

Correctness of concurrent objects on weak memory models

Graeme Smith

The University of Queensland

With various collaborators:

John Derrick, Sheffield, UK

Brijesh Dongol, Brunel, UK

Lindsay Groves, VUW, NZ

Kirsten Winter, UQ

Rob Colvin, UQ

Concurrent objects

- ▶ Data structures, etc., designed to be accessed simultaneously by multiple threads
- ▶ Synchronisation — coarse-grained, fine-grained, non-blocking

Goal: prove correct wrt abstract specification where all operations are atomic.

Example: reader-writer object for a pair of variables x_1 and x_2

```
int x1:=0, x2:=0;
```

```
write(int d1, d2) { x1, x2 := d1, d2; }
```

```
read() { return (x1, x2); }
```

Concurrent objects

Example: reader-writer object for a pair of variables x_1 and x_2

- ▶ can be implemented using a Linux seqlock

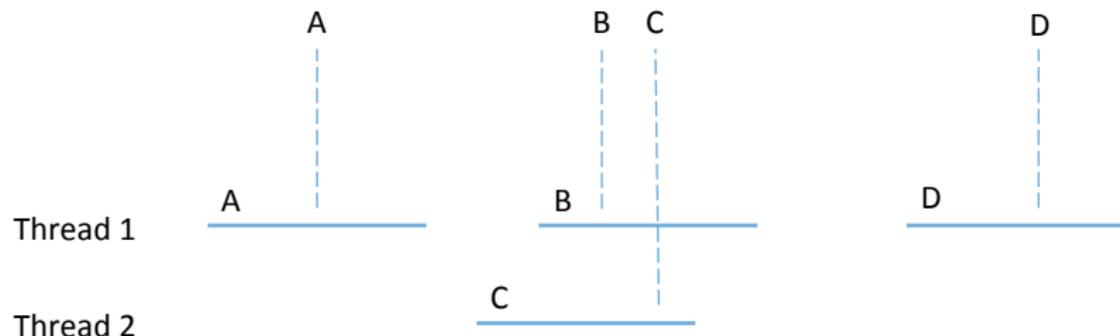
```
int x1:=0,x2:=0,c:=0;

write(int d1,d2){
    acquire;
    c++;
    x1:=d1;
    x2:=d2;
    c++;
    release;
}

read(){
    int c0,d1,d2;
    do{
        do{
            c0:=c
        }while(c0%2!=0)
        d1:=x1;
        d2:=x2;
    }while(c!=c0)
    return(d1,d2);
}
```

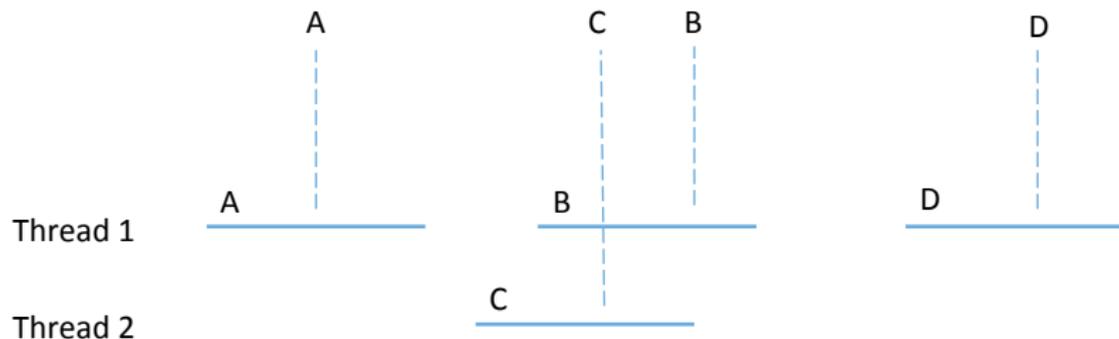
Linearizability

- ▶ Standard correctness criterion for concurrent objects
- ▶ Proposed by Herlihy and Wing (1989)
- ▶ Compositional
- ▶ Relation between **sequential** specification histories and **concurrent** implementation histories



Linearizability

- ▶ Standard correctness criterion for concurrent objects
- ▶ Proposed by Herlihy and Wing (1989)
- ▶ Compositional
- ▶ Relation between **sequential** specification histories and **concurrent** implementation histories



Linearizability

- ▶ A **history** is a finite sequence of invocations and responses, e.g.,

$$h = \langle \text{inv}((p, \text{write}), (1, 2)), \text{inv}((q, \text{read}), \perp), \text{resp}((p, \text{write}), \perp), \text{resp}((q, \text{read}), (1, 2)) \rangle$$

- ▶ Assuming operations are distinct, we define $<_h$ such that

$$op_1 <_h op_2 \hat{=} \exists m, n, in, out \bullet m < n \wedge h(m) = \text{resp}(op_1, out) \wedge h(n) = \text{inv}(op_2, in)$$

