

Weak Memory Concurrency in C/C++11

Ori Lahav

<http://www.cs.tau.ac.il/~orilahav/>



Hydra
Distributed Computing Conference

July 11, 2019

Example: Dekker's mutual exclusion

Initially, $x = y = 0$.

$x := 1;$

$a := y;$

if ($a = 0$) **then**

/ critical section */*

||

$y := 1;$

$b := x;$

if ($b = 0$) **then**

/ critical section */*

Example: Dekker's mutual exclusion

Initially, $x = y = 0$.

$x := 1;$

$a := y;$

if ($a = 0$) **then**

/ critical section */*

||

$y := 1;$

$b := x;$

if ($b = 0$) **then**

/ critical section */*

Is it safe?

Example: Dekker's mutual exclusion

Initially, $x = y = 0$.

```
x := 1;
```

```
a := y; // 0
```

```
if (a = 0) then
```

```
  /* critical section */
```

```
|||
```

```
y := 1;
```

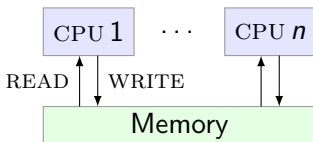
```
b := x; // 0
```

```
if (b = 0) then
```

```
  /* critical section */
```

Is it safe?

Yes, if we assume sequential consistency (SC):



Example: Dekker's mutual exclusion

Initially, $x = y = 0$.

```
x := 1;
```

```
a := y; // 0
```

```
if (a = 0) then
```

```
  /* critical section */
```

```
|||
```

```
y := 1;
```

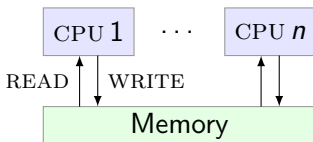
```
b := x; // 0
```

```
if (b = 0) then
```

```
  /* critical section */
```

Is it safe?

Yes, if we assume sequential consistency (SC):



No existing hardware implements SC!

- ▶ SC is very expensive (memory ~ 100 times slower than CPU).
- ▶ SC does not scale to many processors.

Example: Shared-memory concurrency in C++

```
int X, Y, a, b;  
  
void thread1() {  
    X = 1;  
    a = Y;  
}  
  
void thread2() {  
    Y = 1;  
    b = X;  
}
```

```
int main () {  
    int cnt = 0;  
  
    do {  
        X = 0; Y = 0;  
  
        thread first(thread1);  
        thread second(thread2);  
  
        first.join();  
        second.join();  
        cnt++;  
  
    } while (a != 0 || b != 0);  
  
    printf("%d\n",cnt);  
    return 0;  
}
```

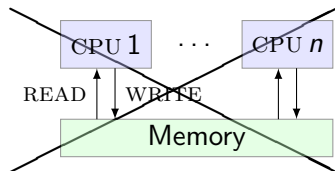
Example: Shared-memory concurrency in C++

```
int X, Y, a, b;  
  
void thread1() {  
    X = 1;  
    a = Y;  
}  
  
void thread2() {  
    Y = 1;  
    b = X;  
}
```

If Dekker's mutual exclusion is safe, this program will not terminate

```
int main () {  
    int cnt = 0;  
  
    do {  
        X = 0; Y = 0;  
  
        thread first(thread1);  
        thread second(thread2);  
  
        first.join();  
        second.join();  
        cnt++;  
  
    } while (a != 0 || b != 0);  
  
    printf("%d\n", cnt);  
    return 0;  
}
```

Weak memory models



We look for a *substitute for SC*:

Unambiguous specification

- ▶ What are the possible outcomes of a multithreaded program?

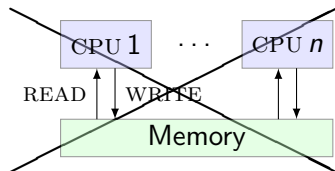
Amenable to formal reasoning

- ▶ Can prove theorems about the model.

Typically called a **weak memory model (WMM)**

- ▶ Allows more behaviors than SC.

Weak memory models



We look for a *substitute for SC*:

Unambiguous specification

- ▶ What are the possible outcomes of a multithreaded program?

Amenable to formal reasoning

- ▶ Can prove theorems about the model.

Typically called a **weak memory model (WMM)**

- ▶ Allows more behaviors than SC.

But it is not easy to get right

- ▶ The Java memory model is flawed.
- ▶ The C/C++11 model is also flawed.

The Problem of Programming Language Concurrency Semantics

Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod,
and Peter Sewell

University of Cambridge

“Disturbingly, 40+ years after the first relaxed-memory hardware was introduced (the IBM 370/158MP), the field still *does not have a credible proposal for the concurrency semantics* of any general-purpose high-level language that includes high performance shared-memory concurrency primitives. This is a *major open problem* for programming language semantics.”

European Symposium on Programming (ESOP) 2015

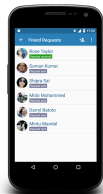
Plan for rest of the talk

1. Challenges for memory models
2. The C/C++11 memory model
3. The “out-of-thin-air” problem
4. A solution: a promising semantics

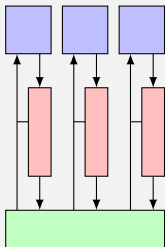
Plan for rest of the talk

1. Challenges for memory models
2. The C/C++11 memory model
3. The “out-of-thin-air” problem
4. A solution: a promising semantics

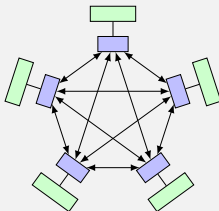
Challenge 1: Various hardware models




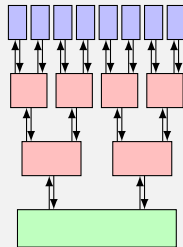
x86-TSO  
(2010)



POWER 
(2011)



ARMv8 
(2016)



Store buffering in x86-TSO



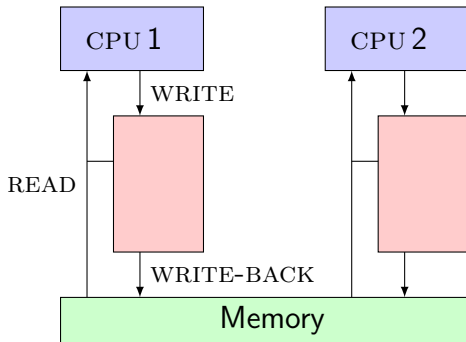
Initially, $x = y = 0$.

```
x := 1;
```

```
y := 1;
```

```
a := y; // 0
```

```
b := x; // 0
```



Store buffering in x86-TSO



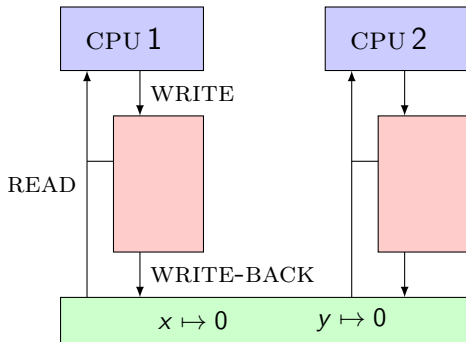
Initially, $x = y = 0$.

► $x := 1;$

► $y := 1;$

$a := y; // 0$

$b := x; // 0$



Store buffering in x86-TSO



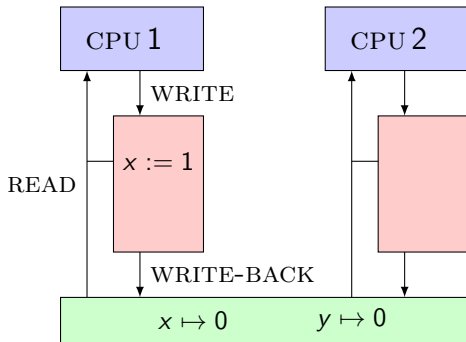
Initially, $x = y = 0$.

$x := 1;$

▶ $y := 1;$

▶ $a := y; \text{// } 0$

$b := x; \text{// } 0$



Store buffering in x86-TSO



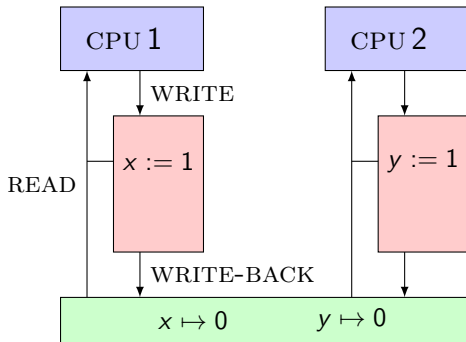
Initially, $x = y = 0$.

