

Software verification under weak memory consistency

Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)

January 2016

Joint work with Soham Chakraborty, Derek Dreyer, Marko Doko, Nick Giannarakis, Ori Lahav, Chinmay Narayan, Joseph Tassarotti, Aaron Turon.

Sequential consistency

Sequential consistency (SC):


- ▶ The standard model for concurrency.
- ▶ Interleave each thread's atomic accesses.
- ▶ Almost all verification work assumes it.

Initially, $x = y = 0$.

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

In SC, this program cannot return $a = b = 0$.

Store buffering in x86-TSO



Initially, $x = y = 0$.

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

This program can return $a = b = 0$.

OG = Hoare logic + rule for parallel composition


$$\frac{\{P_1\} c_1 \{Q_1\} \quad \{P_2\} c_2 \{Q_2\} \quad \text{the two proofs are } \textit{non-interfering}}{\{P_1 \wedge P_2\} c_1 \parallel c_2 \{Q_1 \wedge Q_2\}}$$

Non-interference

$R \wedge P \vdash R\{u/x\}$ for every:

- ▶ assertion R in the proof outline of one thread
- ▶ assignment $x := u$ with precondition P in the proof outline of the other thread

Standard OG is unsound for WM



Standard OG is unsound for WM

$$\begin{array}{c} \{a \neq 0\} \\ \{a \neq 0\} \\ x := 1; \\ \{x \neq 0\} \\ a := y \\ \{x \neq 0\} \\ \{a \neq 0 \vee b \neq 0\} \end{array} \parallel \begin{array}{c} \{a \neq 0\} \\ \{\top\} \\ y := 1; \\ \{y \neq 0\} \\ b := x \\ \{y \neq 0 \wedge (a \neq 0 \vee b = x)\} \end{array}$$

Standard OG is unsound for WM

$$\begin{array}{c} \{a \neq 0\} \\ \{a \neq 0\} \\ x := 1; \\ \{x \neq 0\} \\ a := y \\ \{x \neq 0\} \\ \{a \neq 0 \vee b \neq 0\} \end{array} \parallel \begin{array}{c} \{a \neq 0\} \\ \{\top\} \\ y := 1; \\ \{y \neq 0\} \\ b := x \\ \{y \neq 0 \wedge (a \neq 0 \vee b = x)\} \end{array}$$

To regain soundness, strengthen the non-inference check.
 \implies OGRA: Owicki-Gries for release-acquire (ICALP'15)

- ▶ Introduction to the C11 weak memory model
 - ▶ Release-acquire synchronization
 - ▶ Per-location coherence
- ▶ Reasoning about WMM using program logics
 - ▶ RSL (relaxed separation logic)
 - ▶ FSL (fenced separation logic)
 - ▶ GPS (ghosts, protocols and separation)
- ▶ Avoiding to reason about WMM
 - ▶ Use reduction theorems

The C11 memory model: Atomics

Two types of locations

**Ordinary
(Non-Atomic)**

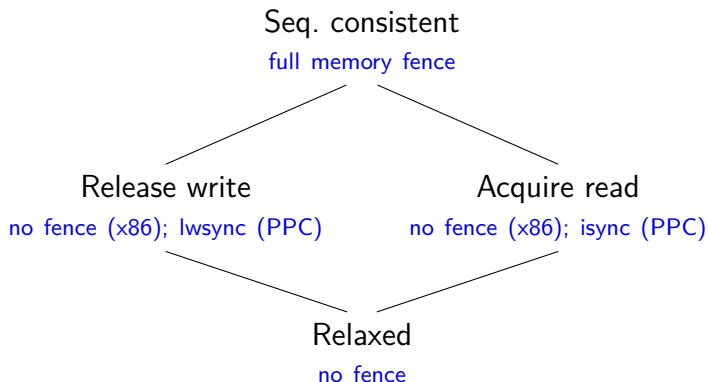
Races are **errors**

Atomic

Welcome to the
expert mode

The C11 memory model: Hierarchy of atomics

A spectrum of atomic accesses:




Explicit primitives for fences

Store buffering in C11

Initially $x = y = 0$.

```
x.store(1, rlx);    || y.store(1, rlx);  
print(y.load(rlx)); || print(x.load(rlx));
```

Can print 00 with the following execution:




Release-acquire synchronization

Initially $a = x = 0$.

```
 $a = 5;$   
 $x.\text{store}(1, \text{release});$  || while ( $x.\text{load}(\text{acq}) == 0$ );  
 $\text{print}(a);$ 
```

One possible execution:




Happens before: $hb = (po \cup sw)^+$

Coherence

Programs with a single shared variable behave as under SC.

$$\begin{array}{l} x.\text{store}(1, rlx); \\ a = x.\text{load}(rlx); \end{array} \parallel \begin{array}{l} x.\text{store}(2, rlx); \\ b = x.\text{load}(rlx); \end{array}$$

The outcome $a = 2 \wedge b = 1$ is forbidden.