- ▶ A history is **completed** when, for each invocation without a response, we
 - ▶ add a response at the end of the history
 - ▶ remove the invocation
- ▶ A object C is linearizable wrt a specification A iff for each history h of C there is a history hs of A such that
 - ▶ hs is a permutation of a completion of h
 - ▶ $<_h \subseteq <_{hs}$

Linearizability vs trace refinement

- ▶ Trace refinement
 - ▶ Back (1990)
 - ▶ Abadi and Lamport (1991)
- ▶ Linearizability is weaker than trace refinement
 - ▶ Only **partial correctness**, e.g., `read` and `write` could be implemented as `while(true) {}`
 - ▶ Only **finite histories** considered, e.g., *C* which outputs *true* an infinite number of times, linearizes with *A* which outputs *true* any finite number of times
 - ▶ **Overlapping operations can be reordered**

```
int a:=0; // shared variable
op1() {
    if (a=0) return 0;
    return 2;
}
op2() {
    a:=1;
    return 1;
}
```

Linearizability vs trace refinement

- ▶ **Linearizability = observational refinement**
 - ▶ Filipović, O'Hearn, Rinetzsky and Yang (2010)
 - ▶ no program calling an object can distinguish its behaviour from the specification

The program

```
x:=op1() || y:=op2()
```

is equivalent to

```
r1:=op1(); x:=r1 || r2:=op2(); y:=r2
```

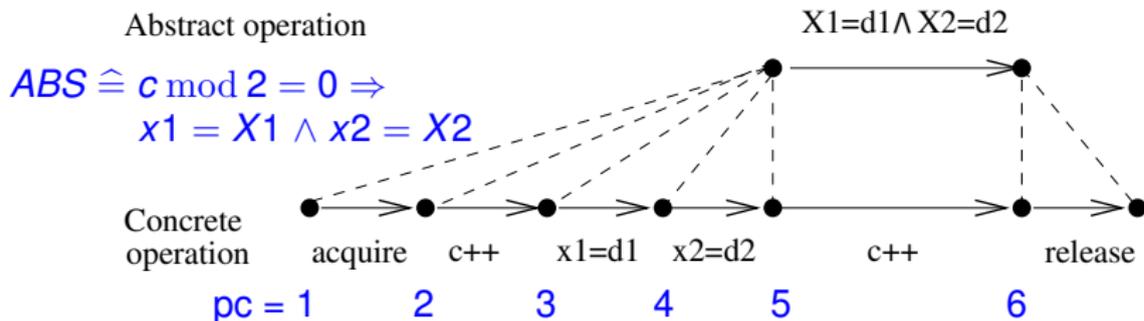
whose behaviour `r2:=op2(); r1:=op1(); x:=r1; y:=r2`
will result in observation of $x = 2$ before $y = 1$.

```
int a:=0; // shared variable
```

```
op1() {  
    if (a=0) return 0;  
    return 2;  
}
```

```
op2() {  
    a:=1;  
    return 1;  
}
```

Simulation-based proof method



- ▶ thread-local — proofs consider 1 thread in isolation
- ▶ step-local — proofs consider 1 program step in isolation

$$INV \hat{=} (pc = 5 \Rightarrow c \bmod 2 \neq 0 \wedge x1 = d1 \wedge x2 = d2) \wedge \dots$$
$$(pc = 4 \Rightarrow c \bmod 2 \neq 0 \wedge x1 = d1) \wedge \dots$$

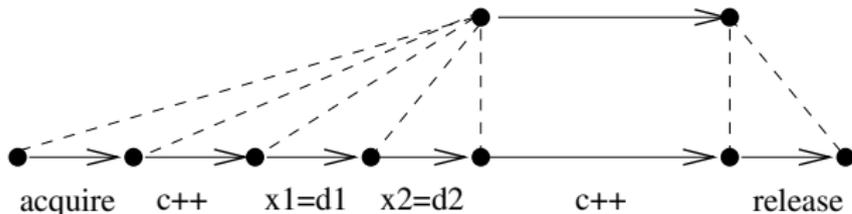
$$D \hat{=} (pc \in 2..6 \Rightarrow pc_q \notin 2..6) \wedge \dots$$

Simulation rules

Abstract operation

$X1=d1 \wedge X2=d2$

Concrete operation



- ▶ non-linearization point

$$\forall a, c, c' \cdot ABS \wedge INV \wedge COp \Rightarrow ABS' \wedge INV'$$

- ▶ linearization point

$$\forall a, a', c, c' \cdot ABS \wedge INV \wedge COp \wedge AOp \Rightarrow ABS' \wedge INV'$$

- ▶ non-interference

$$\forall a, c, c', c_q \cdot ABS \wedge INV \wedge INV_q \wedge D \wedge COp \Rightarrow INV'_q \wedge D'$$