$x := 1;$

$y := 1;$

▶ $a := y; // 0$

▶ $b := x; // 0$



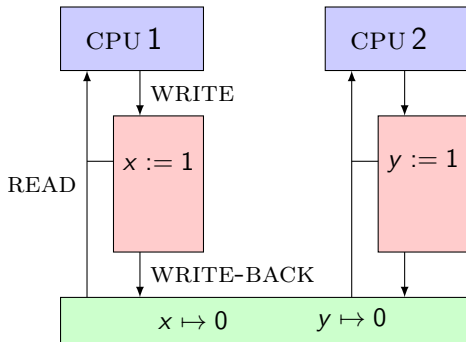
Store buffering in x86-TSO



Initially, $x = y = 0$.

```
x := 1;  
fence;  
a := y; // 0
```

```
y := 1;  
fence;  
b := x; // 0
```

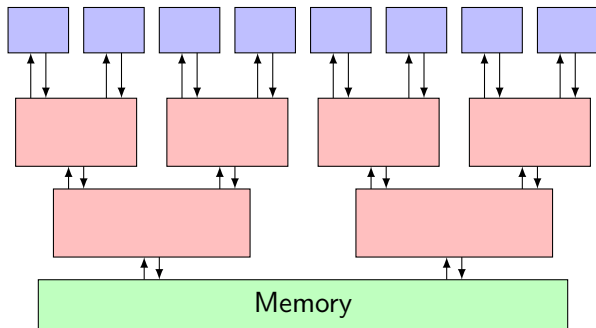


Load buffering in ARM



Initially, $x = y = 0$.

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

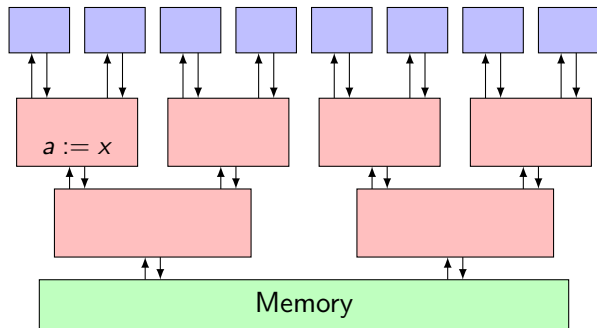


Load buffering in ARM



Initially, $x = y = 0$.

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

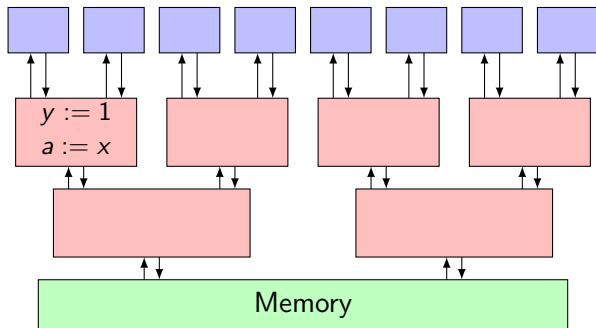


Load buffering in ARM



Initially, $x = y = 0$.

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

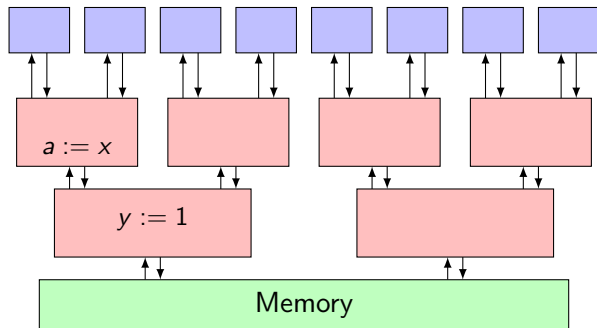


Load buffering in ARM



Initially, $x = y = 0$.

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

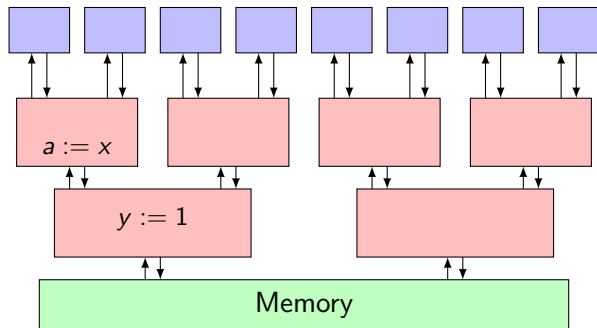


Load buffering in ARM



Initially, $x = y = 0$.

```
    a := x; // 1    ||    b := y; // 1  
    y := 1;        ||    x := b;
```



Challenge 2: Compilers stir the pot

Initially, $x = y = 0$.

```
x := 1; || a := x;  
y := 1; || b := y; // 1  
        || c := x; // 0
```

X forbidden under SC

Challenge 2: Compilers stir the pot

Initially, $x = y = 0$.

$x := 1;$ \parallel $a := x;$
 $y := 1;$ \parallel $b := y; \text{ // } 1$
 \parallel $c := x; \text{ // } 0$

X forbidden under SC



*compiler
optimization*



$x := 1;$ \parallel $a := x;$
 $y := 1;$ \parallel $b := y; \text{ // } 1$
 \parallel $c := a; \text{ // } 0$

✓ allowed under SC

Challenge 3: Transformations do not suffice

Program transformations fail short to explain some weak behaviors.

- ▶ In C/C++:
 - ▶ *Release stores* cannot be reordered.
 - ▶ *Acquire loads* cannot be reordered.

Message passing (MP)

```
x :=rel 1; || a := yacq; // 1  
y :=rel 1; || b := xacq; // 0
```

Challenge 3: Transformations do not suffice

Program transformations fail short to explain some weak behaviors.

- ▶ In C/C++:
 - ▶ *Release stores* cannot be reordered.
 - ▶ *Acquire loads* cannot be reordered.

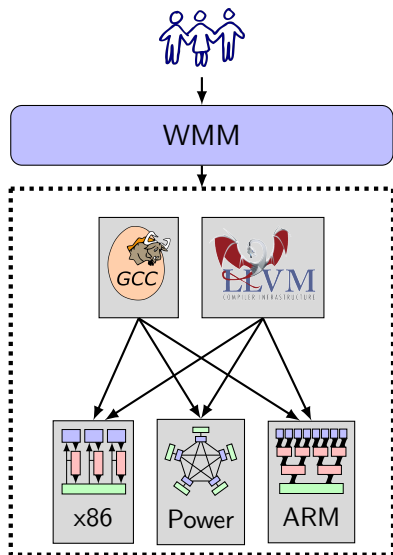
Message passing (MP)

$$\begin{array}{l} x :=_{\text{rel}} 1; \\ y :=_{\text{rel}} 1; \end{array} \parallel \begin{array}{l} a := y_{\text{acq}}; \text{ // } 1 \\ b := x_{\text{acq}}; \text{ // } 0 \end{array}$$

- ▶ And yet, since C/C++ is intended to be compiled to a *non-multi-copy-atomic* architectures:

Independent reads of independent writes (IRIW)

$$\begin{array}{l} a := x_{\text{acq}}; \text{ // } 1 \\ b := y_{\text{acq}}; \text{ // } 0 \end{array} \parallel x :=_{\text{rel}} 1; \parallel y :=_{\text{rel}} 1; \parallel \begin{array}{l} c := y_{\text{acq}}; \text{ // } 1 \\ d := x_{\text{acq}}; \text{ // } 0 \end{array}$$



WMM desiderata

1. Formal and comprehensive
2. Not too weak
(good for programmers)
3. Not too strong
(good for hardware)
4. Admits optimizations
(good for compilers)

Implementability
vs.
Programmability

The C11 memory model

- ▶ Introduced by the ISO C/C++ 2011 standards.
- ▶ Defines the semantics of **concurrent** memory accesses.

