

# Relaxed Separation Logic

Tutorial @ POPL'14

Viktor Vafeiadis

MPI-SWS

20 January 2014

## Part I. Weak memory models

1. Introduction to relaxed concurrency
2. The C11 relaxed memory model

## Part II. Relaxed program logics

3. Concurrent separation logic
4. Relaxed separation logic
5. RSL extensions (ongoing)



<http://www.mpi-sws.org/~viktor/rsl/>

Sequential consistency (SC):

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

Initially,  $x = y = 0$ .

$$\begin{array}{l} x := 1; \\ \textit{print}(y); \end{array} \parallel \begin{array}{l} y := 1; \\ \textit{print}(x); \end{array}$$

In SC, this program can print 01, 10, or 11.

Sequential consistency (SC):

- ▶ Interleave each thread's atomic accesses.
- ▶ The standard model for concurrency.

▶ A But SC is invalidated by:


- ▶ F Hardware implementations
- ▶ Compiler optimisations

Initial

$$\begin{array}{l} x := 1; \\ \textit{print}(y); \end{array} \parallel \begin{array}{l} y := 1; \\ \textit{print}(x); \end{array}$$

In SC, this program can print 01, 10, or 11.

## Store buffering in x86-TSO



Initially,  $x = y = 0$ .

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


This program can also print 00.

## IRIW: Not just reordering

Initially,  $x = y = 0$ .

$$x := 1 \parallel y := 1 \parallel \begin{array}{l} \textit{print}(x); \\ \textit{print}(y); \end{array} \parallel \begin{array}{l} \textit{print}(y); \\ \textit{print}(x); \end{array}$$

Both threads can print 10.



Initially,  $x = y = 0$ .

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

The program can print 010.

## **Justification:**

The compiler may perform CSE:  
Load  $x$  into a temporary  $t$   
and print  $t$ ,  $y$ , and  $t$ .

All *sane* memory models satisfy the DRF property:

### Theorem (DRF-property)

*If  $\llbracket Prg \rrbracket_{SC}$  contains no data races, then*  
 $\llbracket Prg \rrbracket_{Relaxed} = \llbracket Prg \rrbracket_{SC}$ .

- ▶ Program logics that disallow data races are trivially sound.
- ▶ What about *racy* programs?



# The C11 memory model

Two types of locations: ordinary and atomic

- ▶ Races on ordinary accesses  $\rightsquigarrow$  error

A spectrum of atomic accesses:


- ▶ Relaxed  $\rightsquigarrow$  no fence
- ▶ Consume reads  $\rightsquigarrow$  no fence, but preserve deps
- ▶ Release writes  $\rightsquigarrow$  no fence (x86); lwsync (PPC)
- ▶ Acquire reads  $\rightsquigarrow$  no fence (x86); isync (PPC)
- ▶ Seq. consistent  $\rightsquigarrow$  full memory fence

Primitives for explicit fences

- ▶ Execution = set of events & a few relations:
  - ▶ sb: sequenced before
  - ▶ rf: reads-from map
  - ▶ mo: memory order per location
  - ▶ sc: seq.consistency order
  - ▶ sw [derived]: synchronized with
  - ▶ hb [derived]: happens before
- ▶ Axioms constraining the *consistent* executions.
- ▶  $\mathcal{C}(\text{prog})$  = set of all consistent exec's.
- ▶ if all  $\mathcal{C}(\text{prog})$  race-free on ordinary accesses,  $\llbracket \text{prog} \rrbracket = \mathcal{C}(\text{prog})$ ; otherwise,  $\llbracket \text{prog} \rrbracket = \text{"error"}$

# Release-acquire synchronization: message passing in C11

```
atomic_int x = 0; int a = 0;  
( a = 7;  
  x.store(1, release); ||| if (x.load(acq) != 0)  
                             print(a); )
```



happens-before  $\stackrel{\text{def}}{=} (\text{sequenced-before} \cup \text{sync-with})^+$

$\text{sync-with}(a, b) \stackrel{\text{def}}{=} \text{reads-from}(b) = a \wedge \text{release}(a) \wedge \text{acquire}(b)$

## Example (SB)

Initially,  $x = y = 0$ .

```
x.store(1, release);    || y.store(1, release);  
t = y.load(acquire);   || t' = x.load(acquire);
```

This program may produce  $t = t' = 0$ .