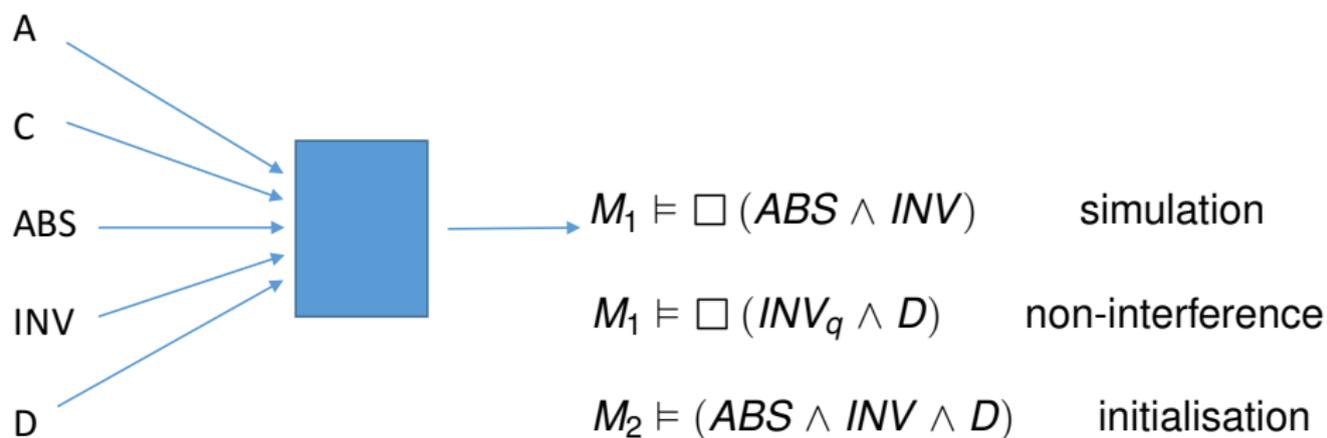
- ▶ initialisation

$$\forall a, c, c_q \cdot CInit \wedge CInit_q \wedge AInit \Rightarrow ABS \wedge INV \wedge D$$

Tool support

- ▶ theorem proving (KIV, Isabelle)
- ▶ model checking (TLC, NuSMV)
- ▶ A number of approaches for model checking linearizability exist
 - ▶ check scenarios consisting of a finite number of threads calling a finite number of operations
 - ▶ number of threads limited to between 2 and 6
 - ▶ number of operations calls can be as limited as 2 or 3
- ▶ Our approach — thread-local and step-local
 - ▶ arbitrary number of threads
 - ▶ (potentially) arbitrary number of operation calls

Model checking approach



where $M1$ is

$$Init \hat{=} \neg stop \wedge ABS \wedge INV \wedge INV_q$$

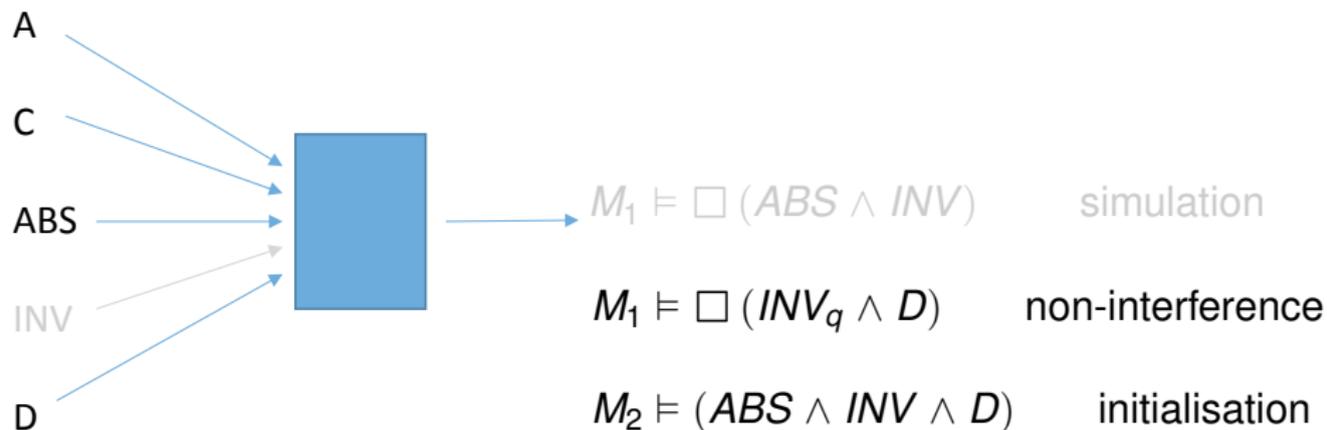
$$Next \hat{=} (COp_1 \vee COp_2 \vee \dots \vee COp_n) \wedge stop'$$

and $M2$ is

$$Init \hat{=} \neg stop \wedge AInit \wedge CInit \wedge CInit_q$$

$$Next \hat{=} (COp_1 \vee COp_2 \vee \dots \vee COp_n) \wedge stop'$$

Model checking approach



where $M1$ is

$$Init \hat{=} \neg stop \wedge ABS \wedge INV \wedge INV_q$$

$$Next \hat{=} (COp_1 \vee COp_2 \vee \dots \vee COp_n) \wedge stop'$$

and $M2$ is

$$Init \hat{=} \neg stop \wedge Ainit \wedge Cinit \wedge Cinit_q$$

$$Next \hat{=} (COp_1 \vee COp_2 \vee \dots \vee COp_n) \wedge stop'$$

Weak memory models

- ▶ High-level programming languages
 - ▶ Java
 - ▶ C/C++
 - ▶ to allow compiler optimisations
- ▶ Multicore architectures
 - ▶ TSO — used by Intel, AMD
 - ▶ Power — used by IBM
 - ▶ ARM — used in most mobile devices
 - ▶ to improve run-time efficiency
- ▶ Allow program instructions to appear as if they occurred out of order