The C11 memory model: Atomics

Two types of accesses

**Ordinary
(Non-Atomic)**

Races are **errors**

Atomic

Welcome to the
expert mode

The C11 memory model: Atomics

Two types of accesses

**Ordinary
(Non-Atomic)**

Races are **errors**

Atomic

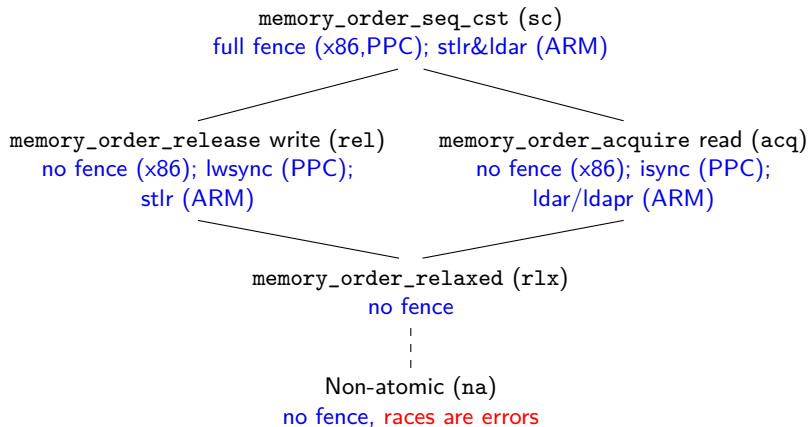
Welcome to the
expert mode

DRF (data race freedom) guarantee

no data races
under SC \implies only
SC behaviors

A spectrum of access modes

non-atomic □ relaxed □ release/acquire □ sc



+ Explicit primitives for fences

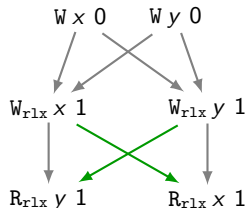
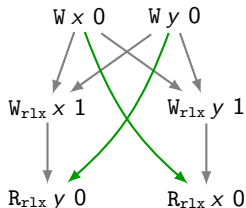
C11: a declarative memory model

Declarative semantics abstracts away from implementation details.

1. a program \rightsquigarrow a set of directed **graphs**.
2. The model defines what executions are **consistent**.
3. C/C++11 also has **catch-fire** semantics (forbidden data races).

Execution graphs

Store buffering (SB)

$$\begin{array}{l} x = y = 0 \\ x :=_{rlx} 1; \parallel y :=_{rlx} 1; \\ a :=_{rlx}; \parallel b :=_{rlx}; \end{array}$$


Relations

- ▶ Program order, po
- ▶ Reads-from, rf

$$\begin{aligned}
[-] &: \text{CExp} \rightarrow \mathbb{P}((\text{res} : \text{Val} \cup \{\perp\}, \mathcal{A} : \mathbb{P}(\text{AName}), \text{lab} : \mathcal{A} \rightarrow \text{Act}, \text{sb} : \mathbb{P}(\mathcal{A} \times \mathcal{A}), \text{fst} : \mathcal{A}, \text{lst} : \mathcal{A})) \\
[v] &\stackrel{\text{def}}{=} \{\langle v, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{skip}\} \\
[\text{alloc}()] &\stackrel{\text{def}}{=} \{\langle \ell, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \ell \in \text{Loc} \wedge \text{lab}(a) = \Lambda(\ell) \} \\
[[v]z ::= v'] &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{Wz}(v, v') \} \\
[[v]z &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge v' \in \text{Val} \wedge \text{lab}(a) = \text{Rz}(v, v') \} \\
[\text{CAS}_{X,Y}(v, v_0, v_n)] &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge v' \in \text{Val} \wedge v' \neq v_0 \wedge \text{lab}(a) = \text{Ry}(v, v') \} \\
&\cup \{\langle v_0, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{RMW}_X(v, v_0, v_n) \} \\
[\text{let } x = E_1 \text{ in } E_2] &\stackrel{\text{def}}{=} \{\langle \perp, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \mid \langle \perp, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in [E_1] \} \\
&\cup \{\langle \text{res}_2, \mathcal{A}_1 \uplus \mathcal{A}_2, \text{lab}_1 \cup \text{lab}_2, \text{sb}_1 \cup \text{sb}_2 \cup \{(\text{lst}_1, \text{fst}_2 \} \}, \text{fst}_1, \text{lst}_2 \} \mid \\
&\quad \langle \text{res}_1, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in [E_1] \wedge \langle \text{res}_2, \mathcal{A}_2, \text{lab}_2, \text{sb}_2, \text{fst}_2, \text{lst}_2 \rangle \in [E_2[v_1/x]] \} \\
[\text{repeat } E \text{ end}] &\stackrel{\text{def}}{=} \{\langle \text{res}_N, \bigcup_{i \in [1..N]} \mathcal{A}_i, \bigcup_{i \in [1..N]} \text{lab}_i, \bigcup_{i \in [1..N]} \text{sb}_i \cup \{(\text{lst}_1, \text{fst}_2, \dots, (\text{lst}_{N-1}, \text{fst}_N) \} \}, \text{fst}_1, \text{lst}_N \rangle \mid \\
&\quad \forall i. \langle \text{res}_i, \mathcal{A}_i, \text{lab}_i, \text{sb}_i, \text{fst}_i, \text{lst}_i \rangle \in [E] \wedge (i \neq N \implies \text{res}_i = 0 \mapsto \text{res}_N \neq 0) \} \\
[[E_1]E_2] &\stackrel{\text{def}}{=} \{\langle \text{combine}(\text{res}_1, \text{res}_2), \mathcal{A}_1 \uplus \mathcal{A}_2 \uplus \{a_{\text{fork}}, a_{\text{join}}\}, \text{lab}_1 \cup \text{lab}_2 \cup \{a_{\text{fork}} \mapsto \text{skip}, a_{\text{join}} \mapsto \text{skip}\}, \\
&\quad \text{sb}_1 \cup \text{sb}_2 \cup \{(a_{\text{fork}}, \text{fst}_1), (a_{\text{fork}}, \text{fst}_2), (\text{lst}_1, a_{\text{join}}), (\text{lst}_2, a_{\text{join}}), a_{\text{fork}}, a_{\text{join}}\} \mid \\
&\quad \langle \text{res}_1, \mathcal{A}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in [E_1] \wedge \langle \text{res}_2, \mathcal{A}_2, \text{sb}_2, \text{fst}_2, \text{lst}_2 \rangle \in [E_2] \wedge a_{\text{fork}}, a_{\text{join}} \in \text{AName} \}
\end{aligned}$$

Figure 2. Semantics of closed program expressions.

$$\begin{aligned}
&\not\exists x. \text{hb}(x, x) && \text{(IrreflexiveHB)} \\
&\forall \ell. \text{totalorder}(\{a \in \mathcal{A} \mid \text{iswrite}_\ell(a)\}, \text{mo}) \wedge \text{hb}_\ell \subseteq \text{mo} && \text{(ConsistentMO)} \\
&\text{totalorder}(\{a \in \mathcal{A} \mid \text{isSeqCst}(a)\}, \text{sc}) \wedge \text{hb}_{\text{SeqCat}} \subseteq \text{sc} \wedge \text{mo}_{\text{SeqCat}} \subseteq \text{sc} && \text{(ConsistentSC)} \\
&\forall b. \text{rf}(b) \neq \perp \iff \exists \ell. a. \text{iswrite}_\ell(a) \wedge \text{isread}_\ell(b) \wedge \text{hb}(a, b) && \text{(ConsistentRFdom)} \\
&\forall a, b. \text{rf}(b) = a \implies \exists \ell. v. \text{iswrite}_{\ell,v}(a) \wedge \text{isread}_{\ell,v}(b) \wedge \neg \text{hb}(a, b) && \text{(ConsistentRF)} \\
&\forall a, b. \text{rf}(b) = a \wedge (\text{mode}(a) = \text{na} \vee \text{mode}(b) = \text{na}) \implies \text{hb}(a, b) && \text{(ConsistentRFna)} \\
&\forall a, b. \text{rf}(b) = a \wedge \text{isSeqCst}(b) \implies \text{isc}(a, b) \vee \neg \text{isSeqCst}(a) \wedge (\forall x. \text{isc}(x, b) \implies \neg \text{hb}(a, x)) && \text{(RestrSCReads)} \\
&\quad \not\exists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), \text{rf}(a)) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentRRR)} \\
&\quad \not\exists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), a) \wedge \text{iswrite}(a) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentWR)} \\
&\quad \not\exists a, b. \text{hb}(a, b) \wedge \text{mo}(b, \text{rf}(a)) \wedge \text{iswrite}(b) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentRW)} \\
&\forall a. \text{isrmw}(a) \wedge \text{rf}(a) \neq \perp \implies \text{mo}(\text{rf}(a), a) \wedge \not\exists c. \text{mo}(\text{rf}(a), c) \wedge \text{mo}(c, a) && \text{(AtomicRMW)} \\
&\quad \forall a, b, \ell. \text{lab}(a) = \text{lab}(b) = \Lambda(\ell) \implies a = b && \text{(ConsistentAlloc)}
\end{aligned}$$

where $\text{iswrite}_{\ell,v}(a) \stackrel{\text{def}}{=} \exists X, v_{\text{old}}. \text{lab}(a) \in \{W_X(\ell, v), \text{RMW}_X(\ell, v_{\text{old}}, v)\}$ $\text{iswrite}_\ell(a) \stackrel{\text{def}}{=} \exists v. \text{iswrite}_{\ell,v}(a)$

$\text{isread}_{\ell,v}(a) \stackrel{\text{def}}{=} \exists X, v_{\text{new}}. \text{lab}(a) \in \{R_X(\ell, v), \text{RMW}_X(\ell, v, v_{\text{new}})\}$ etc.

