

Formal reasoning about the C11 weak memory model

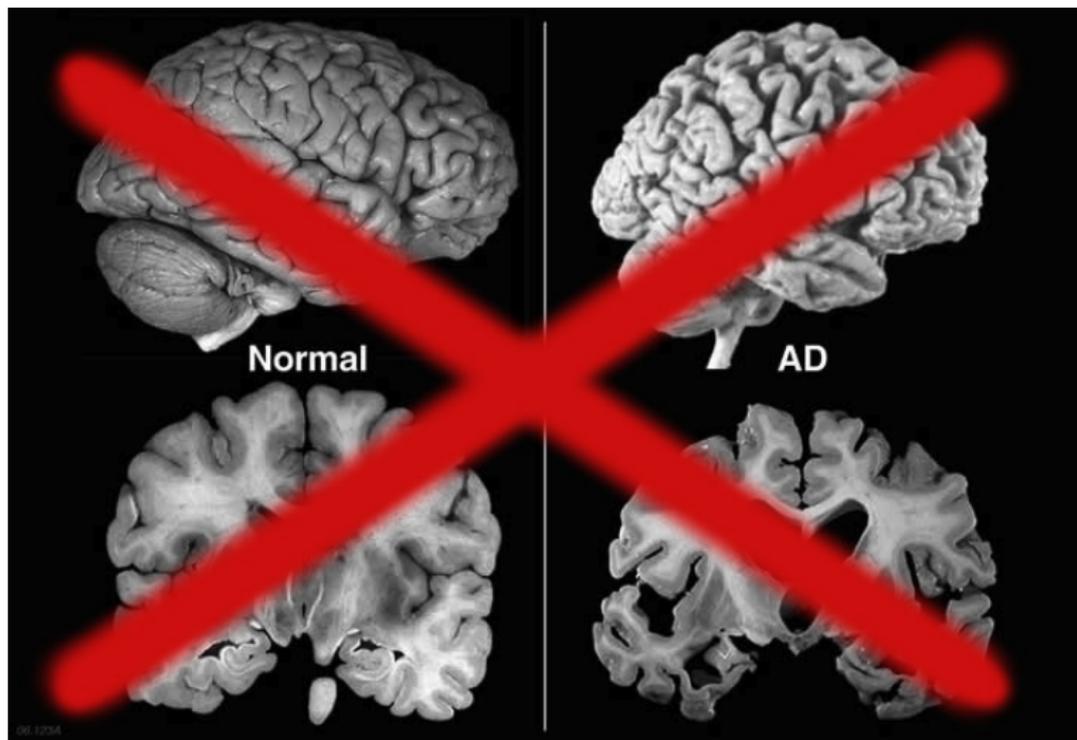
Invited talk @ CPP'15

Viktor Vafeiadis

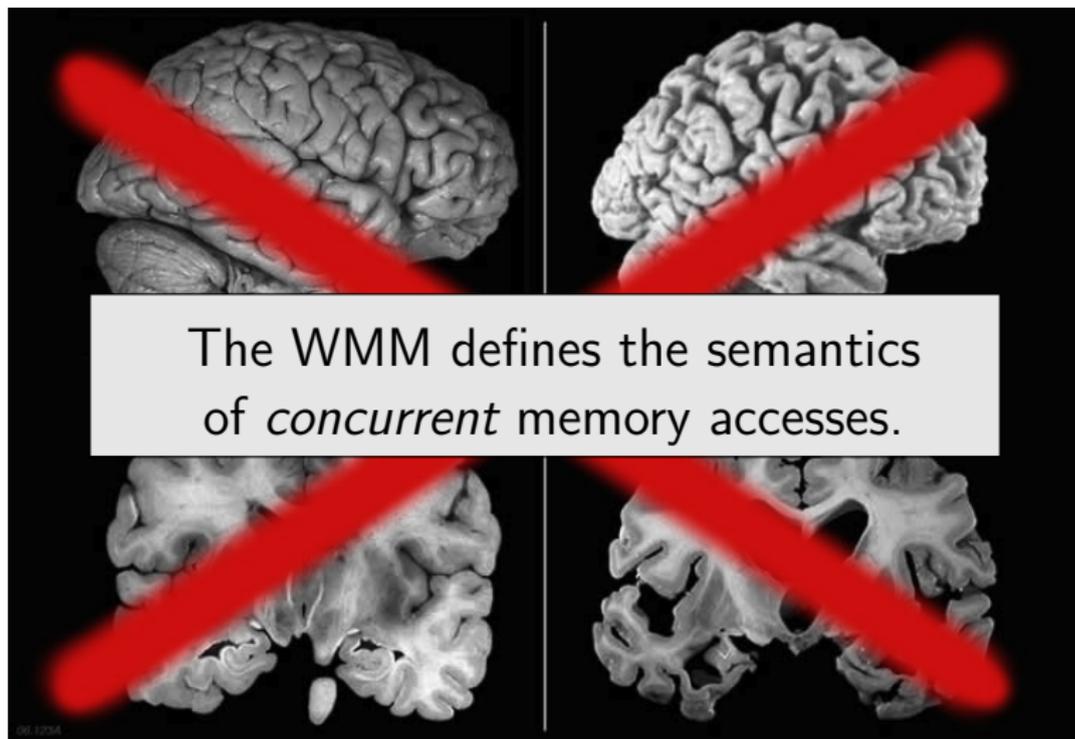
Max Planck Institute for Software Systems (MPI-SWS)

13 January 2015

What is a weak memory model?



What is a weak memory model?



Sequential consistency

Sequential consistency (SC):

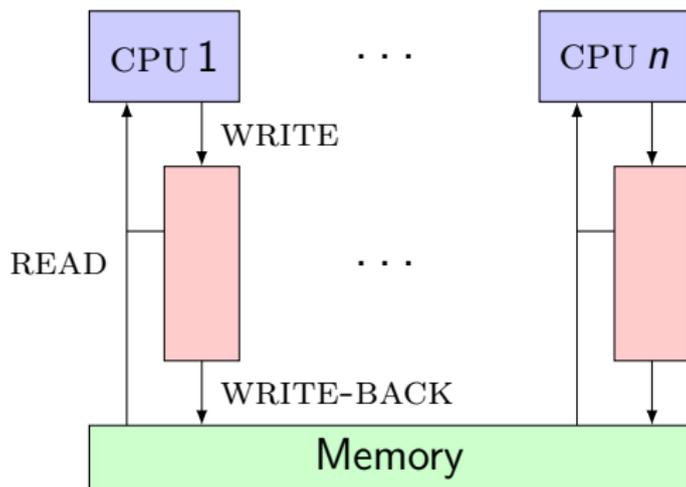
- ▶ The standard model for concurrency.
- ▶ Interleave each thread's atomic accesses.
- ▶ Almost all verification work assumes it.

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, $X = Y = 0$.

$$\begin{array}{l} X := 1; \\ a := Y \end{array} \parallel \begin{array}{l} Y := 1; \\ b := X \end{array}$$

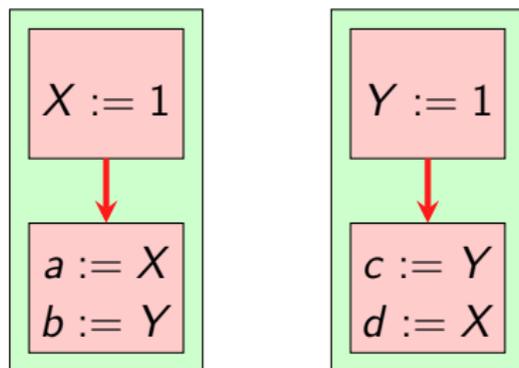
Allowed outcome: $a = b = 0$.

IRIW: Not just store buffering

Initially, $X = Y = 0$.

$$X := 1 \parallel Y := 1 \parallel \begin{array}{l} a := X; \\ b := Y \end{array} \parallel \begin{array}{l} c := Y; \\ d := X \end{array}$$

Allowed outcome: $a = c = 1$ and $b = d = 0$.