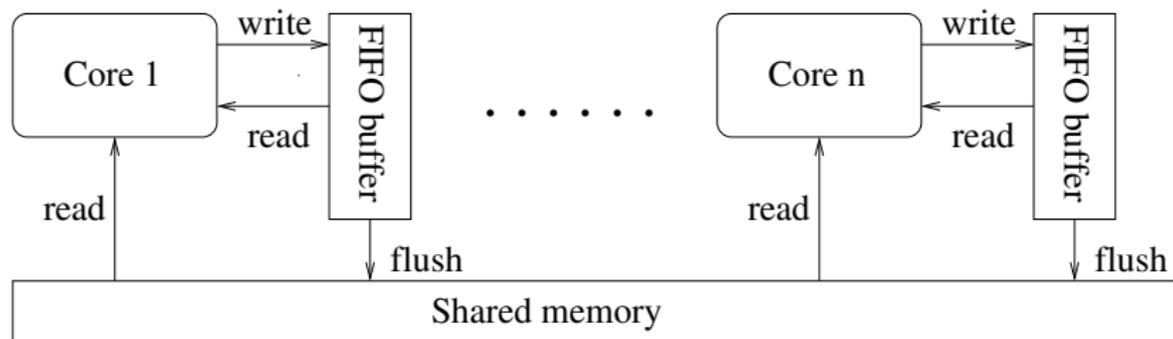
Power and ARM

- ▶ Observable behaviour can be derived via systematic testing (litmus tests)
 - ▶ Maranget, Sarkar, Sewell (2012)

Thread 1	Thread 2
$x := 1$	$r1 := y$
$y := 1$	$r2 := x$
Initial state: $x = 0 \wedge y = 0$	
Allowed: $r1 = 1 \wedge r2 = 0$	

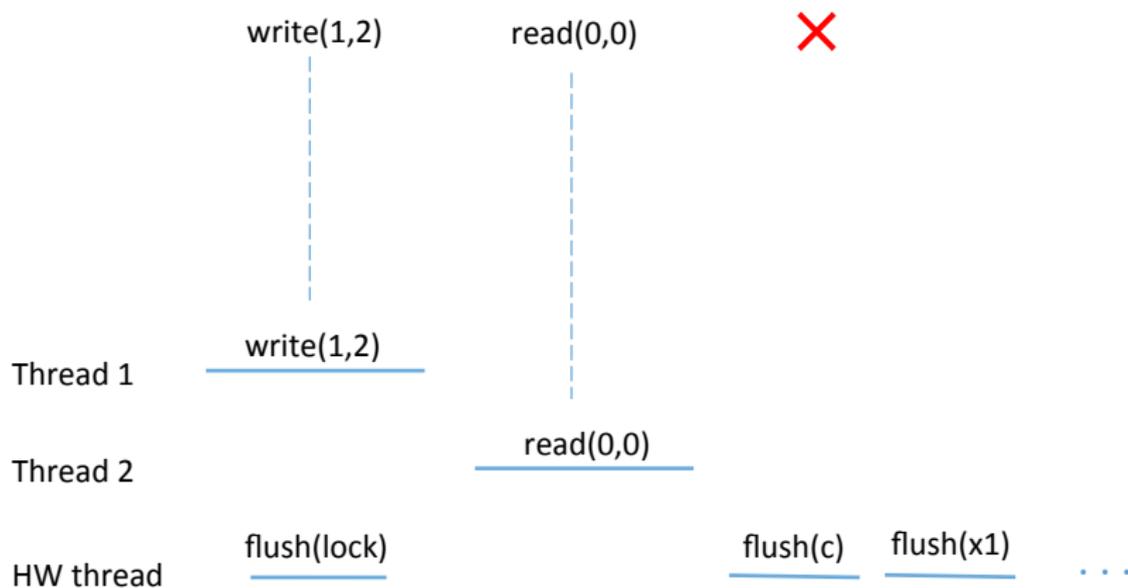
For example, the above result occurred 1.7G/167G times on Power 7 architecture (1%), and 61K/552M times on ARM APQ8060 (0.01%).