$\text{rsElem}(a, b) \stackrel{\text{def}}{=} \text{sameThread}(a, b) \vee \text{isrmw}(b)$

$\text{rseq}(a) \stackrel{\text{def}}{=} \{a\} \cup \{b \mid \text{rsElem}(a, b) \wedge \text{mo}(a, b) \wedge (\forall c. \text{mo}(a, c) \wedge \text{mo}(c, b) \implies \text{rsElem}(a, c))\}$

$\text{sw} \stackrel{\text{def}}{=} \{(a, b) \mid \text{mode}(a) \in \{\text{rel}, \text{rel_acq}, \text{sc}\} \wedge \text{mode}(b) \in \{\text{acq}, \text{rel_acq}, \text{sc}\} \wedge \text{rf}(b) \in \text{rseq}(a)\}$

$\text{hb} \stackrel{\text{def}}{=} (\text{sb} \uplus \text{sw})^+$

$\text{hb}_\ell \stackrel{\text{def}}{=} \{(a, b) \in \text{hb} \mid \text{iswrite}_\ell(a) \wedge \text{iswrite}_\ell(b)\}$

$X_{\text{SeqCst}} \stackrel{\text{def}}{=} \{(a, b) \in X \mid \text{isSeqCst}(a) \wedge \text{isSeqCst}(b)\}$

$\text{isc}(a, b) \stackrel{\text{def}}{=} \text{iswrite}_{\text{locs}(b)}(a) \wedge \text{sc}(a, b) \wedge \not\exists c. \text{sc}(a, c) \wedge \text{sc}(c, b) \wedge \text{iswrite}_{\text{locs}(b)}(c)$

Figure 3. Axioms satisfied by consistent C11 executions, Consistent(\mathcal{A} , lab, sb, rf, mo, sc).

$c : W(\ell, 1) \xrightarrow{\text{rf}} a : R(\ell, 1)$ $\uparrow \text{mo}$ $d : W(\ell, 2) \xrightarrow{\text{rf}} b : R(\ell, 2)$ violates CoherentRR	$c : W(\ell, 2) \xrightarrow{\text{mo}} a : W(\ell, 1)$ $\text{rf} \searrow$ $b : R(\ell, 2)$ violates CoherentWR	$c : W(\ell, 1) \xrightarrow{\text{mo}} a : R(\ell, 1)$ $\text{mo} \searrow$ $b : W(\ell, 2)$ violates CoherentRW	$a \xrightarrow{\text{rf}} b$ means $a = \text{rf}(b)$ $a \xrightarrow{\text{mo}} b$ means $\text{mo}(a, b)$ $a \xrightarrow{\text{hb}} b$ means $\text{hb}(a, b)$
-------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 4. Sample executions violating coherency conditions (Batty et al. 2011).

Basic ingredients of execution graph consistency

1. SC-per-location (a.k.a. coherence)
2. Release/acquire synchronization
3. Global conditions on SC accesses

Basic ingredients of execution graph consistency

1. SC-per-location (a.k.a. coherence)
2. Release/acquire synchronization
3. Global conditions on SC accesses

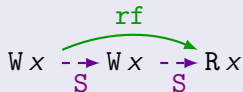
Sequential Consistency (SC)

Definition (Declarative definition of SC, Lamport '79)

G is *SC-consistent* if there exists a relation S s.t. the following hold:

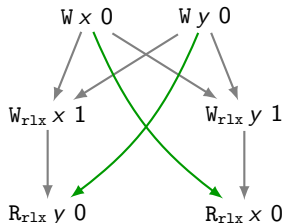
- ▶ S is a total order on the events of G .
- ▶ $(po \cup rf); S = \emptyset$.
- ▶ If $\langle a, b \rangle \in rf$ then there does not exist $c \in W_{loc}(a)$ such that $\langle a, c \rangle \in S$ and $\langle c, b \rangle \in S$.

Namely, the following is *disallowed*:



SC: Example

```
x = y = 0
x :=rlx 1;   || y :=rlx 1;
a :=rlx y; // 0 || b :=rlx x; // 0
```



program order

reads from

not SC-consistent!

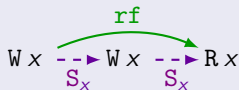
SC-per-location

Definition (SC-per-location)

G satisfies *SC-per-location* if for every location x , there exists a relation S_x s.t. the following hold:

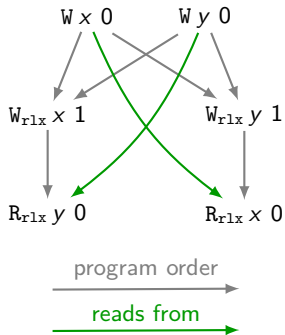
- ▶ S_x is a total order on the events of G that access x .
- ▶ $(po \cup rf); S_x = \emptyset$.
- ▶ If $\langle a, b \rangle \in rf$ then there does not exist $c \in W_x$ such that $\langle a, c \rangle \in S_x$ and $\langle c, b \rangle \in S_x$.

Namely, the following is *disallowed*:



SC-per-location: Example 1

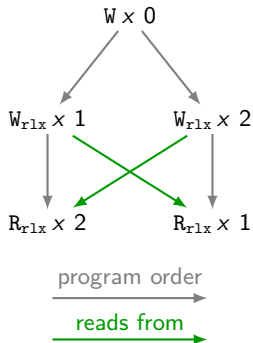
```
x = y = 0
x :=rlx 1;   || y :=rlx 1;
a :=rlx y; // 0 || b :=rlx x; // 0
```



satisfies SC-per-location!

SC-per-location: Example 2

```
x = 0  
x :=rlx 1;      || x :=rlx 2;  
a := xrlx; // 2 || b := xrlx; // 1
```



does not satisfy SC-per-location!

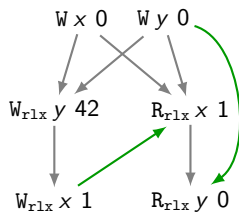
Release/Acquire synchronization

SC-per-location is often *too weak*:

- ▶ It does not support the message passing idiom:

Message passing (MP)

```
y :=rlx 42; || a :=rlx x; // 1  
x :=rlx 1; || b :=rlx y; // 0
```



- ▶ Also: we cannot implement locks.

Synchronization in C/C++11 through examples

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1; ||     print(y);
      || }
```


Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1; ||     print(y);
      || }
```

Synchronization in C/C++11 through examples

1

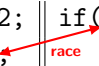
```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;  ← race print(y);
        || }

```

Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1; ||     print(y);
      ||     }
      ||
```



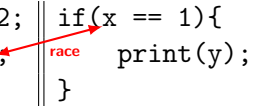
2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
x =rlx 1; ||     print(y);
          ||     }
          ||
```

Synchronization in C/C++11 through examples

1

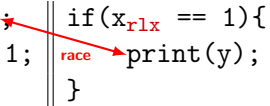
```
int y = 0;
int x = 0;
y = 42;
x = 1;
if(x == 1){
    print(y);
}
```



A diagram illustrating a race condition. Two vertical double lines represent parallel execution paths. The left path contains the code `y = 42;` followed by `x = 1;`. The right path contains `if(x == 1){` followed by `print(y);` followed by `}`. A red arrow labeled "race" points from the `x = 1;` line to the `print(y);` line, indicating that the value of `x` is not guaranteed to be 1 when `print(y)` is executed.