Coherence:

“SC for a single variable”

Initially, $X = 0$.

$$X := 1 \parallel X := 2 \parallel \begin{array}{l} a := X; \\ b := X \end{array} \parallel \begin{array}{l} c := X; \\ d := X \end{array}$$

Forbidden outcome: $a = 1, b = 2, c = 2, d = 1$.

The C11 memory model

Two types of locations: ordinary and atomic

- ▶ Races on ordinary accesses \rightsquigarrow error

A spectrum of atomic accesses:

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

Explicit primitives for fences

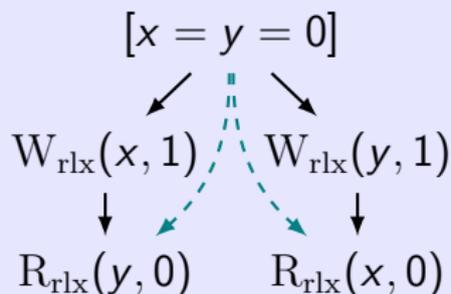
Relaxed behaviour: store buffering

Initially $x = y = 0$.

$$\begin{array}{l} x.\text{store}(1, \text{rlx}); \\ t_1 = y.\text{load}(\text{rlx}); \end{array} \parallel \begin{array}{l} y.\text{store}(1, \text{rlx}); \\ t_2 = x.\text{load}(\text{rlx}); \end{array}$$

This can return $t_1 = t_2 = 0$.

Justification:



Behaviour observed
on x86/Power/ARM

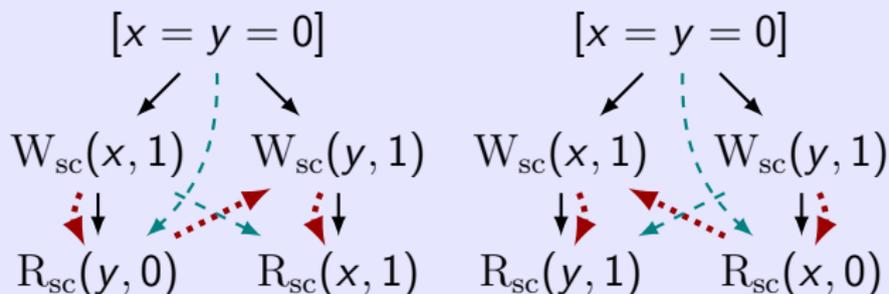
Getting rid of the SB behaviour

Initially $x = y = 0$.

```
x.store(1, sc);    || y.store(1, sc);  
t1 = y.load(sc);  || t2 = x.load(sc);
```

This cannot return $t_1 = t_2 = 0$.

Justification:



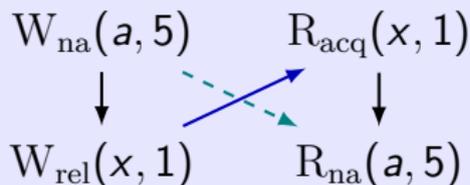
Release-acquire synchronization: message passing

Initially $a = x = 0$.

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

This will always print 5.

Justification:



Release-acquire
synchronization

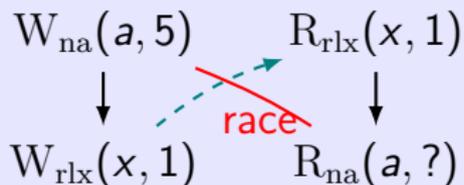
Relaxed accesses don't synchronize

Initially $a = x = 0$.

```
a = 5;           || while (x.load(rlx) == 0);  
x.store(1, rlx); || print(a);
```

The program is racy \rightsquigarrow undefined semantics.