TSO (Total Store Order)



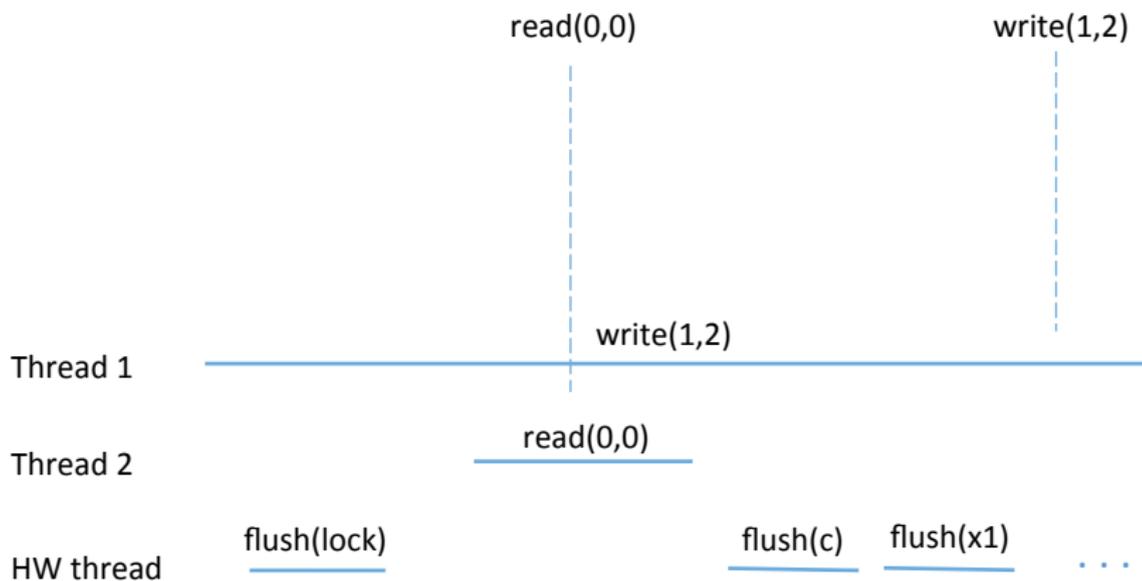
Linearizability

When threads run independently, a client program cannot see relative order of events in different threads



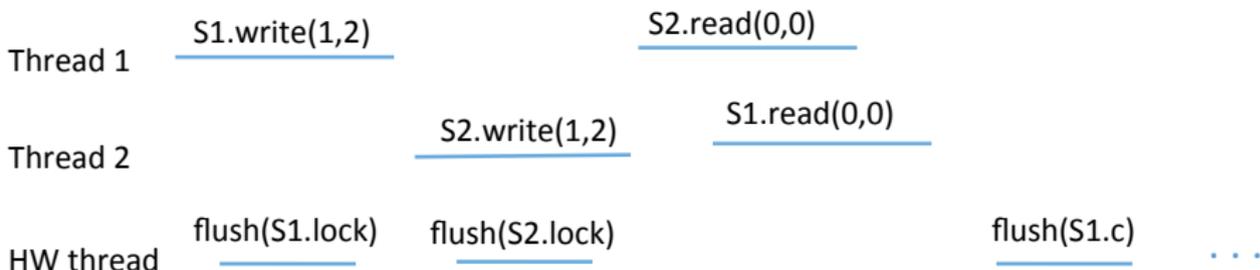
TSO-linearizability

An operation remains active (and hence may take affect) up to the time of the last flush of a value written by the operation