2

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_rlx 1;
if(x_rlx == 1){
    print(y);
}
```



A diagram illustrating a race condition in an atomic context. Two vertical double lines represent parallel execution paths. The left path contains the code `y = 42;` followed by `x =_rlx 1;`. The right path contains `if(x_rlx == 1){` followed by `print(y);` followed by `}`. A red arrow labeled "race" points from the `x =_rlx 1;` line to the `print(y);` line, indicating that the value of `x` is not guaranteed to be 1 when `print(y)` is executed, despite the use of relaxed atomicity.

Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;   ||     print(y);
         ||     }
         ||
```

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
x =rlx 1; ||     print(y);
         ||     }
         ||
```

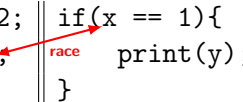
3

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xacq == 1){
x =rel 1; ||     print(y);
         ||     }
         ||
```

Synchronization in C/C++11 through examples

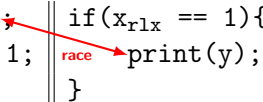
1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;   ||   print(y);
         ||   }
         ||
         ||
```



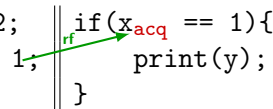
2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
x =rlx 1; ||   print(y);
         ||   }
         ||
```



3

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xacq == 1){
x =rel 1; ||   print(y);
         ||   }
         ||
```



Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;   ||   print(y);
         ||   }
         ||
         ||   race
```

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xr1x == 1){
x =r1x 1; ||   print(y);
         ||   }
         ||
         ||   race
```

3

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xacq == 1){
x =rel 1; ||   print(y);
         ||   }
         ||
         ||   rf
         ||   sw
```

Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1; ||   print(y);
      ||   }
      ||
      ||
```

race

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xr1x == 1){
x =r1x 1; ||   print(y);
      ||   }
      ||
```

race

3

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xacq == 1){
x =rel 1; ||   print(y);
      ||   }
      ||
      ||
```

rf
sw

4

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xr1x == 1){
fencerel; ||   fenceacq;
x =r1x 1; ||   print(y);
      ||   }
      ||
```


Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42; || if(x == 1){
x = 1;   ||   print(y);
         ||   }
         ||
         ||   race
```

2

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
x =rlx 1; ||   print(y);
         ||   }
         ||
         ||   race
```

3

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xacq == 1){
x =rel 1; ||   print(y);
         ||   }
         ||
         ||   rf
         ||   sw
```

4

```
int y = 0;
atomic<int> x = 0;
y = 42; || if(xrlx == 1){
fencerel; ||   fenceacq;
x =rlx 1; ||   print(y);
         ||   }
         ||
         ||   rf
```

Synchronization in C/C++11 through examples

1

```
int y = 0;
int x = 0;
y = 42;
x = 1;
if(x == 1){
    print(y);
}
```

2

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_r1x 1;
if(x_r1x == 1){
    print(y);
}
```

3

```
int y = 0;
atomic<int> x = 0;
y = 42;
x =_rel 1;
if(x_acq == 1){
    print(y);
}
```

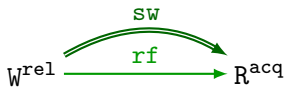
4

```
int y = 0;
atomic<int> x = 0;
y = 42;
fence_rel;
x =_r1x 1;
fence_acq;
if(x_r1x == 1){
    print(y);
}
```

The “synchronizes-with” relation

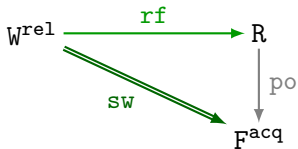
```

y = 42; ||| if(xacq == 1){
x =rel 1; |||   print(y);
                }
    
```



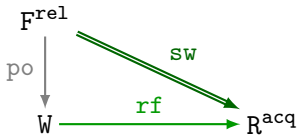
```

y = 42; ||| if(xrlx == 1){
x =rel 1; |||   fenceacq;
                } |||   print(y);
    
```



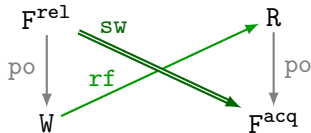
```

y = 42; ||| if(xacq == 1){
fencerel; |||   print(y);
x =rlx 1; ||| }
    
```

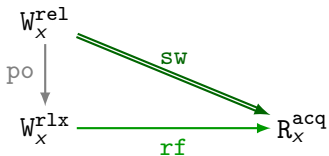
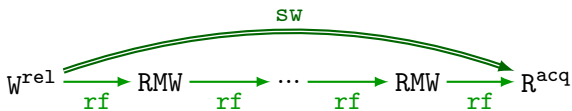


```

y = 42; ||| if(xrlx == 1){
fencerel; |||   fenceacq;
x =rlx 1; |||   print(y);
                }
    
```



The “synchronizes-with” relation: Release sequences



- ▶ Note: the latter case will be deprecated in C++20.

The “happens-before” relation

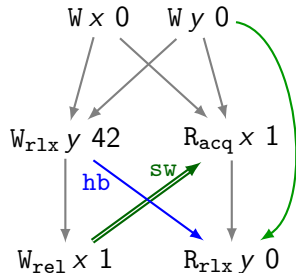
Definition (happens-before)

$$\frac{a \xrightarrow{\text{po}} b}{a \xrightarrow{\text{hb}} b}$$

$$\frac{a \xrightarrow{\text{sw}} b}{a \xrightarrow{\text{hb}} b}$$

$$\frac{a \xrightarrow{\text{hb}} b \quad b \xrightarrow{\text{hb}} c}{a \xrightarrow{\text{hb}} c}$$

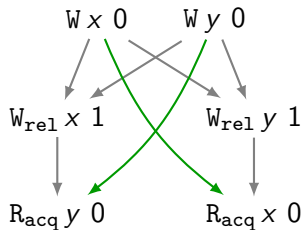
- ▶ **hb** should be acyclic.
- ▶ The SC-per-location orders should never contradict **hb**.



SC accesses and fences

Store buffer

```
x :=rel 1;      || y :=rel 1;  
a := yacq; // 0 || b := xacq; // 0
```

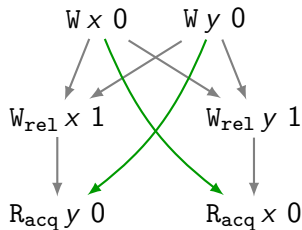


How to guarantee only SC behaviors (*i.e.*, $a = 1 \vee b = 1$)?

SC accesses and fences

Store buffer

```
x :=rel 1;      || y :=rel 1;  
a := yacq; // 0 || b := xacq; // 0
```



How to guarantee only SC behaviors (*i.e.*, $a = 1 \vee b = 1$)?

```
x :=sc 1;      || y :=sc 1;  
a := ysc;      || b := xsc;      ||  
x :=rlx 1;     || y :=rlx 1;     ||  
fencesc;      || fencesc;      ||  
a := yrlx;    || b := xrlx;    ||
```

SC semantics

- ▶ The semantics of SC atomics is the *most complicated* part of the model.

SC semantics

- ▶ The semantics of SC atomics is the *most complicated* part of the model.
- ▶ C/C++11 provides **too strong** semantics (a correctness problem!)

$$\begin{array}{l} a := x_{\text{acq}}; // 1 \\ b := y_{\text{sc}}; // 0 \end{array} \parallel \begin{array}{l} x :=_{\text{sc}} 1; \\ y :=_{\text{sc}} 1; \end{array} \parallel \begin{array}{l} c := y_{\text{acq}}; // 1 \\ d := x_{\text{sc}}; // 0 \end{array}$$

SC semantics

- ▶ The semantics of SC atomics is the *most complicated* part of the model.
- ▶ C/C++11 provides *too strong* semantics (a correctness problem!)

$$\begin{array}{l} a := x_{\text{acq}}; \text{ // } 1 \\ b := y_{\text{sc}}; \text{ // } 0 \end{array} \parallel \parallel x :=_{\text{sc}} 1; \parallel \parallel y :=_{\text{sc}} 1; \parallel \parallel \begin{array}{l} c := y_{\text{acq}}; \text{ // } 1 \\ d := x_{\text{sc}}; \text{ // } 0 \end{array}$$

- ▶ In addition, its semantics for SC fences is *too weak*.

$$\begin{array}{l} a := x_{\text{acq}}; \text{ // } 1 \\ \mathbf{fence}_{\text{sc}}; \\ b := y_{\text{acq}}; \text{ // } 0 \end{array} \parallel \parallel x :=_{\text{rel}} 1; \parallel \parallel y :=_{\text{rel}} 1; \parallel \parallel \begin{array}{l} c := y_{\text{acq}}; \text{ // } 1 \\ \mathbf{fence}_{\text{sc}}; \\ d := x_{\text{acq}}; \text{ // } 0 \end{array}$$

