

# Explaining relaxed memory models with program transformations

Ori Lahav and Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)

FM, November 2016

# Relaxed memory models

- ▶ **Sequential consistency** (*aka “interleaving semantics”*) is the standard memory model for **reasoning** about concurrent programs.
- ▶ **Modern hardware** employs, e.g., local write buffers, hierarchies of caches, and speculative executions, that significantly improve performance, but **invalidate SC** in the presence of data races.
- ▶ To further improve performance, **compilers** perform **concurrency-oblivious optimizations**.

**Relaxed memory models** provide **formal sound** semantics for **realistic high-performance** concurrency.

# Litmus tests

## Store buffering (SB)

$x := 1;$        $\parallel$        $y := 1;$   
 $a := y; // 0$     $\parallel$     $b := x; // 0$

Allowed by x86-TSO, Power,  
ARM, C11 with non-SC accesses

## Load buffering (LB)

$a := x; // 1$        $\parallel$        $b := y; // 1$   
 $y := 1;$              $\parallel$        $x := 1;$

Allowed by Power, ARM,  
C11 with relaxed accesses

## Litmus tests

### Store buffering (SB)

$x := 1;$        $y := 1;$   
 $a := y; // 0$     $b := x; // 0$

Allowed by x86-TSO, Power,  
ARM, C11 with non-SC accesses

### Load buffering (LB)

$a := x; // 1$        $b := y; // 1$   
 $y := 1;$        $x := 1;$

Allowed by Power, ARM,  
C11 with relaxed accesses

Program transformations provide intuitive explanations.

# Litmus tests

## Store buffering (SB)

$x := 1;$        $y := 1;$   
 $a := y; // 0$     $b := x; // 0$

Allowed by x86-TSO, Power,  
ARM, C11 with non-SC accesses

## Load buffering (LB)

$a := x; // 1$        $b := y; // 1$   
 $y := 1;$        $x := 1;$

Allowed by Power, ARM,  
C11 with relaxed accesses

Program transformations provide intuitive explanations.

## Our goal

Formally reconcile relaxed memory models definitions with the transformations account.

# Litmus tests

## Store buffering (SB)

$x := 1;$       $\parallel$       $y := 1;$   
 $a := y; // 0$     $\parallel$     $b := x; // 0$

Allowed by x86-TSO, Power,  
ARM, C11 with non-SC accesses

## Load buffering (LB)

$a := x; // 1$       $\parallel$       $b := y; // 1$   
 $y := 1;$       $\parallel$       $x := 1;$

Allowed by Power, ARM,  
C11 with relaxed accesses


Program transformations provide intuitive explanations.

## Our goal

Formally reconcile relaxed memory models definitions with the transformations account.

TSO  $\succ$  C11-release/acquire  $\succ$  Power  $\succ$  ARM

# Operational account for store buffering in x86-TSO



## Store buffering

```
x := 1;           || y := 1;  
a := y; // 0     || b := x; // 0
```


## Store buffering + fences

```
x := 1;           || y := 1;  
fence;            || fence;  
a := y; // 0     || b := x; // 0
```

# SB and LB in axiomatic models

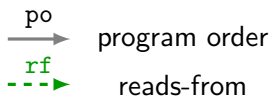

## Store buffering (SB)

```
x := 1;      || y := 1;  
a := y; // 0 || b := x; // 0
```



## Load buffering (LB)

```
a := x; // 1 || b := y; // 1  
y := 1;      || x := 1;
```






# Axiomatic x86-TSO model

## Definition

An execution is *TSO-consistent* if:

- ▶  $po \cup rf$  is acyclic
- ▶ there exists a total ordering  $mo$  of all write events, such that the **none** of the following occurs:




+ conditions on fences and RMWs

# SB and LB in TSO

## Store buffering (SB)


```
x := 1;      || y := 1;  
a := y; // 0 || b := x; // 0
```



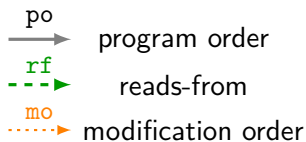
allowed by TSO

## Load buffering (LB)

```
a := x; // 1 || b := y; // 1  
y := 1;      || x := 1;
```



Forbidden by TSO



# Sound optimizations under TSO

## Sound Transformations

$$P_{\text{source}} \rightsquigarrow P_{\text{target}} \implies \llbracket P_{\text{target}} \rrbracket \subseteq \llbracket P_{\text{source}} \rrbracket$$

### Write-read reordering

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

### Read-after-write elimination

$$\begin{array}{l} x := 1; \\ a := x; \end{array} \rightsquigarrow \begin{array}{l} x := 1; \\ a := 1; \end{array}$$

# Sound optimizations under TSO

## Sound Transformations

$$P_{\text{source}} \rightsquigarrow P_{\text{target}} \implies \llbracket P_{\text{target}} \rrbracket \subseteq \llbracket P_{\text{source}} \rrbracket$$