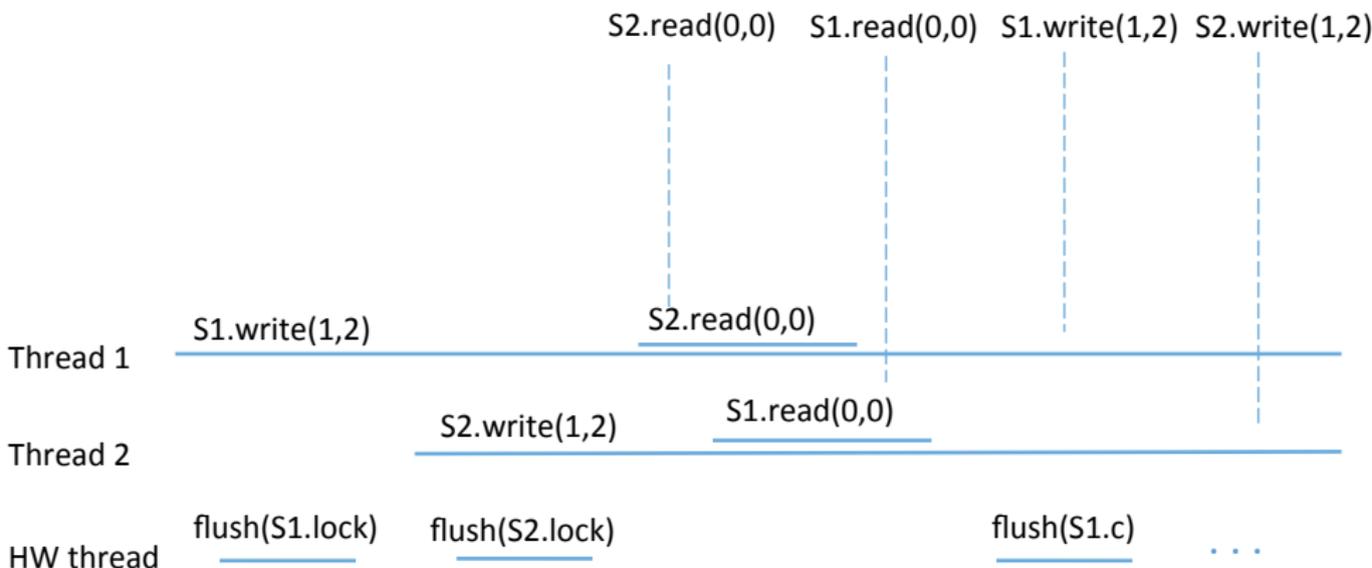
TSO-linearizability

An operation remains active (and hence may take affect) up to the time of the last flush of a value written by the operation



TSO-linearizability

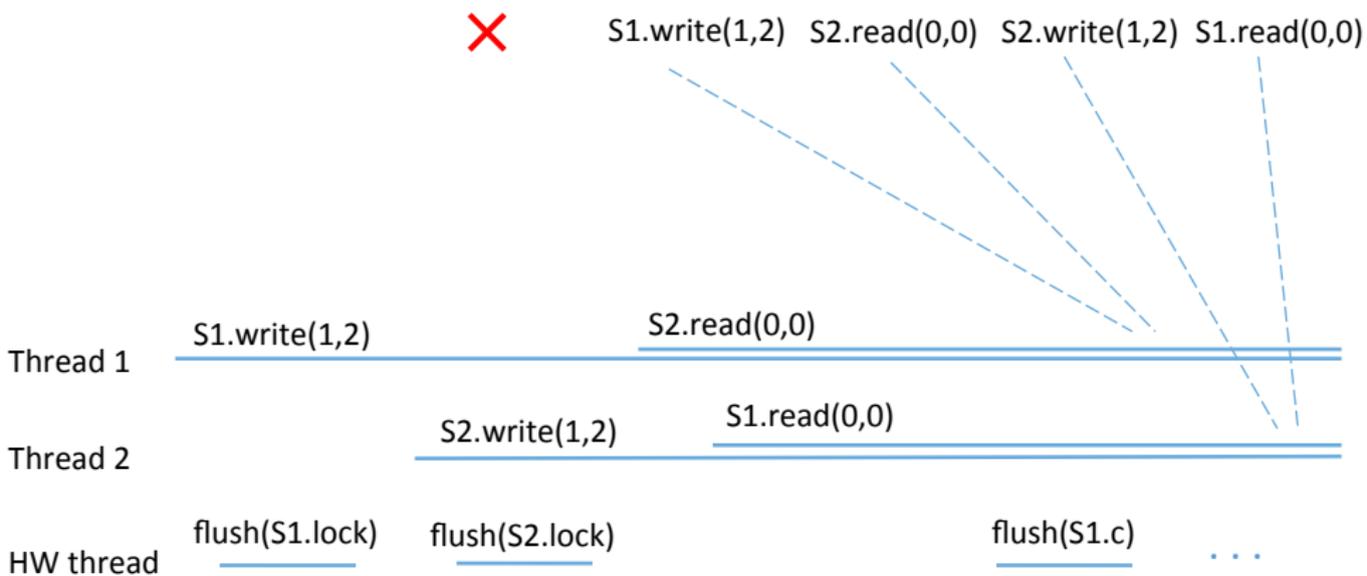
An operation remains active (and hence may take affect) up to the time of the last flush of a value written by the operation



