

Software verification under weak memory consistency

Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)

Dagstuhl, 2015-05-04


Goal:

- ▶ Understand & exploit weak memory consistency (using PLV techniques)

Questions:

- ▶ How do the different models relate?
- ▶ What program transformations are allowed? (Compiler optimisations)
- ▶ What proof techniques are sound? (Robustness theorems, program logics)
- ▶ Is there a good weak memory model?

Store buffering in x86-TSO



Initially, $x = y = 0$.

```
x := 1;    || y := 1;  
print(y); || print(x);
```


This program can also print 00.

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:




Causality 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 $x = y = 1$.



Relaxed accesses don't synchronize.

Sequentialisation is invalid

Initially $a = x = y = 0$.

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

Remarks:

- ▶ Non-atomic read axiom: $rf \cap (na \times na) \subseteq hb$.
- ▶ This program is not racy according to C11.
- ▶ The only possible output is $a = 1$ and $x = y = 0$.

Consequences:

- ▶ The model is not monotone.
- ▶ Sequentialization $C_1 || C_2 \rightsquigarrow C_1; C_2$ is invalid.

Theorem (DRF for a model M)

If $\llbracket C \rrbracket_{\text{sf}}$ is data race free, then $\llbracket C \rrbracket_M = \llbracket C \rrbracket_{\text{sc}}$.

- ▶ Holds for all ‘reasonable’ memory models M .
- ▶ But not for C11 with RLX accesses due to causality cycles.
Recall the example. Initially $x = y = 0$.

$$\begin{array}{l} \mathbf{if} (x.\text{load}(rlx) == 1) \\ \quad y.\text{store}(1, rlx); \end{array} \parallel \begin{array}{l} \mathbf{if} (y.\text{load}(rlx) == 1) \\ \quad x.\text{store}(1, rlx); \end{array}$$

- ▶ Strengthen model for relaxed accesses. . .

Program logics for WM

- ▶ Standard SC logics are unsound.
- ▶ What can be done.

Owicki-Gries method (1976)

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

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

Program logics for C11 release-acquire:

- ▶ Relaxed separation logic (V. & Narayan, OOPSLA'13)
- ▶ GPS (Turon et al., OOPSLA'14)
- ▶ OGRA (Lahav & V., ICALP'15)

Program logics for TSO:

- ▶ Rely/guarantee for TSO (Ridge, VSTTE'10)
- ▶ Verifying TSO programs (Jacobs, 2014)
- ▶ iCAP-TSO (Sieczkowski et al., ESOP'15)

Relaxed separation logic (simplified)

Joint work with Chinmay Narayan, published at OOPSLA'13.

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.

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

- ▶ Acquire read \rightsquigarrow gain permissions.

$$\{\mathbf{R}_Q(x)\} t = x.\text{load}(\text{acq}); \{Q(t)\}$$

Message passing in RSL


Let $Q(v) \triangleq v = 0 \vee \&a \mapsto 7$.

$$\left(\begin{array}{l} \{T\} \\ \mathbf{atomic_int} \ x = 0; \ \mathbf{int} \ a = 0; \\ \{\&a \mapsto 0 * \mathbf{W}_Q(x) * \mathbf{R}_Q(x)\} \\ \left(\begin{array}{l} \{\&a \mapsto 0 * \mathbf{W}_Q(x)\} \\ a = 7; \\ \{\&a \mapsto 7 * \mathbf{W}_Q(x)\} \\ x.\mathbf{store}(1, \mathit{release}); \\ \{T\} \end{array} \right) \parallel \left(\begin{array}{l} t = x.\mathbf{load}(\mathit{acq}); \\ \{t = 0 \vee \&a \mapsto 7\} \\ \mathbf{if} \ (t \neq 0) \{\&a \mapsto 7\} \\ \mathbf{print}(a); \\ \{T\} \end{array} \right) \\ \{T\} \end{array} \right)$$

GPS: A better logic for release-acquire

Joint work with Aaron Turon and Derek Dreyer, at OOPSLA'14.

Three key features:

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

Example (Racy message passing)

Initially, $x = y = 0$.

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

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.

