

# Program logics for relaxed consistency

UPMARC Summer School 2014

Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)

1st Lecture, 28 July 2014

## Part I. Weak memory models

1. Intro to relaxed memory consistency
2. The C11 memory model

## Part II. Program logics

3. Separation logic
4. Relaxed separation logic
5. GPS : ghosts & protocols
6. Advanced features



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

## Sequential consistency

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

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.

# Basic compiler optimisations break SC / TSO

Initially,  $x = y = 0$ .

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

Can the program print 010?

## Justification:


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

## IRIW: Not just store buffering

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.



## From IRIW to the store buffering example

Take the IRIW example:

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

Linearize twice (threads 1-3 and 2-4):

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

That's the store buffering example (with two extra print statements).



Initially,  $x = 0$ .

$$\begin{array}{l} x = 1; \parallel \textit{print}(x); \\ x = 2; \parallel \textit{print}(x); \end{array}$$

Cannot print 10 nor 20 nor 21.

- ▶ Programs with one shared variable have SC semantics.
- ▶ Ensured by the cache coherence protocol.

## Message passing

Initially,  $x = y = 0$ .

$x = 1;$		$print(y);$
[WW fence]		[RR fence]
$y = 1;$		$print(x);$

Cannot print 10.

- ▶ No fences needed on x86-TSO
- ▶ lwsync/isync on Power
- ▶ dmb/isync on ARM

## Understanding weak memory consistency

Read the architecture/language specs?

- ▶ Too informal, often wrong.

Read the formalisations?

- ▶ Fairly complex.

Run benchmarks / Litmus tests?

- ▶ Observe only subset of behaviours.

We need a better methodology...

# Which memory model?

Hardware or language models?

- ▶ Want to reason at “high level”
- ▶ TSO  $\rightsquigarrow$  good robustness theorems

C/C++ or Java?

- ▶ JMM is broken [Ševčík & Aspinall, ECOOP'08]
- ▶ So, only C11 left

Goals:

- ▶ Understand the memory model
- ▶ Verify intricate concurrent 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  $\stackrel{\text{def}}{=} (\text{sequenced-before} \cup \text{sync-with})^+$

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

# Rel-acq synchronization is weaker than SC

## 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,  $x = 0$ .

$$x.store(1, rel); \parallel x.store(2, rel); \parallel \begin{array}{l} a=x.load(acq); \\ b=x.load(acq); \end{array} \parallel \begin{array}{l} c=x.load(acq); \\ d=x.load(acq); \end{array}$$

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

Today:

- ▶ Separation logic
- ▶ Relaxed separation logic

## When should we care about relaxed memory?

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?

## Separation logic assertions

Assertions describe the heap ( $Loc \rightarrow Val$ ):

- ▶ emp: the empty heap
- ▶  $l \mapsto v$ : a cell at address  $l$  containing  $v$

$$h \models l \mapsto v \iff h = \{l \mapsto v\}$$

- ▶  $P * Q$ : separating conjunction

$$h \models P * Q \iff \exists h_1 h_2. h = h_1 \uplus h_2 \wedge h_1 \models P \wedge h_2 \models Q$$

- ▶  $\wedge, \vee, \neg, \text{true}, \text{false}, \forall, \exists$ : as usual

# The separating conjunction

Some basic properties:

- ▶  $*$  is commutative & associative.
- ▶  $P * \text{emp} \iff \text{emp} * P \iff P$
- ▶  $l \mapsto v * l \mapsto v' \implies \text{false}$

Useful for describing inductive data structures:

- ▶  $\text{list}(x) \stackrel{\text{def}}{=} (x = 0 \wedge \text{emp}) \vee \exists y. x \mapsto y * \text{list}(y)$
- ▶  $\text{ls}(x, z) \stackrel{\text{def}}{=} (x = z \wedge \text{emp}) \vee \exists y. x \mapsto y * \text{ls}(y, z)$
- ▶  $\text{tree}(x) \stackrel{\text{def}}{=} (x = 0 \wedge \text{emp}) \vee \exists y, z. x \mapsto y * x+1 \mapsto z * \text{tree}(y) * \text{tree}(z)$

Key concept of *ownership* :

- ▶ Resourceful reading of Hoare triples

$$\frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\}}{\{P_1 * P_2\} C_1 || C_2 \{Q_1 * Q_2\}} \quad (\text{Par})$$

$$\frac{\{P\} C \{Q\}}{\{P * R\} C \{Q * R\}} \quad (\text{Frame})$$

- ▶ Ensure memory safety & race-freedom

## Separation logic rules for non-atomic accesses

- ▶ Allocation gives you permission to access  $x$ .

$$\{\text{emp}\} x = \text{alloc}(); \{\exists v. x \mapsto v\}$$

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

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


## Release-acquire synchronization: message passing

Initially  $a = x = 0$ .

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

This will always print 5.

### Justification:



Release-acquire  
synchronization



# Rules for release/acquire accesses

Relaxed separation logic [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.

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

- ▶ Acquire read  $\rightsquigarrow$  gain permissions.

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

# Message passing in RSL

Let  $Q(v) \stackrel{\text{def}}{=} v = 0 \vee \&a \mapsto 5$ .

$$\left( \begin{array}{c|c} \{true\} & \\ \mathbf{atomic\_int} \ x = 0; \ \mathbf{int} \ a = 0; & \\ \{ \&a \mapsto 0 * \mathbf{W}_Q(x) * \mathbf{R}_Q(x) \} & \\ \left( \begin{array}{c|c} \{ \&a \mapsto 0 * \mathbf{W}_Q(x) \} & \{ \mathbf{R}_Q(x) \} \\ a = 5; & \mathbf{while} \ (x.load(acq) == 0); \\ \{ \&a \mapsto 5 * \mathbf{W}_Q(x) \} & \{ \&a \mapsto 5 \} \\ x.store(1, release); & \mathbf{print}(a); \\ \{ true \} & \{ \&a \mapsto 5 \} \end{array} \right) & \\ \{true\} & \end{array} \right)$$

Write permissions can be duplicated:

$$\mathbf{W}_{Q_1}(l) \iff \mathbf{W}_{Q_1}(l) * \mathbf{W}_{Q_2}(l)$$

Read permissions cannot, but may be split:

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

```
a = 7;
b = 8;
x.store(1, rel);  $\parallel$  t = x.load(acq);
                    if (t  $\neq$  0)
                      print(a);  $\parallel$  t' = x.load(acq);
                    if (t'  $\neq$  0)
                      print(b);
```

Basically, disallow ownership transfer.

- ▶ Relaxed reads:

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

- ▶ Relaxed writes:

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

Basically, disallow ownership transfer.

- ▶ Relaxed reads:

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

- ▶ Relaxed writes:

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

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



Topics covered today:

- ▶ The C11 memory model
- ▶ Separation logic
- ▶ Relaxed separation logic

Tomorrow:

- ▶ Compare and swap
- ▶ GPS
- ▶ Advanced C11 features