- ▶ Modification order, mo_x , total order of writes to x .
- ▶ Reads-before : $rb \triangleq (rf^{-1}; mo) \cap (\neq)$
- ▶ Coherence : $hb \cup rf_x \cup mo_x \cup rb_x$ is acyclic for all x .

Relaxed program logics

- ▶ RSL (relaxed separation logic, OOPSLA'13)
- ▶ FSL (fenced separation logic, VMCAI'16)
- ▶ GPS (ghosts & protocols, OOPSLA'14, PLDI'15)

Separation logic

Key concept of *ownership* :

- ▶ Resourceful reading of Hoare triples.

$$\{P\} C \{Q\}$$

- ▶ To access a non-atomic location, you must own it:

$$\begin{aligned} & \{x \mapsto v\} * x \{t. t = v \wedge x \mapsto v\} \\ & \{x \mapsto v\} * x = v'; \{x \mapsto v'\} \end{aligned}$$

- ▶ Disjoint parallelism:

$$\frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\}}{\{P_1 * P_2\} C_1 || C_2 \{Q_1 * Q_2\}}$$

Rules for release/acquire accesses

Ownership transfer by rel-acq synchronizations.

- ▶ Atomic allocation \rightsquigarrow pick loc. invariant Q .

$$\{Q(v)\} x = \text{alloc}(v); \{\mathbf{W}_Q(x) * \mathbf{R}_Q(x)\}$$

- ▶ Release write \rightsquigarrow give away permissions.

$$\{Q(v) * \mathbf{W}_Q(x)\} x.\text{store}(v, \text{rel}); \{\mathbf{W}_Q(x)\}$$

- ▶ Acquire read \rightsquigarrow gain permissions.

$$\{\mathbf{R}_Q(x)\} t = x.\text{load}(\text{acq}); \{Q(t) * \mathbf{R}_{Q[t:=\text{emp}]}(x)\}$$

Release-acquire synchronization: message passing

Initially $a = x = 0$. Let $J(v) \triangleq v = 0 \vee \&a \mapsto 5$.

$\{\&a \mapsto 0 * \mathbf{W}_J(x)\}$		$\{\mathbf{R}_J(x)\}$
$a = 5;$		while ($x.\text{load}(acq) == 0$);
$\{\&a \mapsto 5 * \mathbf{W}_J(x)\}$		$\{\&a \mapsto 5\}$
$x.\text{store}(\text{release}, 1);$		$\text{print}(a);$
$\{\mathbf{W}_J(x)\}$		$\{\&a \mapsto 5\}$

PL consequences:

Ownership transfer works!

Relaxed accesses

Basically, disallow ownership transfer.

- ▶ Relaxed reads:

$$\{\mathbf{R}_Q(x)\} t := x.\text{load}(rlx) \{\mathbf{R}_Q(x) \wedge (Q(t) \neq \text{false})\}$$

- ▶ Relaxed writes:

$$\frac{Q(v) = \text{emp}}{\{\mathbf{W}_Q(x)\} x.\text{store}(v, rlx) \{\mathbf{W}_Q(x)\}}$$

Unsound because of dependency cycles!


Dependency cycles

Initially $x = y = 0$.

```
if (x.load(rlx) == 1) || if (y.load(rlx) == 1)
    y.store(1, rlx);      x.store(1, rlx);
```

C11 allows the outcome $x = y = 1$.

Justification:



Dependency cycles

Initially $x = y = 0$.

```
if (x.load(rlx) == 1) || if (y.load(rlx) == 1)
    y.store(1, rlx);      x.store(1, rlx);
```

C11 allows the outcome $x = y = 1$.

What goes wrong:

Non-relational invariants are unsound.

$$x = 0 \wedge y = 0$$

The DRF-property does not hold.

Dependency cycles

Initially $x = y = 0$.

```
if (x.load(rlx) == 1) || if (y.load(rlx) == 1)
    y.store(1, rlx);      x.store(1, rlx);
```

C11 allows the outcome $x = y = 1$.