Justification:



Relaxed accesses
don't synchronize

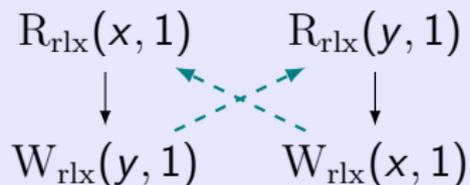
Dependency cycles

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:



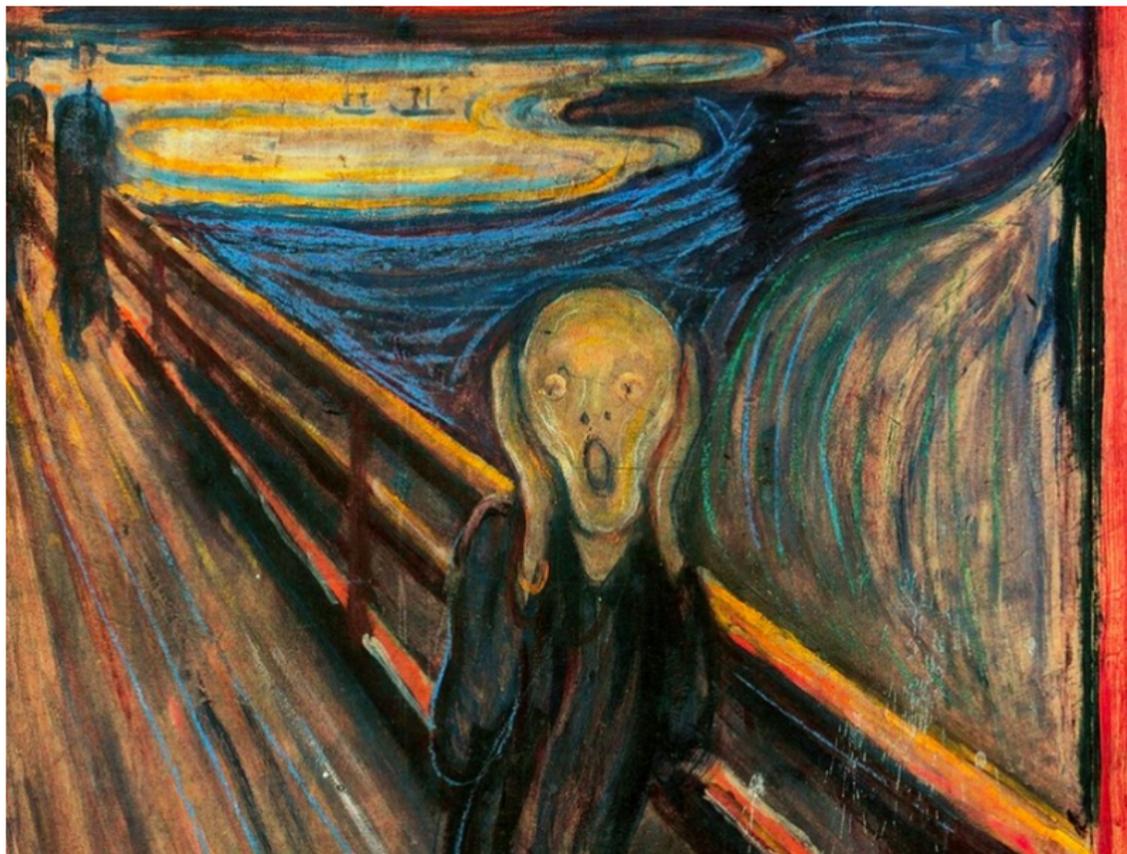
Relaxed accesses
don't synchronize

The C11 weak memory model (simplified)

