

Compositional reasoning
about concurrent libraries
on the axiomatic
TSO memory model

Artem Khyzha
IMDEA Software Institute, Madrid, Spain

Joint work with Alexey Gotsman (IMDEA Software)

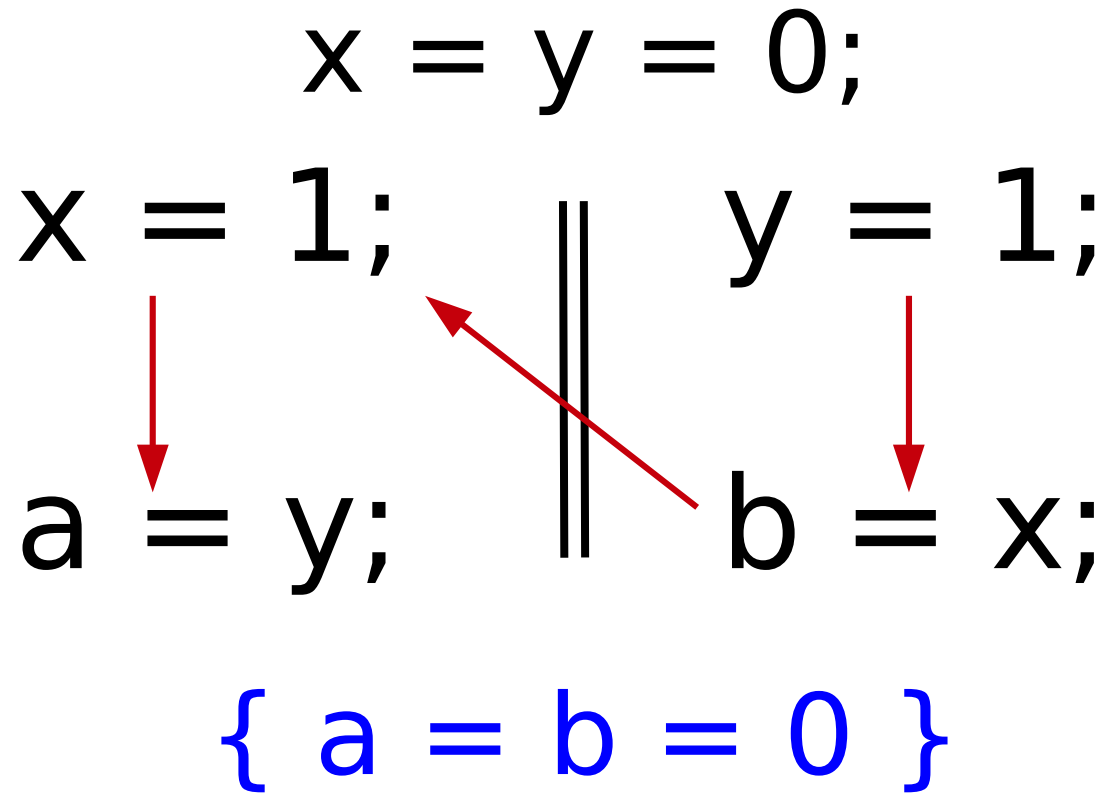
Weak memory

```
    x = y = 0;  
x = 1;    ||    y = 1;  
a = y;    ||    b = x;  
{ a = b = 0 }
```

Weak memory

```
    x = y = 0;  
x = 1;      |      y = 1;  
            |      |  
a = y;      |      b = x;  
            |  
{ a = b = 0 }
```