SC semantics

- ▶ The semantics of SC atomics is the *most complicated* part of the model.
- ▶ C/C++11 provides *too strong* semantics (a correctness problem!)

$$\begin{array}{l} a := x_{\text{acq}}; \text{ // 1} \\ b := y_{\text{sc}}; \text{ // 0} \end{array} \parallel \begin{array}{l} x :=_{\text{sc}} 1; \\ y :=_{\text{sc}} 1; \end{array} \parallel \begin{array}{l} c := y_{\text{acq}}; \text{ // 1} \\ d := x_{\text{sc}}; \text{ // 0} \end{array}$$

- ▶ In addition, its semantics for SC fences is *too weak*.

$$\begin{array}{l} a := x_{\text{acq}}; \text{ // 1} \\ \mathbf{fence}_{\text{sc}}; \\ b := y_{\text{acq}}; \text{ // 0} \end{array} \parallel \begin{array}{l} x :=_{\text{rel}} 1; \\ y :=_{\text{rel}} 1; \end{array} \parallel \begin{array}{l} c := y_{\text{acq}}; \text{ // 1} \\ \mathbf{fence}_{\text{sc}}; \\ d := x_{\text{acq}}; \text{ // 0} \end{array}$$

- ▶ Recently, the standard committee fixed the specification following:
[Repairing Sequential Consistency in C/C++11.
L, Vafeiadis, Kang, Hur, Dreyer. PLDI'17]

The “out-of-thin-air” problem

C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering

```
a := x; // 1    ||    b := y; // 1  
y := 1;        ||    x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**

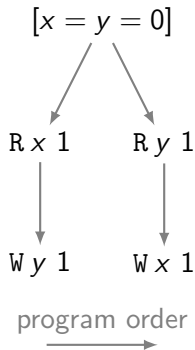
C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering

```
a := x; // 1     ||     b := y; // 1  
y := 1;         ||     x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**



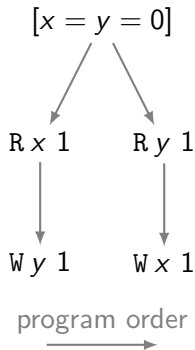
C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering

```
a := x; // 1    ||    b := y; // 1
y := 1;        ||    x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**



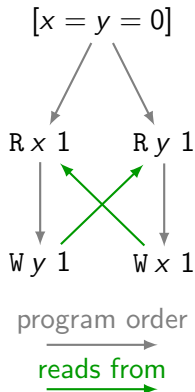
C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering

```
a := x; // 1     ||     b := y; // 1  
y := 1;         ||     x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**



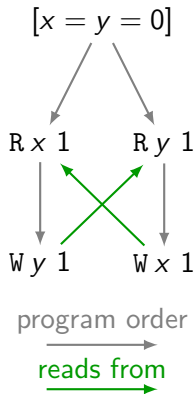
C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering

```
a := x; // 1     ||     b := y; // 1  
y := 1;         ||     x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**



C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

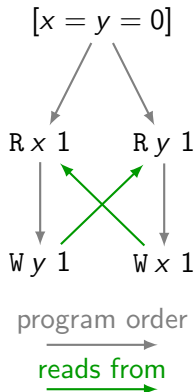
Load-buffering

```
a := x; // 1     ||     b := y; // 1  
y := 1;         ||     x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**

Load-buffering + data dependency

```
a := x; // 1     ||     b := y; // 1  
y := a;         ||     x := b;
```



C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering

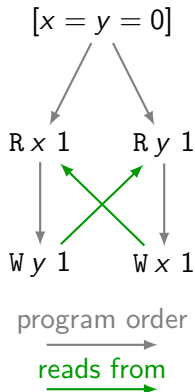
```
a := x; // 1     ||     b := y; // 1  
y := 1;         ||     x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**

Load-buffering + data dependency

```
a := x; // 1     ||     b := y; // 1  
y := a;         ||     x := b;
```

C/C++11 allows this behavior



C/C++11 is too weak

non-atomic □ **relaxed** □ release/acquire □ sc

Load-buffering

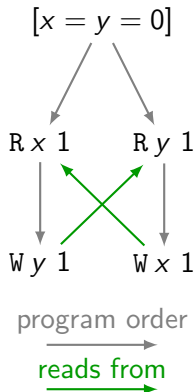
```
a := x; // 1     ||     b := y; // 1  
y := 1;           ||     x := b;
```

C/C++11 allows this behavior
because **POWER & ARM allow it!**

Load-buffering + data dependency

```
a := x; // 1     ||     b := y; // 1  
y := a;           ||     x := b;
```

C/C++11 allows this behavior
Values appear out-of-thin-air!
(no hardware/compiler exhibit this behavior)

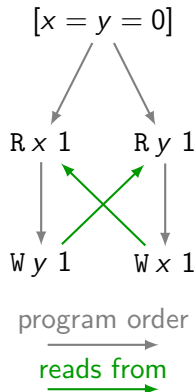


C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering + control dependency

```
a := x; // 1      ||      b := y; // 1
if (a = 1)        ||      if (b = 1)
  y := 1;         ||      x := 1;
```



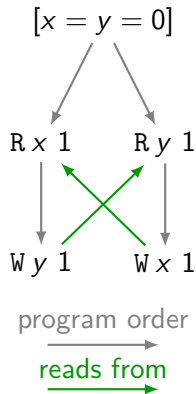
C/C++11 is too weak

non-atomic □ relaxed □ release/acquire □ sc

Load-buffering + control dependency

```
a := x; // 1      ||      b := y; // 1
if (a = 1)        ||      if (b = 1)
  y := 1;         ||      x := 1;
```

C/C++11 allows this behavior



C/C++11 is too weak

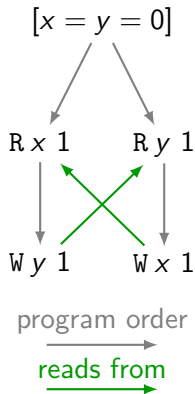
non-atomic □ relaxed □ release/acquire □ sc

Load-buffering + control dependency

```
a := x; // 1      ||      b := y; // 1
if (a = 1)        ||      if (b = 1)
  y := 1;         ||      x := 1;
```

C/C++11 allows this behavior

The DRF guarantee is broken!



C/C++11 is too weak

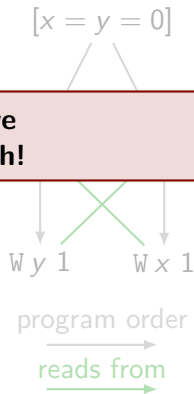
non-atomic □ relaxed □ release/acquire □ sc

Load-buffering + control dependency

```
a := x; // 1    ||    b := y; // 1  
if (a = 1)      ||    if (b = 1)
```

**The three examples have
the same execution graph!**

The DRF guarantee is broken!



The hardware solution

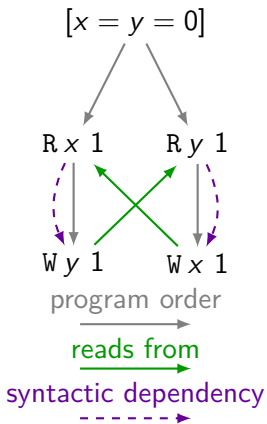
Keep track of **syntactic dependencies** and forbid **dependency cycles**.

Load-buffering

```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

Load-buffering + data dependency

```
a := x; // 1    ||    b := y; // 1
y := a;         ||    x := b;
```



The hardware solution

Keep track of **syntactic dependencies** and forbid **dependency cycles**.