How to fix this:


Don't use relaxed writes



Strengthen the model


Incorrect message passing

```
int a; atomic_int x = 0;
( a = 5;                               || if (x.load(rlx) ≠ 0){
  x.store(1, rlx);                       ||   print(a); } )
```



Message passing with C11 memory fences

```
int a; atomic_int x = 0;
( a = 5;                               || if (x.load(rlx) ≠ 0){
  fence(release);                       ||   fence(acq);
  x.store(1, rlx);                       ||   print(a); } )
```



- ▶ Introduce two ‘modalities’ in the logic

$$\{P\} \text{ fence}(\text{release}) \{\Delta P\}$$

$$\{\mathbf{W}_Q(x) * \Delta Q(v)\} x.\text{store}(v, r/x) \{\mathbf{W}_Q(x)\}$$

$$\{\mathbf{R}_Q(x)\} t := x.\text{load}(r/x) \{\mathbf{R}_{Q[t:=\text{emp}]}(x) * \nabla Q(t)\}$$

$$\{\nabla P\} \text{ fence}(\text{acq}) \{P\}$$


Reasoning about fences

Let $Q(v) \triangleq v = 0 \vee \&a \mapsto 5$.

$$\left(\begin{array}{l} \{ \&a \mapsto 0 * \mathbf{W}_Q(x) * \mathbf{R}_Q(x) \} \\ \{ \&a \mapsto 0 * \mathbf{W}_Q(x) \} \\ a = 5; \\ \{ \&a \mapsto 5 * \mathbf{W}_Q(x) \} \\ \text{fence}(\text{release}); \\ \{ \Delta(\&a \mapsto 5) * \mathbf{W}_Q(x) \} \\ x.\text{store}(1, r/x); \\ \{ \top \} \end{array} \parallel \begin{array}{l} t = x.\text{load}(r/x); \\ \{ \nabla(t = 0 \vee \&a \mapsto 5) \} \\ \text{if } (t \neq 0) \\ \quad \text{fence}(\text{acq}); \\ \quad \{ \&a \mapsto 5 \} \\ \quad \mathbf{print}(a); \\ \{ \top \} \end{array} \right)$$

GPS: A better logic for release-acquire

Three key features:

- ▶ Location protocols
- ▶ Ghost state/tokens 
- ▶ Escrows for ownership transfer

Example (Racy message passing)

Initially, $x = y = 0$.

$$\begin{array}{l} x.\text{store}(1, rlx); \\ y.\text{store}(1, rel); \end{array} \parallel \begin{array}{l} x.\text{store}(1, rlx); \\ y.\text{store}(1, rel); \end{array} \parallel \begin{array}{l} t = y.\text{load}(acq); \\ t' = x.\text{load}(rlx); \end{array}$$

Cannot get $t = 1 \wedge t' = 0$.

Racy message passing in GPS

Protocol for x : **A**: $x = 0$ \longrightarrow **B**: $x = 1$

Protocol for y : **C**: $y = 0$ \longrightarrow **D**: $y = 1 \wedge x.st \geq \mathbf{B}$

Acquire reads gain knowledge, not ownership.

$$\left. \begin{array}{l} \{x.st \geq \mathbf{A} \wedge y.st \geq \mathbf{C}\} \\ x.store(1, rlx); \\ \{x.st \geq \mathbf{B} \wedge y.st \geq \mathbf{C}\} \\ y.store(1, rel); \\ \{x.st \geq \mathbf{B} \wedge y.st \geq \mathbf{D}\} \end{array} \right\| \left\{ \begin{array}{l} \{x.st \geq \mathbf{A} \wedge y.st \geq \mathbf{C}\} \\ t = y.load(acq); \\ \left\{ \begin{array}{l} t = 0 \wedge x.st \geq \mathbf{A} \\ \vee t = 1 \wedge x.st \geq \mathbf{B} \end{array} \right\} \\ t' = x.load(rlx); \\ \{t = 0 \vee (t = 1 \wedge t' = 1)\} \end{array} \right\}$$

GPS ghosts and escrows

To gain ownership, we use ghost state & escrows.

$$\frac{P * P \Rightarrow \text{false}}{Q \Rightarrow \mathbf{Esc}(P, Q)} \quad \frac{}{\mathbf{Esc}(P, Q) * P \Rightarrow Q}$$

Example (Message passing using escrows)

Invariant for x : $x = 0 \vee \mathbf{Esc}(K, \&a \mapsto 7)$.

$\{\&a \mapsto 0\}$		$\{K\}$
$a = 7;$		if ($x.\text{load}(acq) \neq 0$)
$\{\&a \mapsto 7\}$		$\{K * \mathbf{Esc}(K, \&a \mapsto 7)\}$
$\{\mathbf{Esc}(K, \&a \mapsto 7)\}$		$\{\&a \mapsto 7\}$
$x.\text{store}(1, \text{rel});$		print (a);

Avoiding weak memory reasoning

- ▶ DRF theorem
- ▶ Enough fences to guarantee SC

Theorem (DRF)

If $\llbracket \text{Prg} \rrbracket_{\text{SC}}$ contains no data races on non-SC accesses, then $\llbracket \text{Prg} \rrbracket_{\text{C11}} = \llbracket \text{Prg} \rrbracket_{\text{SC}}$.