Weak memory




Weak memory

`x = y = 0;`

`x = 1;` `y = 1;`
↓ ↓
`a = y;` `b = x;`

~~`{ a = b = 0 }`~~

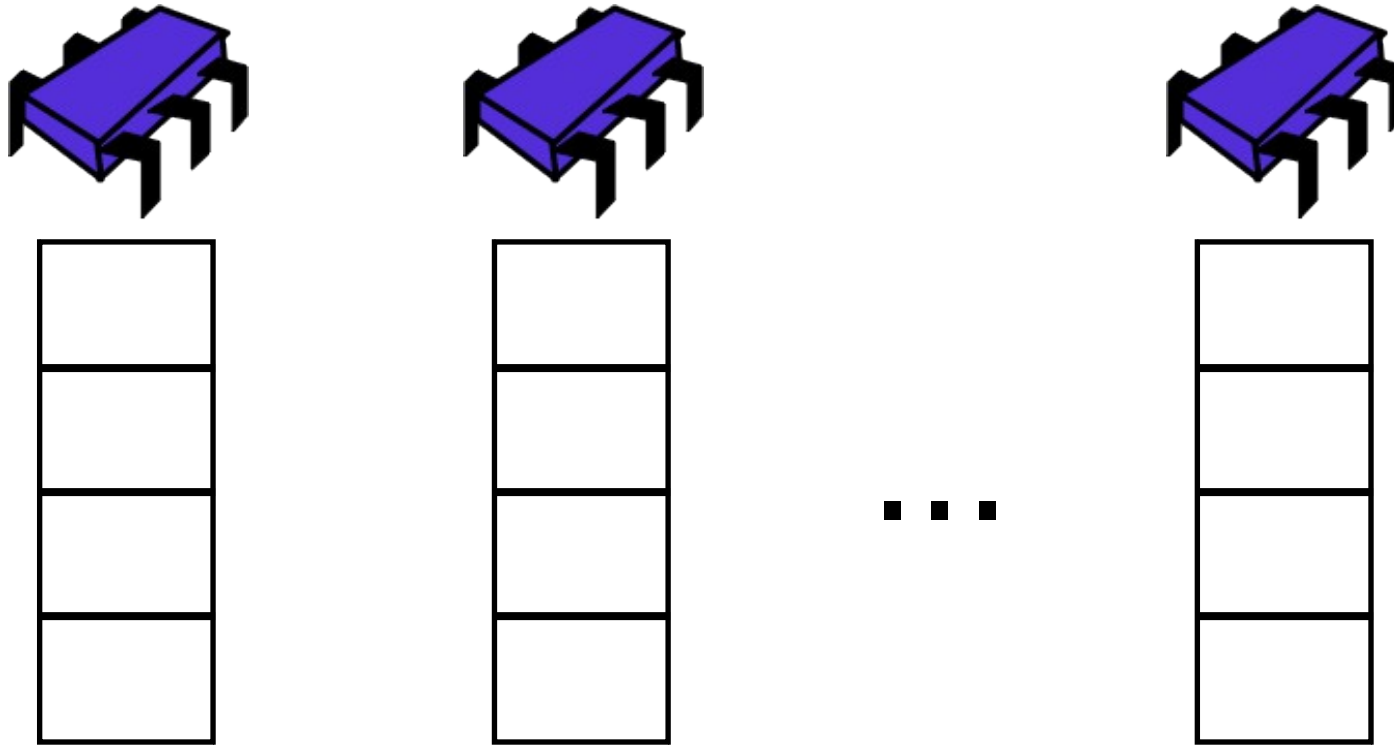
Weak memory



$\{ a = b = 0 \}$



Possible on a weak memory model

The TSO memory model (x86)



RAM

The TSO memory model (x86)





```
*x = a;  
*y = b;
```


Writes stored into
the write buffer in
the order of issue

RAM

The TSO memory model (x86)



```
*x = a;  
*y = b;
```





Writes stored into
the write buffer in
the order of issue




RAM

The TSO memory model (x86)





```
*x = a;  
*y = b;
```



Writes stored into
the write buffer in
the order of issue

RAM

The TSO memory model (x86)





```
*x = a;  
*y = b;
```

Writes flushed in
FIFO order




The TSO memory model (x86)





```
*x = a;  
*y = b;
```

Writes flushed in
FIFO order



The TSO memory model (x86)



```
*x = a;  
*y = b;
```

Writes flushed in
FIFO order



Our goal

- Is compositional reasoning on weak memory possible?

L (*Implementation*)

C

```
struct Node {
  Node *next; int val;
} *Top;

void push(int v) {
  Node *t, *x;
  x = new Node;
  x->val = v;
  do {
    t = Top; x->next = t;
  } while(!CAS(&Top,t,x));
}
```

→
abstracted by

L' (*Specification*)

C

```
Sequence S;

void push(int v) {
  atomic { S = v·S; }
}
```

Our goal

- Is compositional reasoning on weak memory possible?

Linearizability
[Herlihy, Wing]

L (*Implementation*)

C

```
struct Node {
  Node *next; int val;
} *Top;

void push(int v) {
  Node *t, *x;
  x = new Node;
  x->val = v;
  do {
    t = Top; x->next = t;
  } while(!CAS(&Top,t,x));
}
```

→
linearized by


L' (*Specification*)

C

```
Sequence S;

void push(int v) {
  atomic { S = v·S; }
}
```

Challenge




- Writes can be flushed after `f` returns
- Just parameters and return values not enough

Axiomatic model

An execution of $C(L)$ is $(A, po, rf, mo, \dots, hb)$

$$x = y = a = b = 0$$



- Actions A : reads, writes, calls, returns

Axiomatic model

An execution of $C(L)$ is $(A, po, rf, mo, \dots, hb)$


$$x = y = a = b = 0$$

$x = 1;$	\parallel	$y = 1;$
$\text{po} \downarrow$		$\text{po} \downarrow$
$a = y;$		$b = x;$

- **po** – program order
relates actions by the same thread

Axiomatic model


An execution of $C(L)$ is $(A, po, rf, mo, \dots, hb)$



- **mo** – modification order
order in which writes hit the memory
- **rf** – “reads from” relation
relates reads to correspondent writes

Axiomatic model

An execution of $C(L)$ is $(A, po, rf, mo, \dots, hb)$



- **hb** – “happens before” relation
indicates precedence in the whole execution

$$hb = (po \cup rf)^+$$

Axioms

MOWF. mo is total, transitive, irreflexive and relates only store actions in the execution.

POWF. po is total, transitive, irreflexive and relates only actions by the same thread.

SCWF. sc is total, transitive, irreflexive and relates only fences in the execution.

