:  
 $x=0$   
 $y=0$   
Spawn T1, T2  
wait  
print ( $x'$ ,  $y'$ ).

Results:

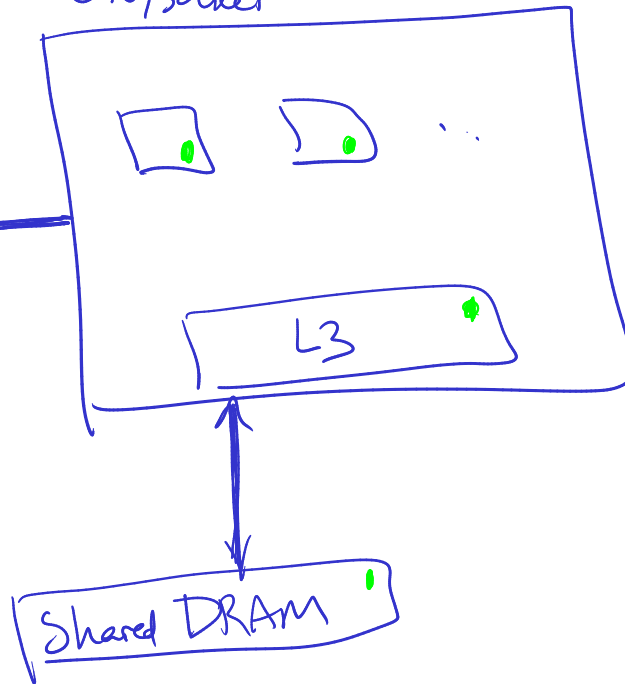
|              |   |                   |   |                       |
|--------------|---|-------------------|---|-----------------------|
| OK<br>in SC. | } | (1, 1)            | } | OK<br>in<br>weak mem. |
|              |   | (0, 1)            |   |                       |
|              |   | (1, 0)            |   |                       |
|              |   | <del>(0, 0)</del> |   |                       |

# Real HW for shared mem.

CPU/socket



CPU/socket



## Out-of-order exec

- | store(x, 1)
- | t := fetch(y)
- store(y', t)

## Speculative exec.

|              |                    |
|--------------|--------------------|
| <u>T1</u>    | <u>T2</u>          |
| *ptr = x     | p = shared         |
| shared = ptr | print( <u>*p</u> ) |

```

if (a) {
    ... ←
} else {
    ... ←
}

```

# Indep. reads of indep. writes (IRIW)

T1

$x := 1 \rightarrow SB$

T2

$y := 1 \rightarrow SB$

T3

$x_3 := x \leftarrow 1 \leftarrow \text{from SB}$

$y_3 := y \leftarrow 0 \leftarrow \text{mem}$

T4

$y_4 := y \leftarrow 1 \leftarrow SB$

$x_4 := x \leftarrow 0 \leftarrow \text{mem}$

|     |       |     |
|-----|-------|-----|
| 0,0 | ————— | 0,0 |
| 1,1 | ————— | 1,1 |
| 1,0 | ————— | 1,0 |
| 0,1 | ————— | 0,1 |

|     |       |     |
|-----|-------|-----|
| 1,0 | — ? — | 0,1 |
|-----|-------|-----|

not on x86 (not on x86-TSO).