### Write-read reordering

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

### Read-after-write elimination

$$\begin{array}{l} x := 1; \\ a := x; \end{array} \rightsquigarrow \begin{array}{l} x := 1; \\ a := 1; \end{array}$$

## Definition

$G \rightsquigarrow_{\text{TSO}} G'$  if  $G'$  is obtained from  $G$  by one of the following:

$$\begin{array}{ccc} W(x, v_x) & R(y, v_y) & W(x, v) \\ \text{po} \downarrow & \rightsquigarrow \text{po} \downarrow & \text{po} \downarrow \\ R(y, v_y) & W(x, v_x) & W(x, v) \rightsquigarrow R(x, v) \end{array}$$

## Theorem (Soundness of transformations)

If  $G \rightsquigarrow_{\text{TSO}} G'$  and  $G'$  is TSO-consistent, so is  $G$ .

# Alternative TSO characterisation

TSO = SC + WR-reordering + RaW-elimination

## Theorem

*$G$  is TSO-consistent iff  
there exists  $G'$  such that  $G \rightsquigarrow_{\text{TSO}}^* G'$  and  $G'$  is SC-consistent.*

- ( $\Leftarrow$ ) By soundness of transformations.
- ( $\Rightarrow$ ) Assume that  $G$  is TSO-consistent and not SC-consistent. Show that some transformation is applicable and that it preserves TSO-consistency.

Induction metrics:  $|\text{po} \cap (\mathbb{W} \times \mathbb{R})|$

# Application: compilation correctness

## Compilation Correctness

$$\llbracket \text{compile}(P) \rrbracket_{\text{target memory model}} \subseteq \llbracket P \rrbracket_{\text{source memory model}}$$

$$\text{TSO} = \text{SC} + \text{WR-reordering} + \text{RaW-elimination}$$

To prove  $\llbracket \text{compile}(P) \rrbracket_{\text{TSO}} \subseteq \llbracket P \rrbracket_{\text{C11}}$ , it remains to show:

- ▶ Compilation is correct for SC.
- ▶ WR-reorderings and RaW-eliminations correspond to C11-sound transformations in the source program.

# Explaining memory models with program transformations

$\text{TSO} = \text{SC} + \text{WR-reordering} + \text{RaW-elimination}$

$\text{C11 release/acquire} = ?$

$\text{Power} = ?$

$\text{ARM} = ?$

## C11 release/acquire

### Independent reads of independent writes (IRIW)

$$\begin{array}{l} a := x; \text{ // } 1 \\ b := y; \text{ // } 0 \end{array} \parallel x := 1; \parallel y := 1; \parallel \begin{array}{l} c := y; \text{ // } 1 \\ d := x; \text{ // } 0 \end{array}$$

- ▶ This behavior is allowed by C11.
- ▶ No sound thread-local transformation can be applied.
- ▶ Sequentialization,  $C_1 \parallel C_2 \rightsquigarrow C_1; C_2$ , is applicable, and then the outcome is possible.

$$\begin{array}{l} a := x; \text{ // } 1 \\ b := y; \text{ // } 0 \end{array} \parallel x := 1; \rightsquigarrow \begin{array}{l} x := 1; \\ a := x; \\ b := y; \end{array} \rightsquigarrow \begin{array}{l} x := 1; \\ a := 1; \\ b := y; \end{array} \rightsquigarrow \begin{array}{l} b := y; \\ x := 1; \\ a := 1; \end{array}$$



## C11 release/acquire counterexample

RA  $\prec$  SC + WR-reordering + RaW-elimination + SEQ

<code>y := 1;</code>	<code>  </code>	<code>x := 3;</code>	<code>  </code>	<code>z := 1;</code>
<code>x := 1;</code>				<code>x := 2;</code>
<code>a := x; // 3</code>				<code>c := x; // 3</code>
<code>b := z; // 0</code>				<code>d := y; // 0</code>

- ▶ This behavior is allowed by C11.
- ▶ No local transformations are possible.
- ▶ Sequentialisation rules out the outcome.

# POWER multiprocessor

- ▶ weaker than C11 release/acquire

Power  $\prec$  SC + reorderings + eliminations + SEQ

An execution  $G$  is *Power-consistent* if the following hold:

1.  $\text{mo}$  is a disjoint union of relations  $\{\text{mo}_x\}_{x \in \text{Loc}}$ , such that each relation  $\text{mo}_x$  is a strict total order on  $W_x \cup x$ .
2.  $\text{hb}$  is acyclic. *(no-thin-air)*
3.  $\text{po}|_x \cup \text{rf} \cup \text{fr} \cup \text{mo}$  is acyclic for every  $x \in \text{Loc}$ . *(SC-per-loc)*
4.  $\text{fre}; \text{prop}; \text{hb}^*$  is irreflexive. *(observation)*
5.  $\text{mo} \cup \text{prop}$  is acyclic. *(propagation)*
6.  $\text{fr}; \text{mo}$  is irreflexive. *(atomicity)*
7.  $\text{mo}; [\text{U}]; \text{po}; [\text{U}]$  is acyclic.