RFWF. $(\forall w_1, w_2, r. w_1 \xrightarrow{rf} r \wedge w_2 \xrightarrow{rf} r \implies w_1 = w_2) \wedge$
 $(\forall w, r. w \xrightarrow{rf} r \implies \exists x, a. w = (-, \text{store}(x, a)) \wedge (r = (-, \text{load}(x, a))) \vee r = (-, \text{call } _ (a)) \vee r = (-, \text{ret } _ (a)))$

RFDET. $\forall x, w, r. (r = (-, \text{load}(x, _)) \vee (\exists t. x = \text{param}_t \wedge r = (t, \text{call } _)) \vee$
 $(\exists t. x = \text{retval}_t \wedge r = (t, \text{ret } _))) \wedge w = (-, \text{store}(x, _)) \wedge (w \xrightarrow{hb} r \vee w \xrightarrow{po} r) \implies \exists w'. w' \xrightarrow{rf} r$

HBDEF. $hb = (po \cup rf)^+$

HBVSMO. $\neg \exists w_1, w_2. w_1 \xrightarrow{hb} w_2$
 $w_2 \xrightarrow{mo} w_1$

RFMR. $\neg \exists w_1, w_2, r. w_1 \xrightarrow{mo} w_2 \xrightarrow{hb} r$
 $w_1 \xrightarrow{rf} r$

where w_1, w_2 and r access the same location.

RFMR'. $\neg \exists w, w_1, w_2, r.$

$w_1 \xrightarrow{mo} w_2 \xrightarrow{mo} w \xrightarrow{rf} r' \xrightarrow{sb} r$
 $w_1 \xrightarrow{rf} r$

where w_1 and w_2 write to the same location,
and w and r' are by different threads.

HBWF. hb is acyclic.

HBVSSC. $\neg \exists f_1, f_2. f_1 \xrightarrow{hb} f_2$
 $f_2 \xrightarrow{sc} f_1$

MOVSSC. $\neg \exists w_1, w_2, f_1, f_2.$

$w_1 \xrightarrow{hb} f_1 \xrightarrow{sc} f_2 \xrightarrow{hb} w_2$
 $w_1 \xrightarrow{mo} w_2$

SCRF. $\neg \exists w, w', f_1, f_2, r.$

$w \xrightarrow{mo} w' \xrightarrow{po} f_1 \xrightarrow{sc} f_2 \xrightarrow{po} r$
 $w \xrightarrow{rf} r$

Axioms

- Filter out impossible executions
- Capture the effect of store buffering

$$\text{HB vs MO. } \neg \exists w_1, w_2. \quad w_1 \begin{array}{c} \xrightarrow{\text{hb}} \\ \xleftarrow{\text{mo}} \end{array} w_2$$

Disallows “writing to the future”


Abstraction Theorem

Assume that $L \sqsubseteq L'$.

Then $\text{client}(\llbracket C(L) \rrbracket) \subseteq \text{client}(\llbracket C(L') \rrbracket)$


L' specifies L :