## Example (IRIW)

Initially,  $x = y = 0$ .

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

May produce  $a = c = 1 \wedge b = d = 0$ .

## Example (Read-Read Coherence)

Initially,  $x = 0$ .

$$x.store(1, rel); \parallel x.store(2, rel); \parallel a=x.load(acq); \parallel c=x.load(acq); \\ b=x.load(acq); \parallel d=x.load(acq);$$

Cannot get  $a = d = 1 \wedge b = c = 2$ .

- ▶ Plus similar WR, RW, WW coherence properties.
- ▶ Ensure SC behaviour for a single variable.
- ▶ Also guaranteed for relaxed atomics  
(the weakest kind of atomics in C11).

## Part II

### Relaxed Program Logics

- ▶ Concurrent separation logic
- ▶ Relaxed separation logic
- ▶ RSL extensions ([ongoing](#))

Key concept of *ownership* :

- ▶ Resourceful reading of Hoare triples.

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

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

$$\{x \mapsto v\} *x \{t. t = v \wedge x \mapsto v\}$$

$$\{x \mapsto v\} *x = v'; \{x \mapsto v'\}$$

- ▶ Disjoint parallelism:

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

# 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.

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

- ▶ Acquire read  $\rightsquigarrow$  gain permissions.

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



# Message passing in RSL

Let  $Q(v) \stackrel{\text{def}}{=} v = 0 \vee \&a \mapsto 7$ .

$$\begin{array}{c} \{ \text{true} \} \\ \text{atomic\_int } x = 0; \text{ int } a = 0; \\ \{ \&a \mapsto 0 * \mathbf{W}_Q(x) * \mathbf{R}_Q(x) \} \\ \left( \begin{array}{l} \{ \&a \mapsto 0 * \mathbf{W}_Q(x) \} \\ a = 7; \\ \{ \&a \mapsto 7 * \mathbf{W}_Q(x) \} \\ x.\text{store}(1, \text{release}); \\ \{ \text{true} \} \end{array} \right) \left\| \begin{array}{l} t = x.\text{load}(\text{acq}); \\ \{ t = 0 \vee \&a \mapsto 7 \} \\ \text{if } (t \neq 0) \{ \&a \mapsto 7 \} \\ \quad \text{print}(a); \\ \quad \{ \text{true} \} \end{array} \right. \\ \{ \text{true} \} \end{array}$$

Write permissions can be duplicated:

$$\mathbf{W}_Q(\ell) \iff \mathbf{W}_Q(\ell) * \mathbf{W}_Q(\ell)$$

Read permissions cannot, but may be split:

$$\mathbf{R}_{Q_1 * Q_2}(\ell) \iff \mathbf{R}_{Q_1}(\ell) * \mathbf{R}_{Q_2}(\ell)$$

```
a = 7;           || t = x.load(acq); || t' = x.load(acq);  
b = 8;           || if (t ≠ 0)           || if (t' ≠ 0)  
x.store(1, rel); ||   print(a);           ||   print(b);
```

Basically, disallow ownership transfer.

- ▶ Relaxed reads:

$$\{\mathbf{R}_Q(x)\} x.load(rlx) \{y. Q(y) \neq \text{false}\}$$

- ▶ Relaxed writes:

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

Basically, disallow ownership transfer.

- ▶ Relaxed reads:

$$\{\mathbf{R}_Q(x)\} x.load(rlx) \{y. Q(y) \neq \text{false}\}$$

- ▶ Relaxed writes:

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

Unfortunately *not sound* because of a bug in the C11 memory model.


## Dependency cycles in C11

Initially  $x = y = 0$ .