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

Some unresolved questions

- ▶ OOTA
- ▶ LLVM memory model
- ▶ MM classification

Resolving out of thin air (OOA) reads

Current status:

- ▶ C11 permits OOTA because of relaxed accesses.
- ▶ OOTA invalidates type safety & DRF.
- ▶ Hardware implementations don't generate OOTA.

Attempts at alternative models:

- ▶ Rule out RLX accesses and/or strengthen C11 axioms.
- ▶ Local transformations + RA.
- ▶ Operational
- ▶ Event structures

The LLVM memory model

Almost C11, but not exactly!

- ▶ C11: non-atomic races \rightsquigarrow program error.
- ▶ Irrelevant read introduction is unsound.
- ▶ LLVM, however, does introduce reads.
- ▶ LLVM: read-write race \rightsquigarrow read gets bogus value.

Example (Unsound sequence of transformations)


```
if (flag)
    a = X;
else
    a = 0;
lock();
b = X;

(1)  $\rightsquigarrow$ 
t = X;
a = flag ? t : 0;
lock();
b = X;


(2)  $\rightsquigarrow$ 
t = X;
a = flag ? t : 0;
lock();
b = t;
```

Hierarchy of relatively strong WMM

Theorems & Conjectures:




Hierarchy of relatively strong WMM



Theorems & Conjectures:

- ▶ $WRA = PSI$ if no WW races.


Hierarchy of relatively strong WMM



Theorems & Conjectures:

- ▶ $WRA = PSI$ if no WW races.
- ▶ $TSO = WB$ and $PSI = SI$ for server-client programs.


Hierarchy of relatively strong WMM



Theorems & Conjectures:

- ▶ $\text{WRA} = \text{PSI}$ if no WW races.
- ▶ $\text{TSO} = \text{WB}$ and $\text{PSI} = \text{SI}$ for server-client programs.
- ▶ $\text{Power-RA} = \text{WB}$.


Hierarchy of relatively strong WMM



Theorems & Conjectures:

- ▶ $WRA = PSI$ if no WW races.
- ▶ $TSO = WB$ and $PSI = SI$ for server-client programs.
- ▶ $Power-RA = WB$.
- ▶ Operational model for MP?


Hierarchy of relatively strong WMM



Theorems & Conjectures:

- ▶ $WRA = PSI$ if no WW races.
- ▶ $TSO = WB$ and $PSI = SI$ for server-client programs.
- ▶ $Power-RA = WB$.
- ▶ Operational model for MP?
- ▶ OGRA sound for RA, not for WRA.


Hierarchy of relatively strong WMM



Theorems & Conjectures:

- ▶ $WRA = PSI$ if no WW races.
- ▶ $TSO = WB$ and $PSI = SI$ for server-client programs.
- ▶ $Power-RA = WB$.
- ▶ Operational model for MP?
- ▶ OGRA sound for RA, not for WRA.
- ▶ RSL and GPS sound for WRA?


Hierarchy of relatively strong WMM



Theorems & Conjectures:

- ▶ $WRA = PSI$ if no WW races.
- ▶ $TSO = WB$ and $PSI = SI$ for server-client programs.
- ▶ $Power-RA = WB$.
- ▶ Operational model for MP?
- ▶ OGRA sound for RA, not for WRA.
- ▶ RSL and GPS sound for WRA?
- ▶ Restricted forms of auxiliary state sound for (W)RA and WB.

Hierarchy of relatively strong WMM



Theorems & Conjectures:

- ▶ $WRA = PSI$ if no WW races.
- ▶ $TSO = WB$ and $PSI = SI$ for server-client programs.
- ▶ $Power-RA = WB$.
- ▶ Operational model for MP?
- ▶ OGRA sound for RA, not for WRA.
- ▶ RSL and GPS sound for WRA?
- ▶ Restricted forms of auxiliary state sound for (W)RA and WB.

Q: Explain the MM differences with program logic?

Summary:

- ▶ C11 is broken, but largely fixable.
- ▶ OG is unsound for TSO.

Further work:

- ▶ Resolving OOTA
- ▶ Formalizing the LLVM model
- ▶ MM hierarchy