- ▶ Requires strengthened semantics for relaxed accesses.
- ▶ Program logics that disallow data races are trivially sound.
- ▶ What about racy programs?

C11's SC-fences

- ▶ The strongest fence instruction provided by C11 is **SC-fence**.
- ▶ Can also be used to regain sequential consistency.

Example (Store Buffering)

```
x = y = 0
x.store(1, rlx); || y.store(1, rlx);
a=y.load(rlx); || b=x.load(rlx);
```




C11's SC-fences

- ▶ The strongest fence instruction provided by C11 is **SC-fence**.
- ▶ Can also be used to regain sequential consistency.

Example (Store Buffering)

```
x = y = 0
x.store(1, rlx); || y.store(1, rlx);
fence(sc);      || fence(sc);
a=y.load(rlx); || b=x.load(rlx);
```




Inconsistent: $(F_{sc} \times F_{sc}) \cap (po^?; (hb \cup rf \cup mo \cup rb); po^?)$ is cyclic.

SC-fences are overly weak

Initially, $x = y = 0$.

```
x.store(1, rlx); || a = x.load(rlx); || c = y.load(rlx); || y.store(1, rlx);  
                || b = y.load(rlx); || d = x.load(rlx);
```

The outcome $a = c = 1 \wedge b = d = 0$ is allowed.




SC-fences are overly weak

Initially, $x = y = 0$.

$x.\text{store}$	$\left\ \begin{array}{l} a = x.\text{load}(rlx); \\ \text{fence}(sc); \\ b = y.\text{load}(rlx); \end{array} \right\ $	$\left\ \begin{array}{l} c = y.\text{load}(rlx); \\ \text{fence}(sc); \\ d = x.\text{load}(rlx); \end{array} \right\ $	$y.\text{store}$
$(1, rlx);$			$(1, rlx);$


The outcome $a = c = 1 \wedge b = d = 0$ is allowed.



$(F_{sc} \times F_{sc}) \cap (po^?; (hb \cup rf \cup mo \cup rb); po^?)$ is acyclic.


Our suggestion

- ▶ Model SC-fences as **release-acquire atomic updates** of a distinguished **fence location**.
- ▶ RA semantics enforces all fence events to be ordered by *hb*.




Our suggestion

- ▶ Model SC-fences as **release-acquire atomic updates** of a distinguished **fence location**.
- ▶ RA semantics enforces all fence events to be ordered by *hb*.




Our suggestion

- ▶ Model SC-fences as **release-acquire atomic updates** of a distinguished **fence location**.
- ▶ RA semantics enforces all fence events to be ordered by *hb*.



Our suggestion

- ▶ Model SC-fences as **release-acquire atomic updates** of a distinguished **fence location**.
- ▶ RA semantics enforces all fence events to be ordered by *hb*.



Inconsistent: $R_{rlx}(x, 0)$ reads an overwritten value.

Theorem

If in a program P ,

- ▶ *all shared accesses are **atomic** (relaxed or stronger), and*
- ▶ *there is a **fence** between every two shared accesses to different shared variables,*

then P has only SC behaviors.

Advanced reduction theorem

- ▶ For x86-TSO, it suffices to have a fence between every **racy write** and subsequent **racy read**.
- ▶ Generally, C11 requires more fences.

```
x := e;  
fence();  
r := y
```

Theorem (Advanced reduction to SC, simplified)

If in a *client-server* program P ,

- ▶ all shared accesses are *release/acquire*, and
- ▶ there is a *fence* between every store to a shared location and subsequent shared location load,

then P has only SC behaviors.

Applying the theorem to RCU

```
rcu_quiescent_state() :  
  rc[get_my_tid()] := gc;  
  fence()
```

```
rcu_thread_offline() :  
  rc[get_my_tid()] := 0;  
  fence()
```

```
rcu_thread_online() :  
  rc[get_my_tid()] := gc;  
  fence()
```

```
synchronize_rcu() :  
  gc := gc + 1;  
  fence();  
  for i := 1 to N do  
    wait (rc[i] ∈ {0,gc})
```

Reasoning about weak memory is challenging and often unavoidable.

Two approaches:

- ▶ Use **relaxed program logics** to reason about weak memory.
- ▶ Use **reduction theorems** to avoid such reasoning.

Relaxed program logics also useful for understanding weak memory.