**if** ( $x.load(rlx) == 1$ )  $y.store(1, rlx);$     **if** ( $y.load(rlx) == 1$ )  $x.store(1, rlx);$

The formal C11 model allows  $x = y = 1$ .

### Justification:



Relaxed accesses  
don't synchronize

## Dependency cycles in C11

Initially  $x = y = 0$ .

**if** ( $x.load(rlx) == 1$ )  $y.store(1, rlx);$     **if** ( $y.load(rlx) == 1$ )  $x.store(1, rlx);$

The formal C11 model allows  $x = y = 1$ .

### **What goes wrong:**

Non-relational invariants are unsound.

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

The DRF-property does not hold.

## Dependency cycles in C11

Initially  $x = y = 0$ .

**if** ( $x.load(rlx) == 1$ )  $y.store(1, rlx);$     **if** ( $y.load(rlx) == 1$ )  $x.store(1, rlx);$

The formal C11 model allows  $x = y = 1$ .

### How to fix this:

Don't use relaxed writes

∨

Require *acyclic*( $sb \cup rf$ ).  
(Disallow RW reordering.)

# Compare and swap (CAS)

- ▶ New assertion form,  $P := \dots \mid \mathbf{R}_Q^U(x)$ .
- ▶ Duplicable,  $\mathbf{R}_Q^U(x) \iff \mathbf{R}_Q^U(x) * \mathbf{R}_Q^U(x)$ .

$$X \in \{rel, rlx\} \Rightarrow Q(v) \equiv \text{emp}$$

$$X \in \{acq, rlx\} \Rightarrow Q(v') \equiv \text{emp}$$

$$P * Q(v) \Rightarrow Q(v') * R[v/z]$$

$$\{P\} \text{ x.load}(Y) \{z. z \neq v \Rightarrow R\}$$

---

$$\{\mathbf{R}_Q^U(x) * \mathbf{W}_Q(x) * P\} \text{ x.CAS}(v, v', X, Y) \{z. R\}$$



# Mutual exclusion locks

**Let**  $Q_J(v) \stackrel{\text{def}}{=} (v = 0 \wedge \text{emp}) \vee (v = 1 \wedge J)$   
 $\text{Lock}(x, J) \stackrel{\text{def}}{=} \mathbf{W}_{Q_J}(x) * \mathbf{R}_{Q_J}^U(x)$

$\text{new-lock}() \stackrel{\text{def}}{=}$   
 $\{J\}$   
 $\text{res} = \mathbf{alloc}(1)$   
 $\{\text{Lock}(\text{res}, J)\}$


$\text{unlock}(x) \stackrel{\text{def}}{=}$   
 $\{J * \text{Lock}(x, J)\}$   
 $x.\text{store}(1, \text{rel})$   
 $\{\text{Lock}(x, J)\}$

$\text{lock}(x) \stackrel{\text{def}}{=}$   
 $\{\text{Lock}(x, J)\}$   
**repeat**  
 $\{\text{Lock}(x, J)\}$   
 $y = x.\text{CAS}(1, 0, \text{acq}, \text{rlx})$   
 $\left\{ \text{Lock}(x, J) * \begin{pmatrix} y=0 \wedge \text{emp} \\ \vee y=1 \wedge J \end{pmatrix} \right\}$   
**until**  $y \neq 0$   
 $\{J * \text{Lock}(x, J)\}$

# GPS: A better logic for release-acquire

Joint work with Aaron Turon and Derek Dreyer.

Three key features:

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

## Example (Racy message passing)

Initially,  $x = y = 0$ .

```
x.store(1, rlx); || x.store(1, rlx); || t = y.load(acq);  
y.store(1, rel); || y.store(1, rel); || t' = x.load(rlx);
```

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

# Racy message passing in GPS

Protocol for  $x$ :  $\mathbf{A}: x = 0$   $\longrightarrow$   $\mathbf{B}: x = 1$

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

Acquire reads gain knowledge, not ownership.