## How to program on weak mem?

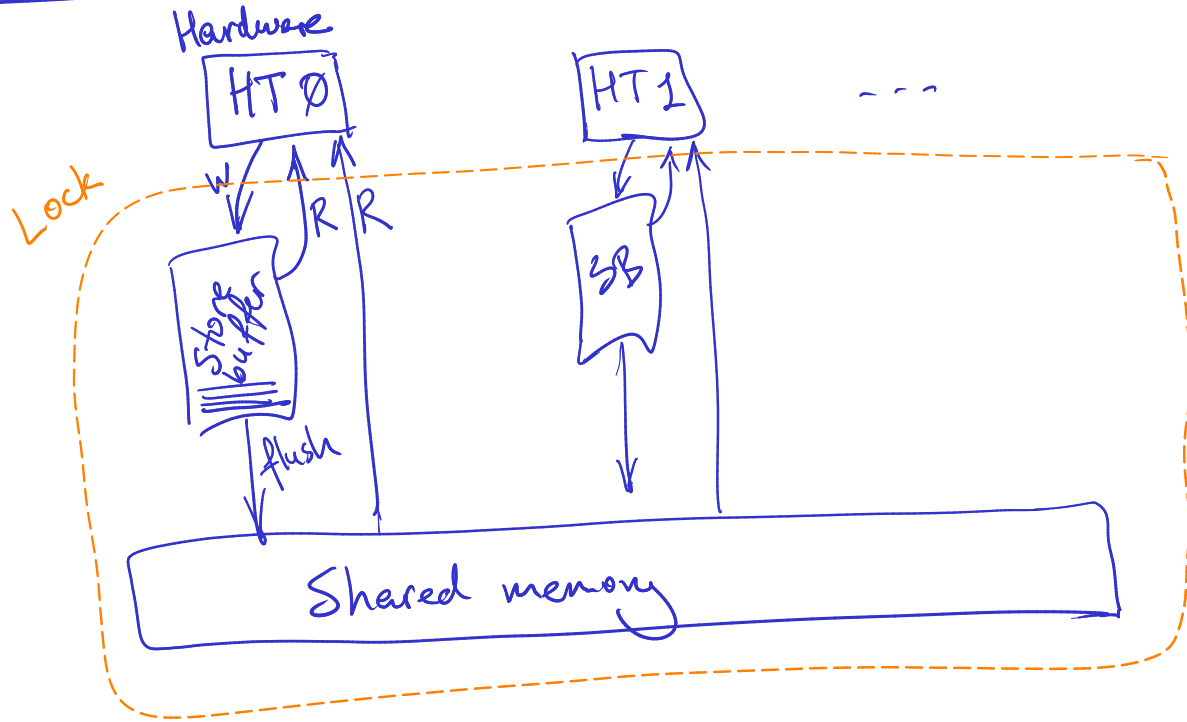
Simulate SC: locks, etc.

General: lock impl?

lock-free: Linux RCU

x86-TSO.

x86-TSO model : Abstract machine.



T1  
 $x = 1 \rightarrow SB$   
 $y' = y \leftarrow mem$

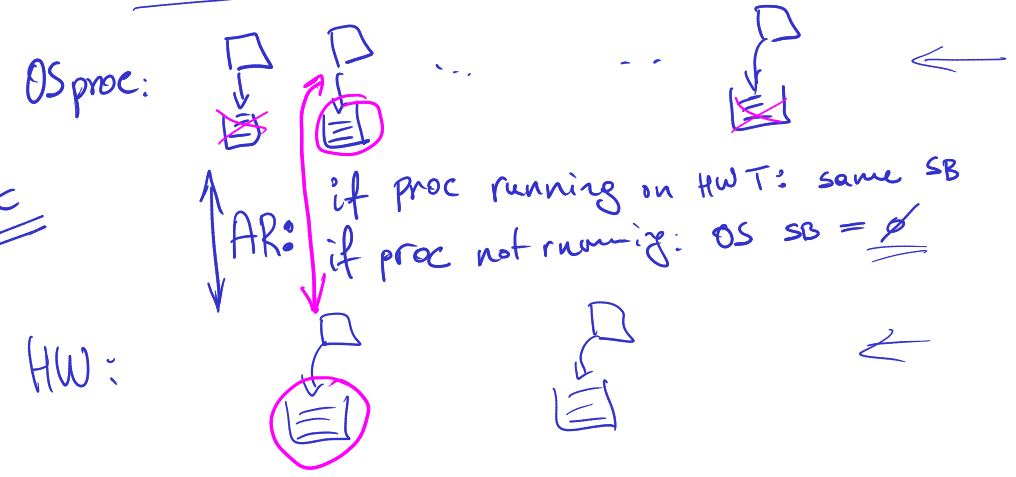
T2  
 $y = 1 \rightarrow SB$   
 $x' = x \leftarrow mem$

0,0: not in SC.

$\frac{x = 1 \rightarrow SB}{y = x \leftarrow \emptyset}$

Atomic steps: R from SB or mem  
★ W to SB  
Flush from SB: fence  
→ ↻ ←  
Locked op: + flush SB.

Breakout Q: x86-TSO for OS thread?



# Intel/AMD specs

Principles: "causality" —


Litmus tests. —

Imprecise (prose)

Hard to understand.

Incomplete  $\Leftarrow$

"Axiomatic": predicates on traces.



|   |                                    |
|---|------------------------------------|
| On entry the address of spinlock is in register EAX and the spinlock is unlocked iff its value is 1 |                                    |
| acquire: LOCK;DEC [EAX]   | ; LOCK'd decrement of [EAX]        |
| JNS enter   | ; branch if [EAX] was $\geq 1$     |
| spin: CMP [EAX],0   | ; test [EAX]                       |
| JLE spin  | ; branch if [EAX] was $\leq 0$     |
| JMP acquire   | ; try again                        |
| enter:  | ; the critical section starts here |
| release: MOV [EAX] $\leftarrow$ 1   |                                    |

Locked?

release in mem or later?