Load-buffering

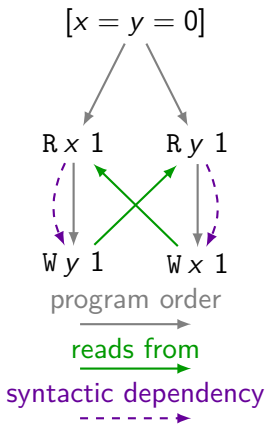
```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

Load-buffering + data dependency

```
a := x; // 1    ||    b := y; // 1
y := a;         ||    x := b;
```

Load-buffering + fake dependency

```
a := x; // 1    ||    b := y; // 1
y := a + 1 - a; ||    x := b;
```



The hardware solution

Keep track of **syntactic dependencies** and forbid **dependency cycles**.

Load-buffering

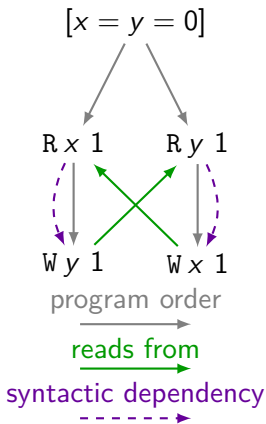
```
a := x; // 1    ||    b := y; // 1
y := 1;         ||    x := b;
```

Load-buffering + data dependency

```
a := x; // 1    ||    b := y; // 1
y := a;         ||    x := b;
```

Load-buffering + fake dependency

```
a := x; // 1    ||    b := y; // 1
y := a + 1 - a; ||    x := b;
```



This approach is not suitable for a programming language:
Compilers do not preserve syntactic dependencies.

The “out-of-thin-air” problem

C/C++11 is too weak

- ▶ Values might appear *out-of-thin-air*.
- ▶ The *DRF guarantee* is broken.

The C++14 standard states:

“Implementations should ensure that no “out-of-thin-air” values are computed that circularly depend on their own computation.”

- ▶ “Defined” by examples.

A straightforward solution

- ▶ Disallow $po \cup rf$ cycles!
- ▶ On weak hardware it carries a certain *implementation cost*.
[Ou and Demsky. Towards understanding the costs of avoiding out-of-thin-air results. OOPSLA'18]

A straightforward solution

- ▶ Disallow $po \cup rf$ cycles!
- ▶ On weak hardware it carries a certain *implementation cost*.
[Ou and Demsky. Towards understanding the costs of avoiding out-of-thin-air results. OOPSLA'18]

RC11 (Repaired C11) model

[L, Vafeiadis, Kang, Hur, Dreyer. PLDI'17]

- ▶ (Modified) compilation schemes are correct.
- ▶ DRF holds and no OOTA-values.
- ▶ Model checking [Kokologiannakis, L, Sagonas, Vafeiadis. POPL'18]
<http://plv.mpi-sws.org/rcmc/>

A straightforward solution

- ▶ Disallow $po \cup rf$ cycles!
- ▶ On weak hardware it carries a certain *implementation cost*.
[Ou and Demsky. Towards understanding the costs of avoiding out-of-thin-air results. OOPSLA'18]

RC11 (Repaired C11) model

[L, Vafeiadis, Kang, Hur, Dreyer. PLDI'17]

- ▶ (Modified) compilation schemes are correct.
 - ▶ DRF holds and no OOTA-values.
 - ▶ Model checking [Kokologiannakis, L, Sagonas, Vafeiadis. POPL'18]
<http://plv.mpi-sws.org/rcmc/>
-
- ▶ Solving the problem without changing the compilation schemes will require a *major revision* of the standard.

A 'promising' solution to OOTA

[Kang, Hur, L, Vafeiadis, Dreyer. POPL'17]

A 'promising' solution to OOTA

[Kang, Hur, L, Vafeiadis, Dreyer. POPL'17]

Key idea: Start with an operational interleaving semantics, but allow threads to **promise** to write in the future.

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; //0 || b = x; //0
```

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
┆
▶ x = 1;  || ▶ y = 1;
a = y; //0 || b = x; //0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
```

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
      x = 1;  ||  ▶ y = 1;
▶ a = y; // 0 ||  ▶ b = x; // 0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@5⟩
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
      |
x = 1; | y = 1;
      |
▶ a = y; // 0 | ▶ b = x; // 0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@5⟩
⟨y : 1@5⟩
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
	5

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; // 0  ▶ b = x; // 0
▶
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@5⟩
⟨y : 1@5⟩
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
	5

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; // 0 || b = x; // 0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	1
	5

- ▶ Global memory is a pool of messages of the form

$\langle location : value @ timestamp \rangle$

- ▶ Each thread maintains a *thread-local view* recording the last observed timestamp for every location

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; // 0 || b = x; // 0
```

▶

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@5⟩
⟨y : 1@5⟩
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
	5

Coherence Test

```
      x = 0
x := 1; || x := 2;
a = x; // 2 || b = x; // 1
```

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      ||      y = 1;
a = y; // 0 ||      b = x; // 0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
	5

Coherence Test

```
x = 0
x := 1; || x := 2;
a = x; // 2 || b = x; // 1
```

Memory

```
<x : 0@0>
```

T_1 's view

x
0

T_2 's view

x
0

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
      ||
  x = 1;  ||  y = 1;
  a = y; // 0 ||  b = x; // 0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@5⟩
⟨y : 1@5⟩
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
	5

Coherence Test

```
      x = 0
      ||
  x := 1;  ||  x := 2;
  a = x; // 2 ||  b = x; // 1
```

Memory

```
⟨x : 0@0⟩
⟨x : 1@5⟩
```

T_1 's view

x
0
5

T_2 's view

x
0

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
      x = y = 0
x = 1;  ||  y = 1;
a = y; // 0 || b = x; // 0
```

Memory

```
⟨x : 0@0⟩
⟨y : 0@0⟩
⟨x : 1@5⟩
⟨y : 1@5⟩
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	1
	5

Coherence Test

```
      x = 0
x := 1;  ||  x := 2;
a = x; // 2 || b = x; // 1
```

Memory

```
⟨x : 0@0⟩
⟨x : 1@5⟩
⟨x : 2@7⟩
```

T_1 's view

x
0
5

T_2 's view

x
1
7

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      ||      y = 1;
a = y; // 0 ||      b = x; // 0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
	5

Coherence Test

```
x = 0
x := 1;      ||      x := 2;
a = x; // 2  ▶ b = x; // 1
```

Memory

```
<x : 0@0>
<x : 1@5>
<x : 2@7>
```

T_1 's view

x
0
1
7

T_2 's view

x
0
7

Simple operational semantics for C11's relaxed accesses

Store-buffering

```
x = y = 0
x = 1;      ||      y = 1;
a = y; // 0 ||      b = x; // 0
```

Memory

```
<x : 0@0>
<y : 0@0>
<x : 1@5>
<y : 1@5>
```

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
	5

Coherence Test

```
x = 0
x := 1;  ||      x := 2;
a = x; // 2 ||      b = x; // 1
```

Memory

```
<x : 0@0>
<x : 1@5>
<x : 2@7>
```

T_1 's view

x
0
1
7

T_2 's view

x
0
7

Load-buffering

```
a := x; // 1 || x := y;  
y := 1;
```

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
▶ a := x; // 1 || ▶ x := y;  
y := 1;
```

Memory

$\langle x : 0@0 \rangle$

$\langle y : 0@0 \rangle$

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
▶ a := x; // 1 || ▶ x := y;  
y := 1;
```

Memory

$\langle x : 0@0 \rangle$

$\langle y : 0@0 \rangle$

$\langle y : 1@5 \rangle$

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
▶ a := x; // 1 || ▶ x := y;  
y := 1;
```

Memory

$\langle x : 0@0 \rangle$

$\langle y : 0@0 \rangle$

$\langle y : 1@5 \rangle$

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0
	5

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
▶ a := x; // 1 || x := y;  
y := 1; ▶
```

Memory

$\langle x : 0@0 \rangle$

$\langle y : 0@0 \rangle$

$\langle y : 1@5 \rangle$

$\langle x : 1@5 \rangle$

T_1 's view

x	y
0	0

T_2 's view

x	y
0	0
5	5

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
a := x; // 1
▶ y := 1;  ||  x := y;
▶
```

Memory

$\langle x : 0@0 \rangle$

$\langle y : 0@0 \rangle$

$\langle y : 1@5 \rangle$

$\langle x : 1@5 \rangle$

T_1 's view

x	y
0	0
5	

T_2 's view

x	y
0	0
5	5

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
a := x; // 1  
y := 1;      |      x := y;
```

Memory

$\langle x : 0@0 \rangle$

$\langle y : 0@0 \rangle$

$\langle y : 1@5 \rangle$

$\langle x : 1@5 \rangle$

T_1 's view

x	y
0	0
5	5

T_2 's view

x	y
0	0
5	5

- ▶ To model load-store reordering, we allow **“promises”**.
- ▶ At any point, a thread may promise to write a message in the future, allowing other threads to read from the promised message.