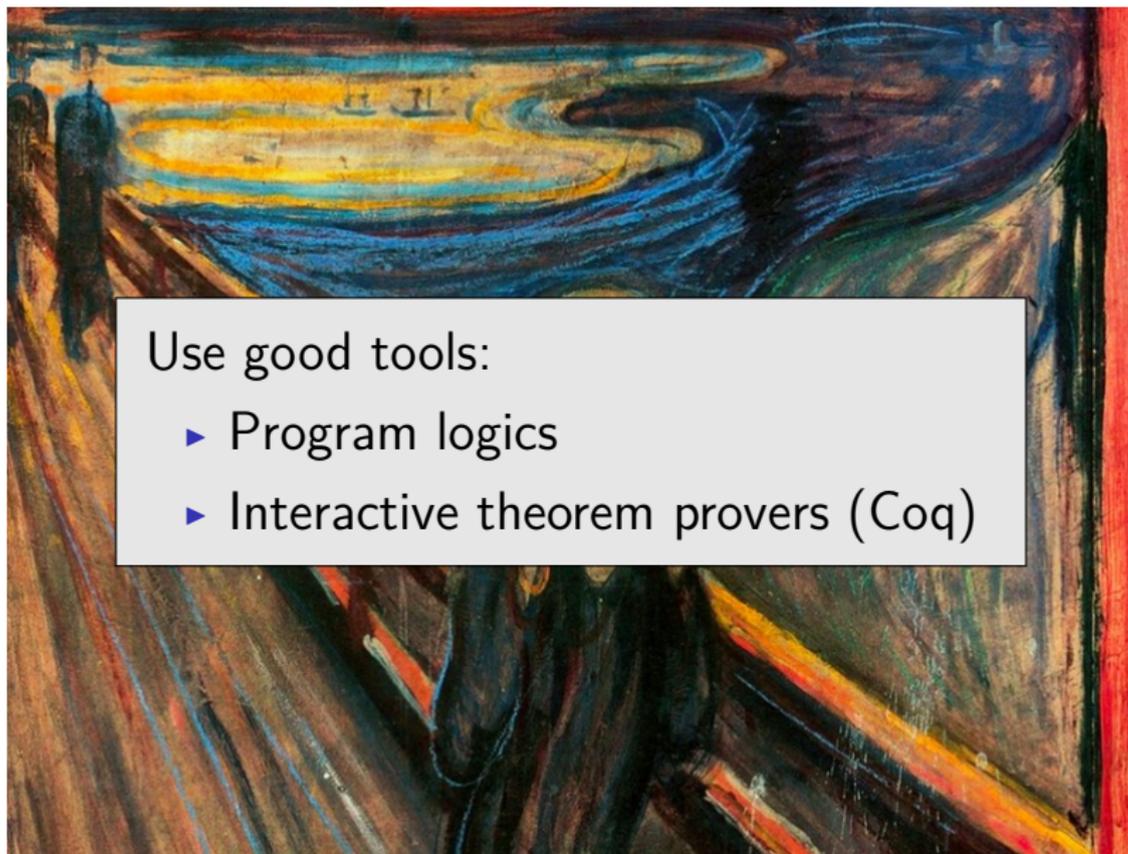
| | |
|--|--|
| $\text{isread}_{\ell,v}(a) \stackrel{\text{def}}{=} \exists X, v'. \text{lab}(a) \in \{\mathbf{R}_X(\ell, v), \mathbf{C}_X(\ell, v, v')\}$ $\text{iswrite}_{\ell,v}(a) \stackrel{\text{def}}{=} \exists X, v'. \text{lab}(a) \in \{\mathbf{W}_X(\ell, v), \mathbf{C}_X(\ell, v', v)\}$ $\text{isfence}(a) \stackrel{\text{def}}{=} \text{lab}(a) \in \{\mathbf{F}_{\text{ACQ}}, \mathbf{F}_{\text{REL}}\}$ $\text{sameThread}(a, b) \stackrel{\text{def}}{=} \text{tid}(a) = \text{tid}(b)$ $\text{rsElem}(a, b) \stackrel{\text{def}}{=} \text{sameThread}(a, b) \vee \text{isrmw}(b)$ $\text{rseq}(a, b) \stackrel{\text{def}}{=} a = b \vee \text{rsElem}(a, b) \wedge \text{mo}(a, b) \wedge (\forall c. \text{mo}(a, c) \wedge \text{mo}(c, b) \Rightarrow \text{rsElem}(a, c))$ $\text{sw}(a, b) \stackrel{\text{def}}{=} \exists c, d. \neg \text{sameThread}(a, b) \wedge \text{isRel}(a) \wedge \text{isAcq}(b) \wedge \text{rseq}(c, \text{rf}(d))$ $\quad \wedge (a = c \vee \text{isfence}(a) \wedge \text{sb}^+(a, c)) \wedge (d = b \vee \text{isfence}(d) \wedge \text{sb}^+(d, b))$ $\text{hb} \stackrel{\text{def}}{=} (\text{sb} \cup \text{sw} \cup \text{asw})^+$ $\text{Racy} \stackrel{\text{def}}{=} \exists a, b. \text{isaccess}(a) \wedge \text{isaccess}(b) \wedge \text{loc}(a) = \text{loc}(b) \wedge a \neq b$ $\quad \wedge (\text{iswrite}(a) \vee \text{iswrite}(b)) \wedge (\text{isNA}(a) \vee \text{isNA}(b)) \wedge \neg(\text{hb}(a, b) \vee \text{hb}(b, a))$ $\text{Observation} \stackrel{\text{def}}{=} \{(a, b) \mid \text{mo}(a, b) \wedge \text{loc}(a) = \text{loc}(b) = \text{world}\}$ | $\text{isread}_{\ell}(a) \stackrel{\text{def}}{=} \exists v. \text{isread}_{\ell,v}(a)$ $\text{iswrite}_{\ell}(a) \stackrel{\text{def}}{=} \exists v. \text{iswrite}_{\ell,v}(a)$ $\text{isaccess}(a) \stackrel{\text{def}}{=} \text{isread}_{\ell}(a) \vee \text{iswrite}_{\ell}(a)$ $\text{isrmw}(a) \stackrel{\text{def}}{=} \text{isread}_{\ell}(a) \wedge \text{iswrite}_{\ell}(a)$ $\text{isAcq}(a) \stackrel{\text{def}}{=} \text{mode}(a) \sqsupseteq \text{ACQ}$ $\text{isread}(a) \stackrel{\text{def}}{=} \exists \ell. \text{isread}_{\ell}(a)$ $\text{iswrite}(a) \stackrel{\text{def}}{=} \exists \ell. \text{iswrite}_{\ell}(a)$ $\text{isNA}(a) \stackrel{\text{def}}{=} \text{mode}(a) = \text{NA}$ $\text{isSC}(a) \stackrel{\text{def}}{=} \text{mode}(a) = \text{SC}$ $\text{isRel}(a) \stackrel{\text{def}}{=} \text{mode}(a) \sqsupseteq \text{REL}$ |
|--|--|