$$L \sqsubseteq L'$$

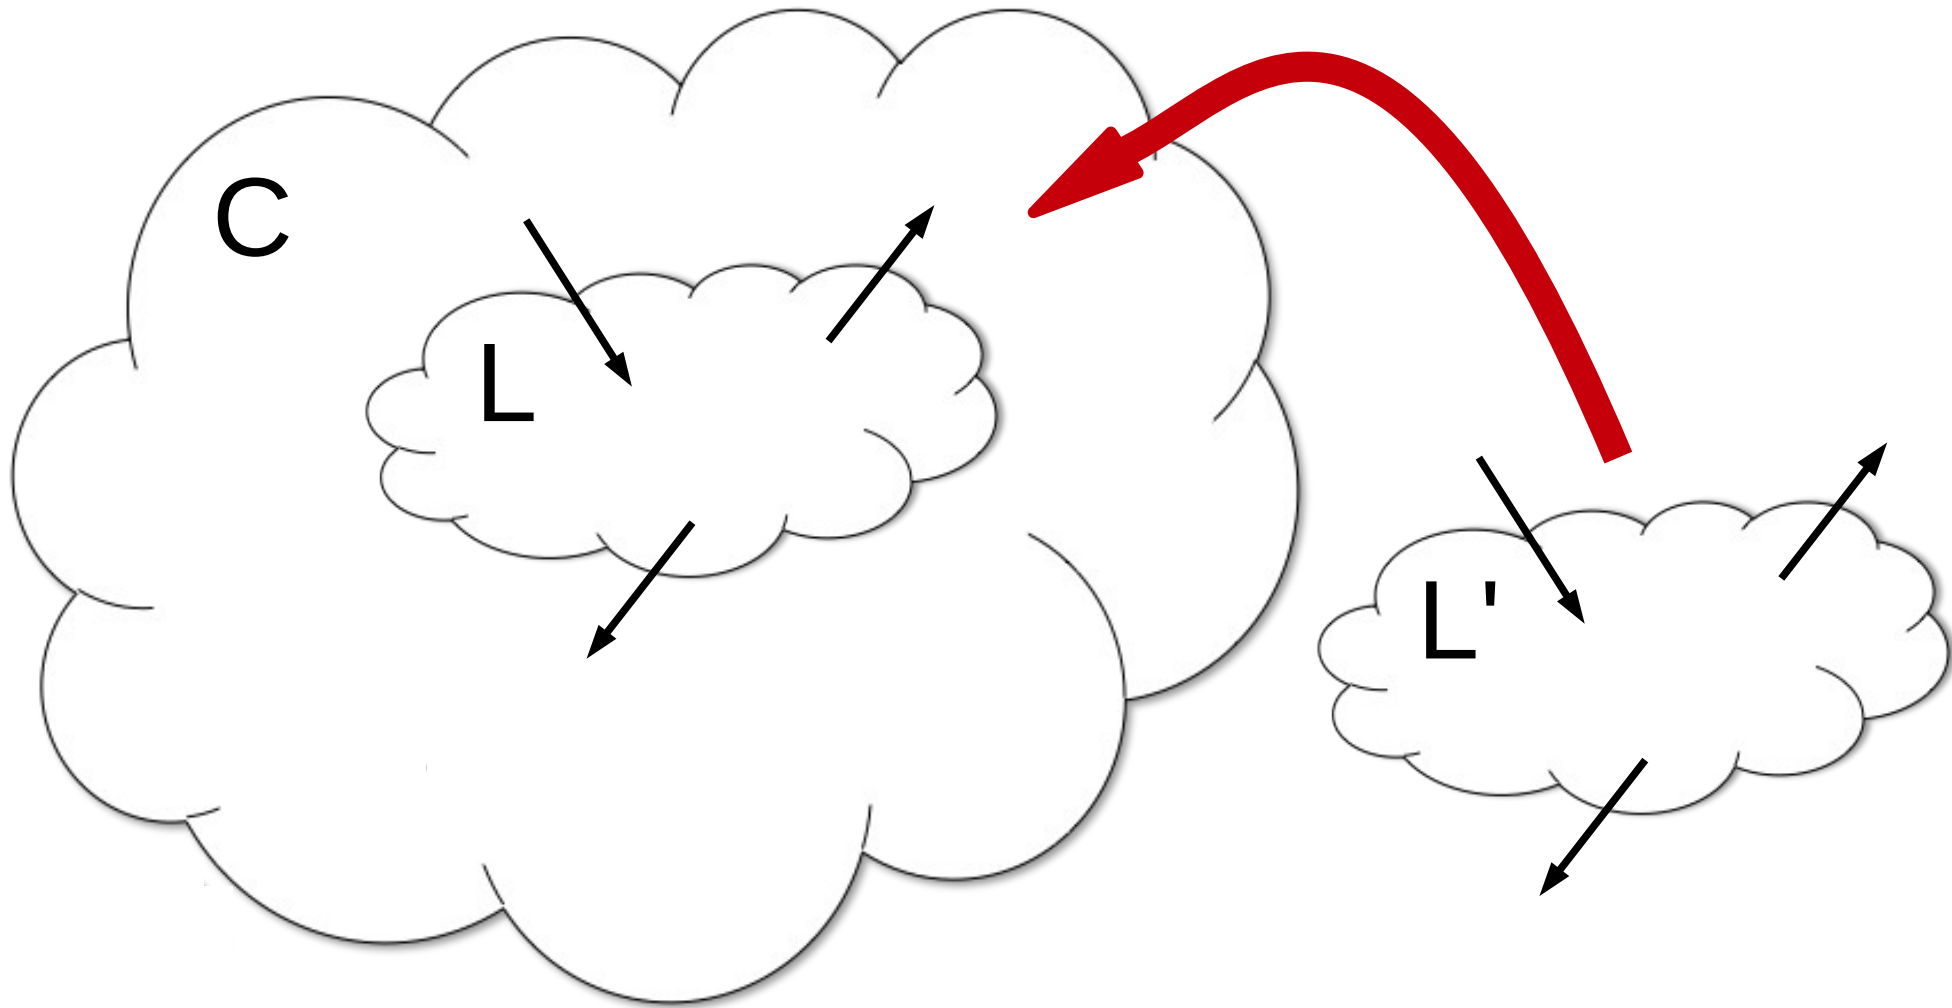


$$C(L') \models P \Rightarrow C(L) \models P$$

Subgraph replacement

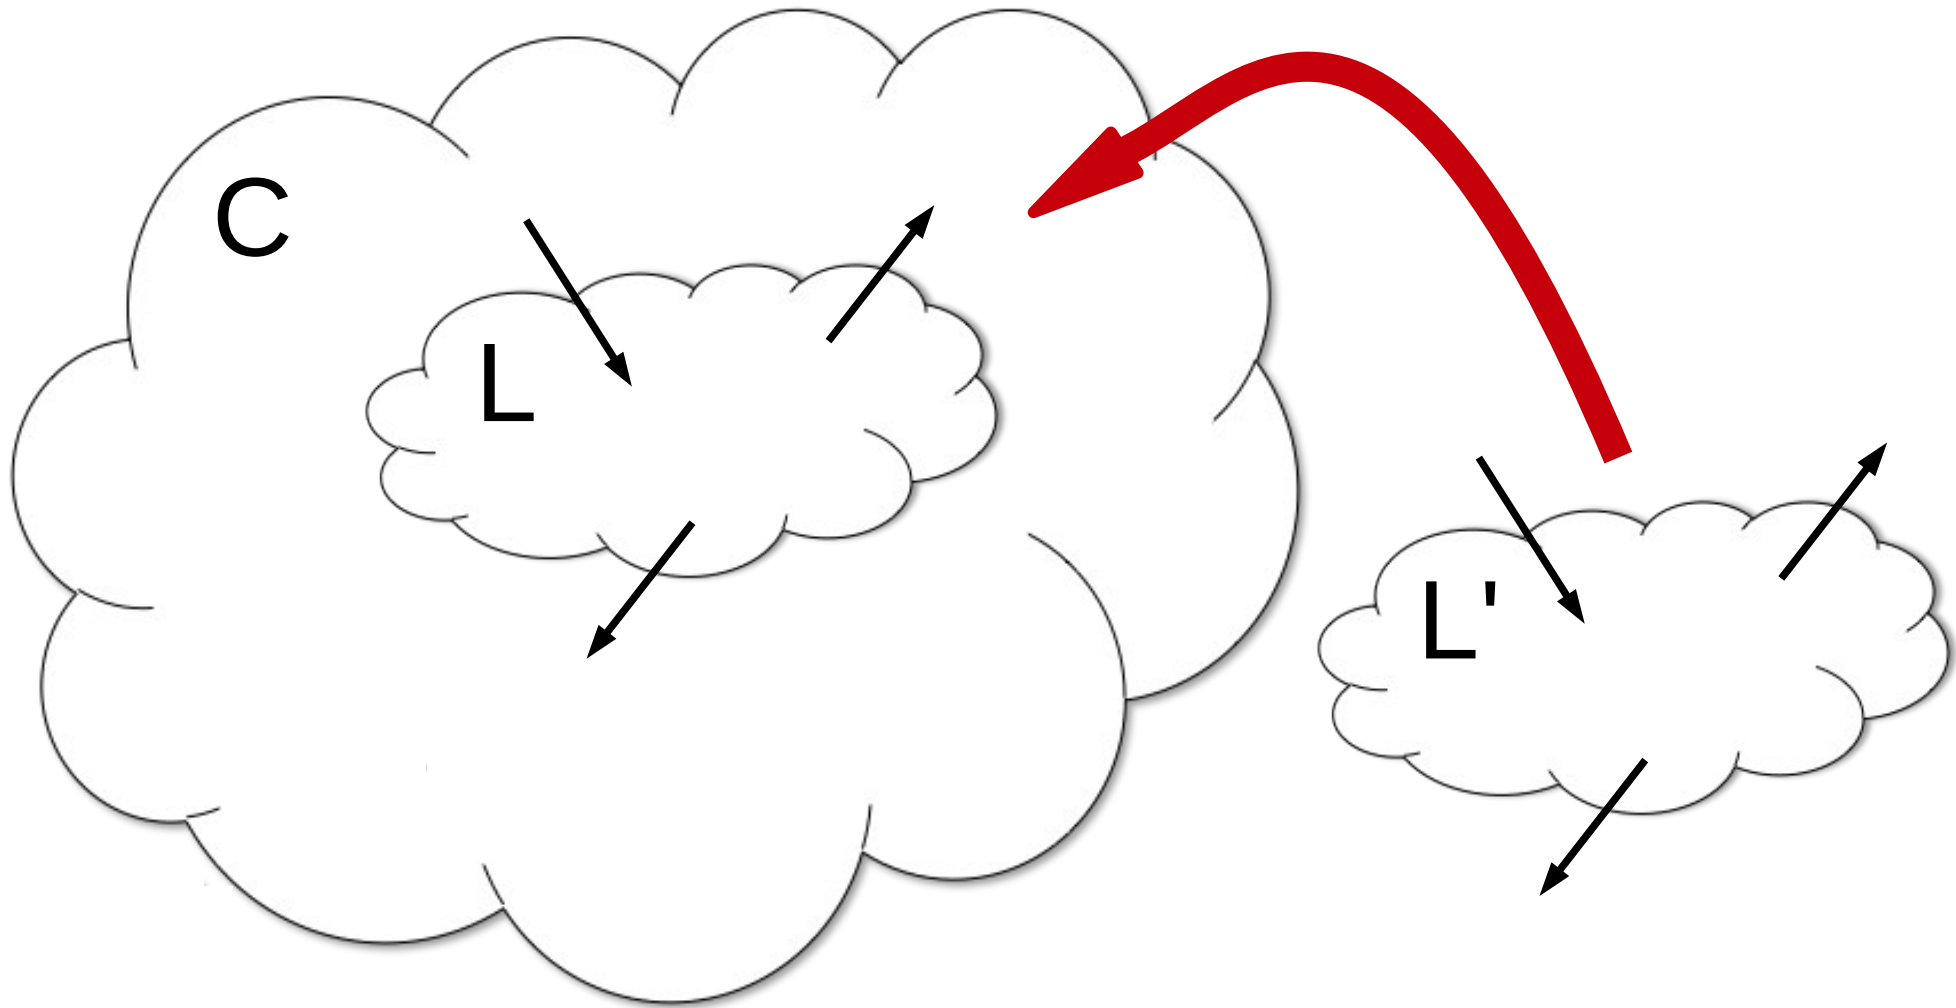


Subgraph replacement



History - several relations on boundary actions

Subgraph replacement



$$L \sqsubseteq L' \iff \forall H \in \llbracket L \rrbracket . \exists H' \in \llbracket L' \rrbracket . H \sqsubseteq H'$$