... not sound

TSO-linearizability + sequential consistency

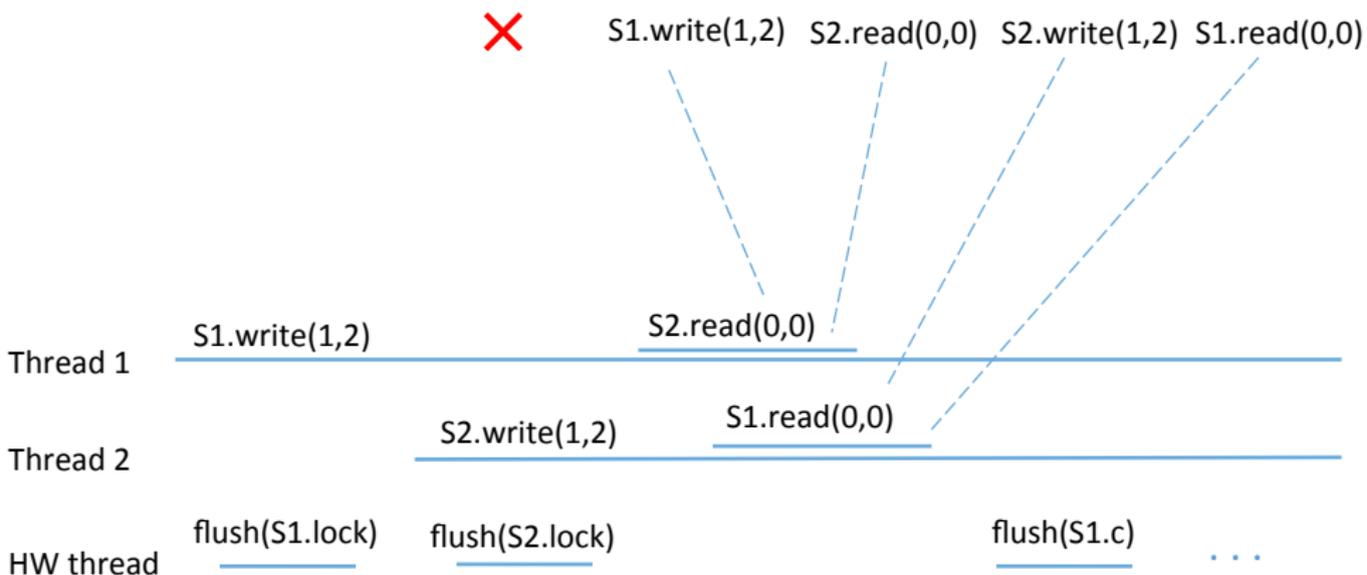
Sequential consistency requires same-thread operations linearize in the order that they are invoked



... but not compositional

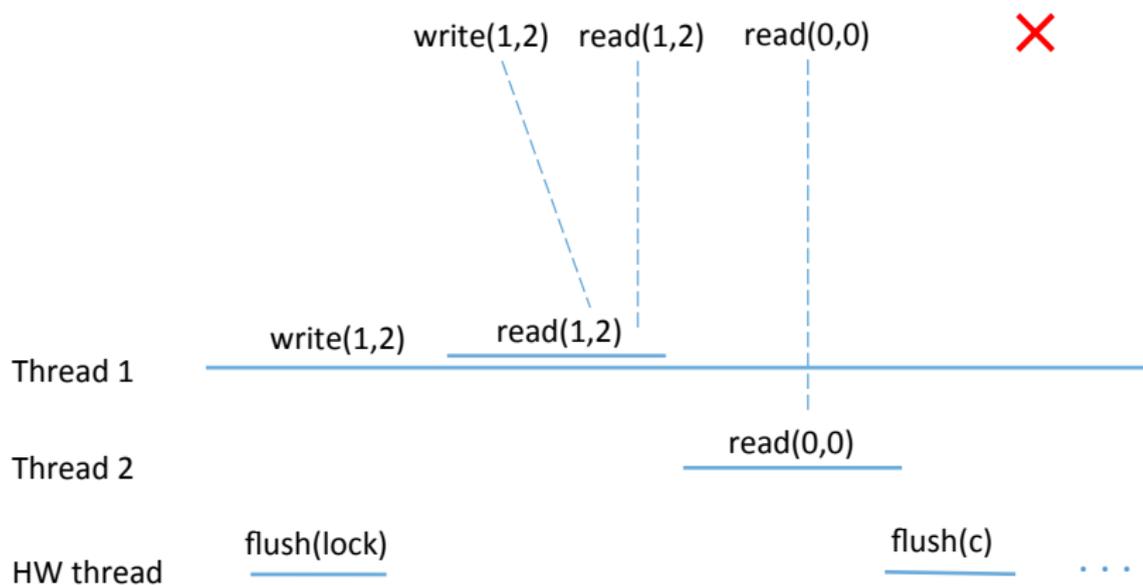
Latest approach – last Tuesday

Operations linearize before the linearization of the next same-thread operation which can be observed by the program



Latest approach – last Tuesday

Seems to be compositional – failure of running example to linearize is due to seqlock not linearizing



ARM and Power

- ▶ Reordering
 - ▶ instructions in a thread may be reordered
 - ▶ sequential semantics should be preserved
 - ▶ order of accesses to a given shared variable should be preserved
- ▶ Speculative execution
 - ▶ branches may be executed before the branch condition is evaluated
 - ▶ results of such execution cannot be committed until the condition is evaluated
 - ▶ allows reordering of instructions after the branch with instructions before
- ▶ Lack of multi-copy atomicity
 - ▶ thread A may see thread B's writes to global variables before thread C sees them

Reordering

- ▶ Reordering
 - ▶ instructions in a thread may be reordered
 - ▶ sequential semantics should be preserved
 - ▶ order of accesses to a given shared variable should be preserved

$$(a; c) \xrightarrow{a} c \qquad \frac{c \xrightarrow{b} c' \quad a \stackrel{R}{\leftarrow} b}{(a; c) \xrightarrow{\tilde{b}^a} (a; c')}$$

where $a \stackrel{R}{\leftarrow} b$ means b can be reordered before a
and \tilde{b}^a means b takes into account the effect of a (even though a has not yet occurred globally) — *forwarding*