| | |
|--|--|
| $\forall a, b. \text{sb}(a, b) \Longrightarrow \text{tid}(a) = \text{tid}(b) \quad (\text{ConsSB})$ $\text{order}(\text{iswrite}, \text{mo}) \wedge \forall \ell. \text{total}(\text{iswrite}_{\ell}, \text{mo}) \quad (\text{ConsMO})$ $\text{order}(\text{isSC}, \text{sc}) \wedge \text{total}(\text{isSC}, \text{sc})$ $\quad \wedge (\text{hb} \cup \text{mo}) \cap (\text{isSC} \times \text{isSC}) \subseteq \text{sc} \quad (\text{ConsSC})$ $\forall b. (\exists c. \text{rf}(b) = c) \iff$ $\quad \exists \ell, a. \text{iswrite}_{\ell}(a) \wedge \text{isread}_{\ell}(b) \wedge \text{hb}(a, b) \quad (\text{ConsRFdom})$ $\forall a, b. \text{rf}(b) = a \Longrightarrow \exists \ell, v. \text{iswrite}_{\ell,v}(a) \wedge \text{isread}_{\ell,v}(b) \quad (\text{ConsRF})$ $\forall a, b. \text{rf}(b) = a \wedge (\text{isNA}(a) \vee \text{isNA}(b)) \Longrightarrow \text{hb}(a, b) \quad (\text{ConsRFna})$ $\forall a, b. \text{rf}(b) = a \wedge \text{isSC}(b) \Longrightarrow$ $\quad \text{imm}(\text{scr}, a, b) \vee \neg \text{isSC}(a) \wedge \nexists x. \text{hb}(a, x) \wedge \text{imm}(\text{scr}, x, b) \quad (\text{SCRreads})$ $\text{where } \text{order}(P, R) \stackrel{\text{def}}{=} (\nexists a. R(a, a)) \wedge (R^+ \subseteq R) \wedge (R \subseteq P \times P)$ $\text{total}(P, R) \stackrel{\text{def}}{=} (\forall a, b. P(a) \wedge P(b) \Longrightarrow a = b \vee R(a, b) \vee R(b, a))$ | $\nexists a. \text{hb}(a, a) \quad (\text{IrrHB})$ $\nexists a, b. \text{rf}(b) = a \wedge \text{hb}(b, a) \quad (\text{ConsRFhb})$ $\nexists a, b. \text{hb}(a, b) \wedge \text{mo}(b, a) \quad (\text{CohWW})$ $\nexists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), \text{rf}(a)) \quad (\text{CohRR})$ $\nexists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), a) \quad (\text{CohWR})$ $\nexists a, b. \text{hb}(a, b) \wedge \text{mo}(b, \text{rf}(a)) \quad (\text{CohRW})$ $\forall a, b. \text{isrmw}(a) \wedge \text{rf}(a) = b \Longrightarrow$ $\quad \text{imm}(\text{mo}, b, a) \quad (\text{AtRMW})$ $\forall a, b, \ell. \text{lab}(a) = \text{lab}(b) = \mathbf{A}(\ell) \Longrightarrow a = b \quad (\text{ConsAlloc})$ $\text{imm}(R, a, b) \stackrel{\text{def}}{=} R(a, b) \wedge \nexists c. R(a, c) \wedge R(c, b)$ $\text{scr}(a, b) \stackrel{\text{def}}{=} \text{sc}(a, b) \wedge \text{iswrite}_{\text{loc}(b)}(a)$ |
|--|--|