Computing histories

- Take a library L on TSO or SC
- Take the **most general client**:

$$\begin{array}{c} n \\ \parallel \\ k=1 \end{array} \left(\begin{array}{l} \text{while (true) \{ } \\ \quad \text{if (nondet())} \quad m_1(\text{nondet()}); \\ \quad \text{else if (nondet())} \quad m_2(\text{nondet()}); \\ \quad \dots \\ \quad \text{else} \quad m_l(\text{nondet()}); \\ \quad \} \end{array} \right)$$

- Get all possible library executions: $\llbracket L \rrbracket_{\text{TSO}}$

$$L \sqsubseteq L' \iff \forall H \in \llbracket L \rrbracket . \exists H' \in \llbracket L' \rrbracket . H \sqsubseteq H'$$

Computing histories

- Take a library L on TSO or SC
- Take the **most general client**:

Any methods,
in any order,
with any parameters

Any number of
threads

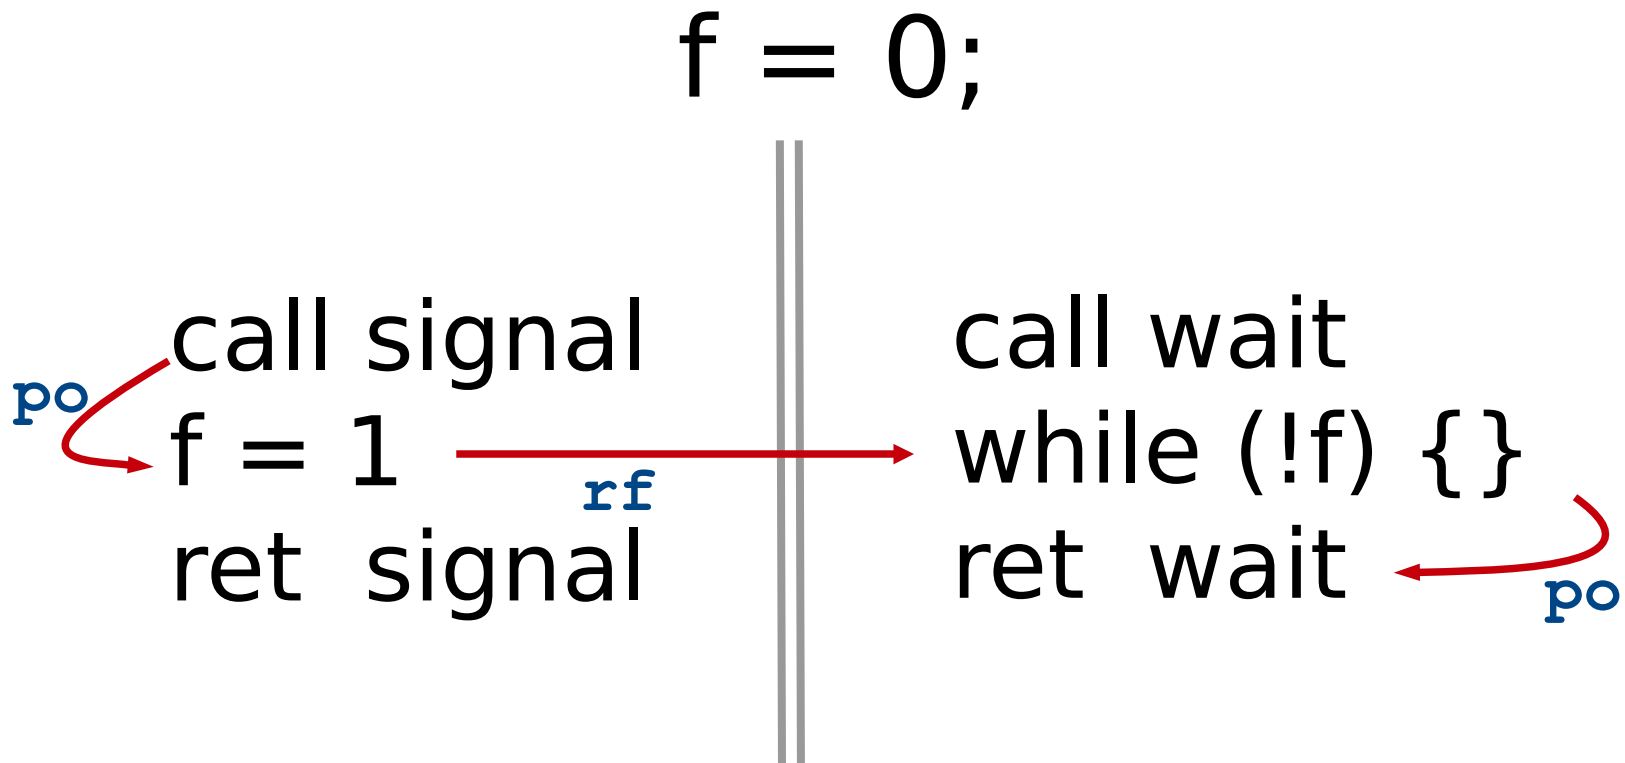
$$\begin{array}{l} n \\ \parallel \\ k=1 \end{array} \left(\begin{array}{l} \text{while (true) \{ \\ \quad \text{if (nondet())} \quad m_1(\text{nondet()}); \\ \quad \text{else if (nondet())} \quad m_2(\text{nondet()}); \\ \quad \dots \\ \quad \text{else} \quad m_l(\text{nondet()}); \\ \} \end{array} \right)$$