Reordering in TSO

store $x, v \stackrel{R}{\Leftarrow} \mathbf{load} \ r, y$

store $x, v \stackrel{R}{\Leftarrow} \mathbf{reg} \ r, v'$

$a \stackrel{R}{\neq} b$ otherwise

$\overbrace{\mathbf{load} \ r, x}^{\mathbf{store} \ x, v} = \mathbf{reg} \ r, v$

Reordering in ARM and Power

store $x, v \stackrel{R}{\Leftarrow} \text{load } r, y$

store $x, v \stackrel{R}{\Leftarrow} \text{store } y, w$ iff $x \neq y$

load $r_1, x \stackrel{R}{\Leftarrow} \text{load } r_2, y$ iff $r_1 \neq r_2, x \neq y$

load $r, x \stackrel{R}{\Leftarrow} \text{store } y, v$ iff $x \neq y$

guard $b \stackrel{R}{\Leftarrow} \text{load } r, x$ iff $r \text{ nfi } b$

store $x, v \stackrel{R}{\Leftarrow} \text{guard } b$

load $r, x \stackrel{R}{\Leftarrow} \text{guard } b$ iff $r \text{ nfi } b$

guard $b_1 \stackrel{R}{\Leftarrow} \text{guard } b_2$

Speculative execution

- ▶ Speculative execution
 - ▶ branches may be executed before the branch condition is evaluated
 - ▶ results of such execution cannot be committed until the condition is evaluated
 - ▶ allows reordering of instructions after the branch with instructions before

Branch point: $(\mathbf{guard} \ b; \ c_1) \parallel (\mathbf{guard} \ \neg b; \ c_2)$

$$\frac{c \xrightarrow{\mathbf{guard} \ b} c' \quad eval_{\sigma}(b) \equiv true}{\sigma \bullet c \xrightarrow{\tau} \sigma \bullet c'}$$

Lack of multi-copy atomicity

- ▶ Lack of multi-copy atomicity
 - ▶ thread A may see thread B's writes to global variables before thread C sees them

$$w ::= (x, v, nl)$$

$$Q ::= w^*$$

Load:

$$\frac{c \xrightarrow{n: \text{load } r, x} c' \quad \forall (y, w, nl) \in Q_1 \bullet y \neq x \vee n \notin nl}{Q_1 \wedge (x, v, nl) \wedge Q_2 \bullet c \xrightarrow{n: \text{reg } r, v} Q_1 \wedge (x, v, nl') \wedge Q_2 \bullet c'}$$

where $n \in nl \Rightarrow nl' = nl$ and $n \notin nl \Rightarrow nl' = nl \wedge n$

Lack of multi-copy atomicity

- ▶ Lack of multi-copy atomicity
 - ▶ thread A may see thread B's writes to global variables before thread C sees them

$$w ::= (x, v, nl)$$

$$Q ::= w^*$$

Store:

$$\frac{c \xrightarrow{n: \text{store } x, v} c' \quad \forall w \in Q_1 \bullet (x, v, n) \stackrel{R}{\Leftarrow} w}{Q_1 \wedge Q_2 \bullet c \xrightarrow{n: \text{store } x, v} Q_1 \wedge (x, v, n) \wedge Q_2 \bullet c'}$$

where $(x, v, n) \stackrel{R}{\Leftarrow} (y, w, nl)$
iff $\text{head } nl \neq n \wedge (x = y \Rightarrow n \notin nl)$

Fences

fence $\stackrel{R}{\not\equiv} a$

$a \stackrel{R}{\not\equiv}$ **fence**

cfence $\stackrel{R}{\not\equiv}$ **load** r, x

guard $b \stackrel{R}{\not\equiv}$ **cfence**

$$\frac{c \stackrel{n: \text{fence}}{\longrightarrow} c'}{Q \bullet c \stackrel{n: \text{fence}}{\longrightarrow} \text{flush}_n(Q) \bullet c'}$$

$$\text{flush}_n(\langle \rangle) = \langle \rangle$$

$$\text{flush}_n((x, v, nl) \hat{\ } Q) = (x, v, nl) \hat{\ } \text{flush}_n(Q) \quad \text{if } n \notin nl$$

$$\text{flush}_n((x, v, nl) \hat{\ } Q) = (x, v, nl \hat{\ } nl') \hat{\ } \text{flush}_n(Q) \quad \text{if } n \in nl$$

where nl' contains all thread ids not in nl .

Validation

Passed 347 / 348 litmus tests for ARMv8

Passed 757 / 758 litmus tests for Power 7

Failed test (on both architectures):

```
store x, 1; fence; store y, 1 ||  
load r0, y; eor r1, r0, r0; add r1, r1, 1; store z, r1; store z, 2;  
load r3, z; (if r3 = r3 then noOp else noOp); cfence; load r4, x
```

Our semantics allows result $z = 2 \wedge r0 = 1 \wedge r4 = 0$. That is, last statement of process 2 occurs before first statement.