The C11 weak memory model (simplified)



The C11 weak memory model (simplified)



Use good tools:

- ▶ Program logics
- ▶ Interactive theorem provers (Coq)

Two research directions

Verify compilation of C11:

- ▶ Compilation of the atomics to hardware
(Batty et al.'11, Sarkar et al.'12)
- ▶ Source-to-source transformations (see POPL'15)
- ▶ An actual compiler (future work)

Verify concurrent C11 programs:

- ▶ **Using program logics**
- ▶ By reduction to SC (robustness)
- ▶ *Don't verify, just find bugs.*

Understanding C11 using
relaxed program logics

When should we care about relaxed memory?

C11 satisfies the DRF-SC property:

Theorem (DRF-SC)

If $\llbracket \text{Prg} \rrbracket_{\text{SC}}$ contains no data races and no weak atomics, then $\llbracket \text{Prg} \rrbracket_{\text{C11}} = \llbracket \text{Prg} \rrbracket_{\text{SC}}$.

- ▶ Program logics that disallow data races are trivially sound for the NA+SC fragment of C11.

Key concept of *ownership* :

- ▶ Resourceful reading of Hoare triples.

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

- ▶ Disjoint parallelism:

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

Separation logic rules for non-atomic accesses

- ▶ Allocation gives you permission to access x .

$$\{\text{emp}\} x = \text{alloc}(); \{x \mapsto _ \}$$

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

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

Reasoning about SC accesses

- ▶ Model SC accesses as non-atomic accesses inside a CCR.
- ▶ Use concurrent separation logic (CSL)

$$J \vdash \{P\} C \{Q\}$$

- ▶ Rule for SC-atomic reads:

$$\frac{\text{emp} \vdash \{J * P\} \ t = *x; \ \{J * Q\}}{J \vdash \{P\} \ t = x.\text{load}(sc); \ \{Q\}}$$

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.

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

- ▶ 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 $J(v) \stackrel{\text{def}}{=} v = 0 \vee \&a \mapsto 5$.

| | |
|---------------------------------------|---|
| $\{\&a \mapsto 0 * \mathbf{W}_J(x)\}$ | $\{\mathbf{R}_J(x)\}$ |
| $a = 5;$ | 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!

Basically, disallow ownership transfer.

- ▶ Relaxed reads:

$$\{\mathbf{R}_Q(x)\} t := x.\text{load}(r/x) \{\mathbf{R}_Q(x)\}$$

- ▶ Relaxed writes:

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

Unsound because of dependency cycles!