- Get all possible library executions: $\llbracket L \rrbracket_{\text{TSO}}$

$$L \sqsubseteq L' \iff \forall H \in \llbracket L \rrbracket . \exists H' \in \llbracket L' \rrbracket . H \sqsubseteq H'$$

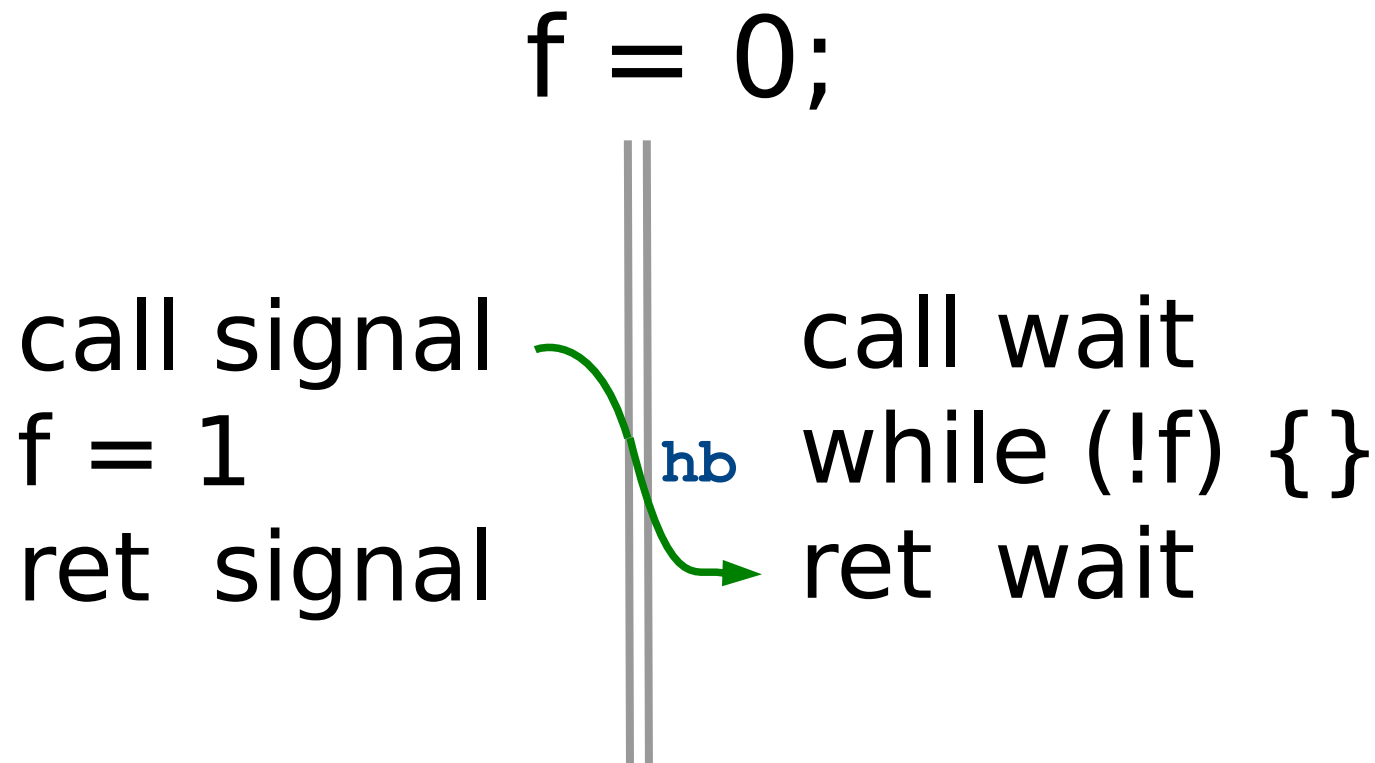
Guarantee relation

- Histories include the projection of happens-before to calls and returns:



Guarantee relation

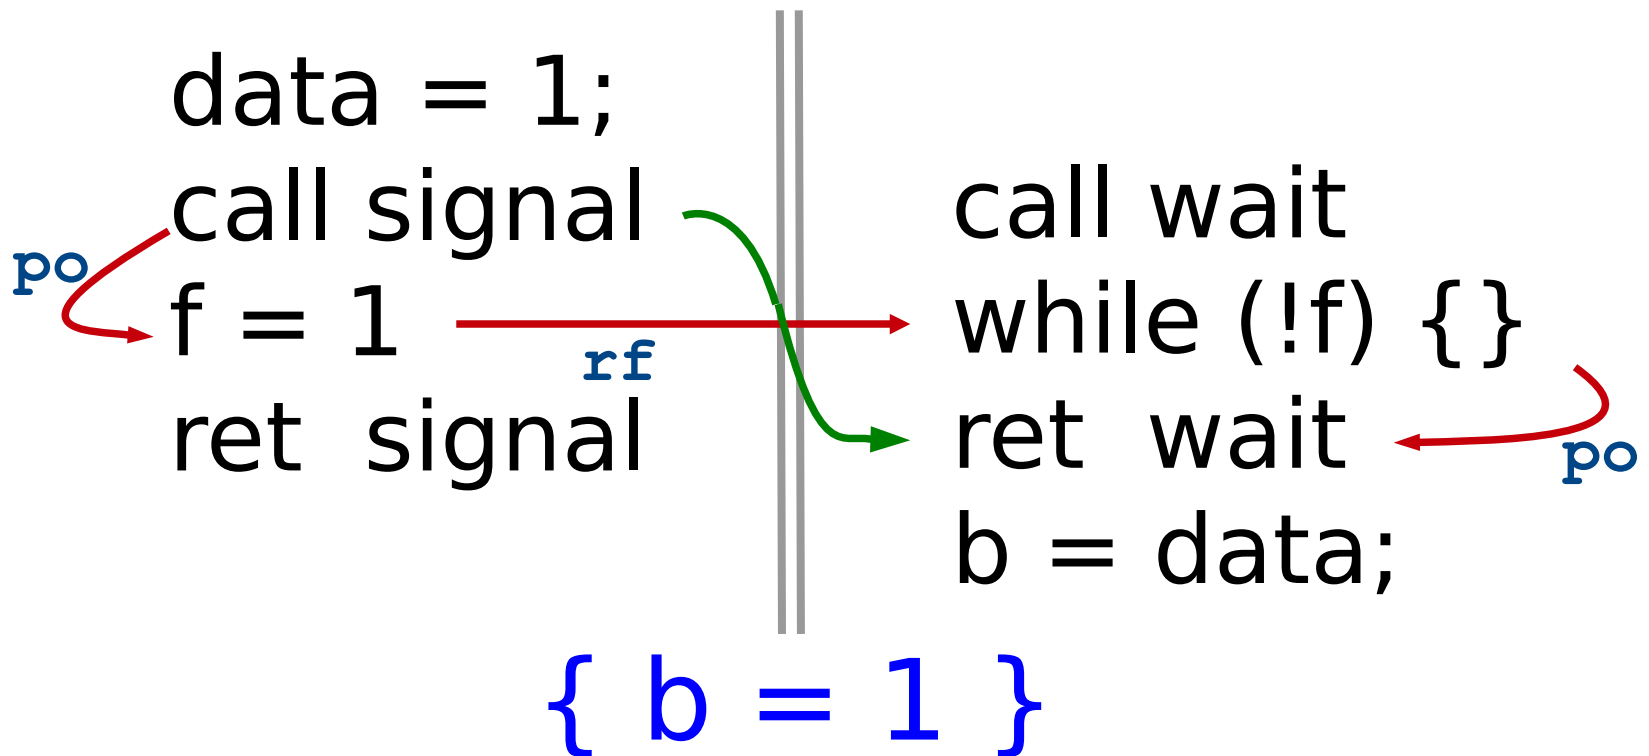
- Histories include the projection of happens-before to calls and returns:



Guarantee relation

- Histories include the projection of happens-before to calls and returns:

data = f = 0;



Conclusions

- Through thorns to the stars – compositional reasoning is possible on TSO.
- Axiomatic semantics can be straightforwardly implemented in SAT solvers
- Can be applied to more complicated models (IBM Power, ARM)

Related work

- **Concurrent library correctness on the TSO memory model**
- *S. Burckhardt et al. (ESOP'12)*

- **Library abstraction for C/C++ concurrency**
- *M. Batty et al. (POPL'13)*

Linearizability

- History: (Interface, Guarantee, Deny)
- $(I_1, G_1, D_1) \sqsubseteq (I_2, G_2, D_2) \iff (I_1 = I_2) \wedge (G_2 \subseteq G_1) \wedge (D_2 \subseteq D_1)$
- Specification can guarantee and deny less to the client
- So client has more behaviours when using specification instead of implementation
- $L_1 \sqsubseteq L_2 \iff \forall H \in \llbracket L \rrbracket . \exists H' \in \llbracket L' \rrbracket . H \sqsubseteq H'$

Linearizability

H t1: (t1, call push(42)) (t1, ret push) (t1, call isEmpty) (t1, ret isEmpty(yes))

t2: (t2, call pop) (t2, ret pop(42)) (t2, call push(11)) (t2, ret push)

H' (t1, call push(42)) (t2, call pop) (t1, call isEmpty) (t2, call push(11))
 (t1, ret push) (t2, ret pop(42)) (t1, ret isEmpty(yes)) (t2, ret push)

$H \sqsubseteq H'$

- We can permute calls and returns by different threads
- But non-overlapping method invocations can't be rearranged

Linearizability

- How it fails:

$x = y = 0;$

$x = 1;$		$y = 1;$
$m();$		$m();$
$a = y;$		$b = x;$

$\{ a = b = 0 \}$

Linearizability

- How it fails:

```
x = y = 0;
x = 1;      |      y = 1;
m();        |      m();
a = y;      |      b = x;
           { a = b = 0 }
```

- Not possible, if *m* contains a fence.