Promises

Load-buffering

```
a := x; // 1 || x := y;  
y := 1;      ▶
```

Memory

$\langle x : 0@0 \rangle$

$\langle y : 0@0 \rangle$

$\langle y : 1@5 \rangle$

$\langle x : 1@5 \rangle$

T_1 's view

x	y
0	0
5	5

T_2 's view

x	y
0	0
5	5

Load-buffering + dependency

```
a := x; // 1 || x := y;  
y := a;      ▶
```

Must not admit the same execution!

Promises

Load-buffering

```
a := x; // 1 || x := y;  
y := 1;      ▶
```

Load-buffering + dependency

```
a := x; // 1 || x := y;  
y := a;      ▶
```

Key Idea

A thread can only promise if it can perform the write anyway (even without having made the promise)

Certified promises

Thread-local certification

A thread can promise to write a message, if it can *thread-locally certify* that its promise will be fulfilled.

Certified promises

Thread-local certification

A thread can promise to write a message, if it can *thread-locally certify* that its promise will be fulfilled.

Load-buffering

$$\begin{array}{l} a := x; \text{ // } 1 \\ y := 1; \end{array} \parallel x := y;$$

T_1 **may promise** $y := 1$, since it is able to write $y := 1$ by itself.

Load buffering + fake dependency

$$\begin{array}{l} a := x; \text{ // } 1 \\ y := a + 1 - a; \end{array} \parallel x := y;$$

T_1 **may NOT promise** $y := 1$, since it is not able to write $y := 1$ by itself.

Load buffering + dependency

$$\begin{array}{l} a := x; \text{ // } 1 \\ y := a; \end{array} \parallel x := y;$$

Is this behavior possible?

```
a := x; // 1  
x := 1;
```

Is this behavior possible?

```
a := x; // 1  
x := 1;
```

No.

Suppose the thread promises $x := 1$. Then, once $a := x$ reads 1, the thread view is increased and so the promise cannot be fulfilled.

Is this behavior possible?

```
a := x; // 1 ||| y := x; ||| x := y;  
x := 1;
```

Is this behavior possible?

```
a := x; // 1 || y := x; || x := y;  
x := 1;
```

Yes. And the ARMv7 model allows it!

Is this behavior possible?

$$\begin{array}{l} a := x; \text{ // } 1 \\ x := 1; \end{array} \parallel \begin{array}{l} y := x; \\ x := y; \end{array}$$

Yes. And the ARMv7 model allows it!

This behavior can be also explained by sequentialization:

$$\begin{array}{l} a := x; \text{ // } 1 \\ x := 1; \\ y := x; \end{array} \parallel x := y; \quad \rightsquigarrow \quad \begin{array}{l} a := x; \text{ // } 1 \\ x := 1; \\ y := 1; \end{array} \parallel x := y;$$

Is this behavior possible?

$$\begin{array}{l}
 a := x; \text{ // } 1 \\
 x := 1;
 \end{array}
 \parallel
 \begin{array}{l}
 y := x; \\
 x := y;
 \end{array}$$

Yes. And the ARMv7 model allows it!

This behavior can be also explained by sequentialization:

$$\begin{array}{l}
 a := x; \text{ // } 1 \\
 x := 1; \\
 y := x;
 \end{array}
 \parallel
 x := y;
 \quad \rightsquigarrow \quad
 \begin{array}{l}
 a := x; \text{ // } 1 \\
 x := 1; \\
 y := 1;
 \end{array}
 \parallel
 x := y;$$

Is this behavior possible?

```
a := xrlx; // 42
y :=rlx a;

b := yrlx;
if (b = 42)
  c := "if";
else
  c := "else";
  b := 42;
x :=rlx b;
print (c); // prints "if"
```


Is this behavior possible?

```
a := xrlx; // 42
y :=rlx a;

b := yrlx;
if (b = 42)
    c := "if";
else
    c := "else";
b := 42;
x :=rlx b;
print (c); // prints "if"
```

Yes. And it can obtained by compiler optimizations!

We have extended this basic idea to handle:

- ▶ Atomic updates (e.g., CAS, fetch-and-add)
- ▶ Release/acquire fences and accesses
- ▶ Release sequences
- ▶ SC fences
- ▶ Plain accesses
(C11's non-atomics & Java's normal accesses)

Results

- ▶ No “out-of-thin-air” values
- ▶ DRF guarantees
- ▶ Compiler optimizations (incl. reorderings, eliminations)
- ▶ Efficient h/w mappings (x86-TSO, Power, ARM)

We have extended this basic idea to handle:

- ▶ Atomic updates (e.g., CAS, fetch-and-add)
- ▶ Release/acquire fences and accesses
- ▶ Release sequences
- ▶ SC fences
- ▶ Plain accesses
(C11's non-atomics & Java's normal accesses)



The **Coq**
proof assistant



Results

- ▶ No “out-of-thin-air” values
- ▶ DRF guarantees
- ▶ Compiler optimizations (incl. reorderings, eliminations)
- ▶ Efficient h/w mappings (x86-TSO, Power, ARM)

We have extended this basic idea to handle:

- ▶ Atomic updates (e.g., CAS, fetch-and-add)
- ▶ Release/acquire fences and accesses
- ▶ Release sequences
- ▶ SC fences
- ▶ Plain accesses
(C11's non-atomics & Java's normal accesses)

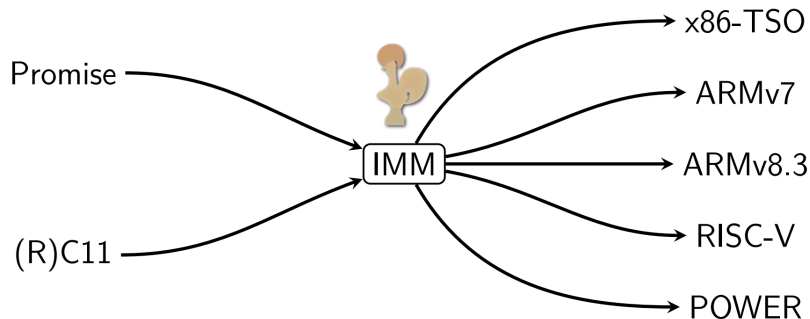


The **Coq**
proof assistant



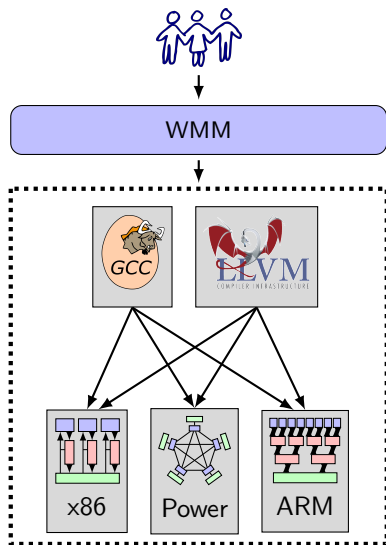
Results

- ▶ No “out-of-thin-air” values
- ▶ DRF guarantees
- ▶ Compiler optimizations (incl. reorderings, eliminations)
- ▶ Efficient h/w mappings (x86-TSO, Power, ARM)



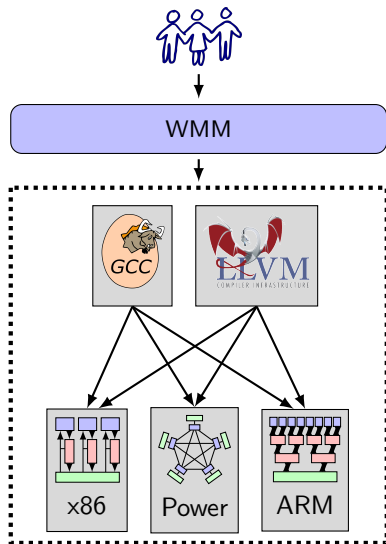
- ▶ A *common denominator* of existing models
- ▶ Formulated in the *declarative style*
- ▶ Simplifies *compilation correctness* proofs

Summary



- ▶ The challenges in designing a WMM.
- ▶ The C/C++11 model.
- ▶ C/C++11 is **broken**:
 - ▶ Most problems are **locally fixable**.
 - ▶ But ruling out **OOTA** requires an entirely different approach.
- ▶ The **promising model** may be the solution.

Summary



- ▶ The challenges in designing a WMM.
- ▶ The C/C++11 model.
- ▶ C/C++11 is **broken**:
 - ▶ Most problems are **locally fixable**.
 - ▶ But ruling out **OOTA** requires an entirely different approach.
- ▶ The **promising model** may be the solution.

Thank you!

<http://www.cs.tau.ac.il/~orilahav/>