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

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

<pre> {&amp;a ↦ 0} a = 7; {&amp;a ↦ 7} {Esc(K, &amp;a ↦ 7)} x.store(1, rel); </pre>	<pre> {K} if (x.load(acq) ≠ 0)   {K * Esc(K, &amp;a ↦ 7)}   {&amp;a ↦ 7}   print(a); </pre>
---	---


# Incorrect message passing

```
int a; atomic_int x = 0;
( a = 7;
  x.store(1, rlx); || if (x.load(rlx) ≠ 0){
                       print(a); } )
```



# Message passing with C11 memory fences

```
int a; atomic_int x = 0;
( a = 7;           || if (x.load(rlx) ≠ 0){
  fence(release); ||   fence(acq);
  x.store(1, rlx); ||   print(a); } )
```



# Reasoning about fences

Joint work with Marko Doko (in progress)

Let  $Q(v) \stackrel{\text{def}}{=} v = 0 \vee \&a \mapsto 7$ .

$$\left( \begin{array}{l} \{ \&a \mapsto 0 * \mathbf{W}_Q(x) * \mathbf{R}_Q(x) \} \\ \{ \&a \mapsto 0 * \mathbf{W}_Q(x) \} \\ a = 7; \\ \{ \&a \mapsto 7 * \mathbf{W}_Q(x) \} \\ \text{fence}(\text{release}) \\ \{ \Delta(\&a \mapsto 7) * \mathbf{W}_Q(x) \} \\ x.\text{store}(1, rlx); \\ \{ \text{true} \} \end{array} \parallel \begin{array}{l} t = x.\text{load}(rlx); \\ \{ \nabla(t = 0 \vee \&a \mapsto 7) \} \\ \text{if } (t \neq 0) \\ \quad \text{fence}(\text{acq}); \\ \quad \{ \&a \mapsto 7 \} \\ \quad \mathbf{print}(a); \} \\ \{ \text{true} \} \end{array} \right)$$

$$\begin{array}{lll} \{ P \} & \text{fence}(\text{release}) & \{ \Delta P \} \\ \{ \nabla P \} & \text{fence}(\text{acq}) & \{ P \} \\ \{ \mathbf{R}_Q(x) \} & x.\text{load}(rlx) & \{ y. \nabla Q(y) \} \\ \{ \mathbf{W}_Q(x) * \Delta Q(v) \} & x.\text{store}(v, rlx) & \{ \text{true} \} \end{array}$$


# Release-consume synchronization

Initially  $a = x = 0$ .

```
 $a = 5;$   
 $x.\text{store}(\text{release}, \&a);$  ||  $t = x.\text{load}(\text{consume});$   
 $\text{if } (t \neq 0) \text{ print}(*t);$ 
```

This program cannot crash nor print 0.

## Justification:





## Release-consume synchronization

Initially  $a = x = 0$ . Let  $J(t) \stackrel{\text{def}}{=} t = 0 \vee t \mapsto 5$ .

$\{\&a \mapsto 0 * \mathbf{W}_J(x)\}$		$\{\mathbf{R}_J(x)\}$
$a = 5;$		$t = x.\text{load}(\text{consume});$
$\{\&a \mapsto 5 * \mathbf{W}_J(x)\}$		$\{\nabla_t(t = 0 \vee t \mapsto 5)\}$
$x.\text{store}(\text{release}, \&a);$		<b>if</b> $(t \neq 0)$ $\text{print}(*t);$

This program cannot crash nor print 0.

Index the  $\nabla$  with program variable  $t$ .  
 $t$  data dependence  $\implies$  locally open  $\nabla_t$ .

Again, work in progress...

Take away:

Formal reasoning about C11 programs  
is not so difficult after all.

We're not quite there yet; there's still a lot to do:

Liveness, refinement, tool support, ...

Topics that were not covered:

- ▶ The soundness proof:  
Really interesting and fully mechanized in Coq.
- ▶ Found 4+1 bugs in the C11 model.  
Program logic as a debugging tool for WMM.