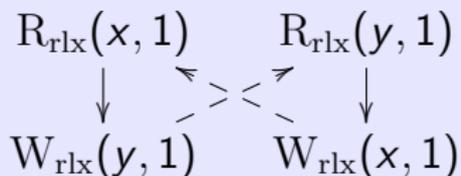
Dependency cycles

Initially $x = y = 0$.

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

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

Justification:



Relaxed accesses
don't synchronize

Dependency cycles

Initially $x = y = 0$.

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

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:

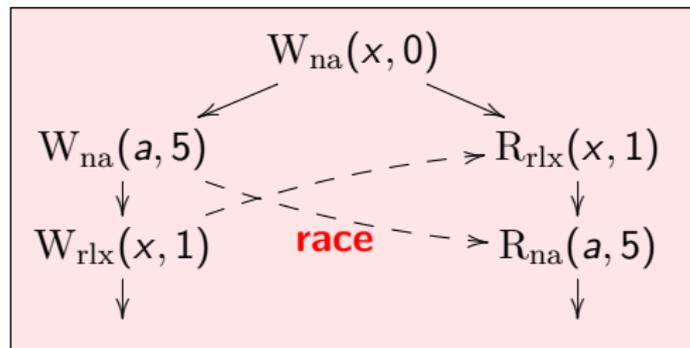
Don't use relaxed writes

∨

Strengthen the model

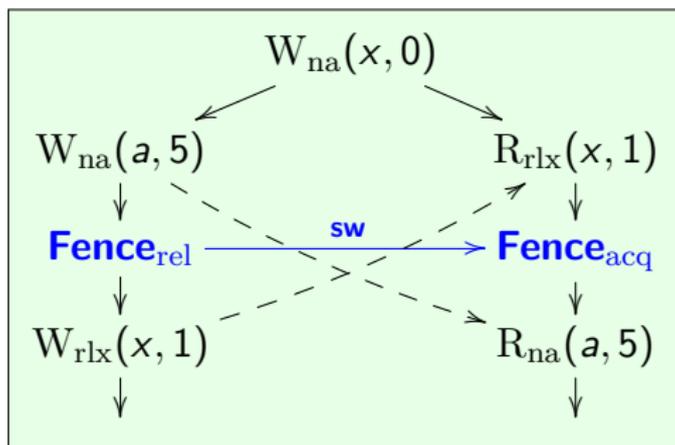
Incorrect message passing

```
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); }
)
```



Reasoning about fences

Joint work with Marko Doko. In progress.

- ▶ Introduce two ‘modalities’ in the logic

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

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

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

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

Reasoning about fences

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

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

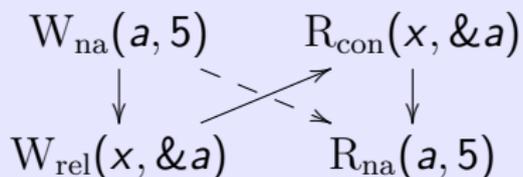
Release-consume synchronization

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

Release-consume synchronization

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

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

This program cannot crash nor print 0.

PL consequences:

Needs funny modality, but otherwise OK.

Proposed rules for consume accesses

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

$$\frac{\begin{array}{c} \{P\} C \{Q\} \\ C \text{ is basic command mentioning } t \end{array}}{\{\nabla_t P\} C \{\nabla_t Q\}}$$

Mutual exclusion locks

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

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

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

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

Summary of program logic features

| Access kind | Program logic features |
|---------------------|---|
| non-atomic | normal SL \mapsto |
| SC-atomic | normal CSL invariants |
| release/ acquire | single-location invariants unidirectional ownership transfer |
| relaxed | send only ΔP ; receive ∇P |
| consume | receive only $\nabla_t P$ |

| Fence kind | Program logic effect |
|------------|--|
| release | introduces ΔP |
| acquire | eliminates ∇P and $\nabla_t P$ |

Relaxed program logics
are good tools for
understanding
weak memory models