where:

- $\text{sync} = \text{po}; [\text{F}_{\text{sync}}]; \text{po}$  and  $\text{lwsync} = \text{po}; [\text{F}_{\text{lwsync}}]; \text{po}$
- $\text{fence} = \text{sync} \cup ([\text{RU}]; \text{lwsync}; [\text{RWU}] \cup ([\text{W}]; \text{lwsync}; [\text{WU}]))$  *(fence order)*
- $\text{fr} = (\text{rf}^{-1}; \text{mo}) \setminus [\text{E}]$  *(read before)*
- $\text{rfe} = \text{rf} \setminus \text{po}$  and  $\text{fre} = \text{fr} \setminus \text{po}$  *(external relations)*
- $\text{ppo} = \dots$  *(preserved program order)*
- $\text{hb} = \text{ppo} \cup \text{fence} \cup \text{rfe}$  *(happens-before)*
- $\text{prop}_1 = [\text{WU}]; \text{rfe}^?; \text{fence}; \text{hb}^*; [\text{WU}]$
- $\text{prop}_2 = ((\text{mo} \cup \text{fr}) \setminus \text{po})^?; \text{rfe}^?; (\text{fence}; \text{hb}^*)^?; \text{sync}; \text{hb}^*$
- $\text{prop} = \text{prop}_1 \cup \text{prop}_2$  *(propagation relation)*

An execution  $G$  is *Power-consistent* if the following hold:

1.  $\text{mo}$  is a disjoint union of relations  $\{\text{mo}_x\}_{x \in \text{Loc}}$ , such that each relation  $\text{mo}_x$  is a strict total order on  $W_x \cup x$ .
2.  $\text{hb}$  is acyclic. *(no-thin-air)*
3.  $\text{po}|_x \cup \text{rf} \cup \text{fr} \cup \text{mo}$  is acyclic for every  $x \in \text{Loc}$ . *(SC-per-loc)*
4.  $\text{fre}; \text{prop}; \text{hb}^*$  is irreflexive. *(observation)*
5.  $\text{mo} \cup \text{prop}$  is acyclic. *(propagation)*
6.  $\text{fr}; \text{mo}$  is irreflexive. *(atomicity)*

The model allows cycles in  $\text{po} \cup \text{rf}$ !

- ▶ The definition is not *prefix-closed*.
- ▶ *Speculation* is required to operationally construct executions.

- $\text{ppo} = \dots$  *(preserved program order)*
- $\text{hb} = \text{ppo} \cup \text{fence} \cup \text{rfe}$  *(happens-before)*
- $\text{prop}_1 = [\text{WU}]; \text{rfe}^?; \text{fence}; \text{hb}^*; [\text{WU}]$
- $\text{prop}_2 = ((\text{mo} \cup \text{fr}) \setminus \text{po})^?; \text{rfe}^?; (\text{fence}; \text{hb}^*)^?; \text{sync}; \text{hb}^*$
- $\text{prop} = \text{prop}_1 \cup \text{prop}_2$  *(propagation relation)*


# Load buffering

The model allows cycles in  $po \cup rf$ !

- ▶ The definition is not *prefix-closed*.
- ▶ *Speculation* is required to operationally construct executions.

## Load buffering (LB)

```
a := x; // 1 ||| x := y;  
y := 1;
```



allowed by Power

## Definition (Strong Power model)


An execution is *StrongPower-consistent* if it is Power-consistent and  $po \cup rf$  is acyclic.

# Reduction to StrongPower-consistency

$$\text{Power} = \text{StrongPower} + \text{reorderings}$$

## Definition

$G \rightsquigarrow_{\text{Power}} G'$  if  $G'$  is obtained from  $G$  by reordering two independent adjacent memory accesses to different locations:



## Theorem

An execution  $G$  is Power-consistent iff  $G \rightsquigarrow_{\text{Power}}^* G'$  for some StrongPower-consistent execution  $G'$ .


- ▶ Limitation: Power's `isync` fences are excluded

# A strange ARM behaviour

ARM  $\prec$  StrongARM + reorderings + eliminations

ARM weak

```
a := x; // 1 || y := x; || x := y;  
x := 1;
```



allowed by ARM

- ▶ No local transformation can be applied.

# Summary

## High-level points

- ▶ Some memory models can be defined via transformations.
- ▶ But there is more to weak memory than transformations.

## Technical results

- ▶  $\text{TSO} = \text{SC} + \text{WR-reordering} + \text{RaW-elimination}$
- ▶  $\text{RA} \prec \text{SC} + \text{WR-reordering} + \text{RaW-elimination} + \text{SEQ}$
- ▶  $\text{Power} = \text{StrongPower} + \text{reorderings}$
- ▶  $\text{ARM} \prec \text{StrongARM} + \text{reorderings} + \text{eliminations}$

## Application

- ▶ Simplify compilation correctness proofs

See <http://plv.mpi-sws.org/trns/> for more details and Coq proofs.



# Summary

## High-level points

- ▶ Some memory models can be defined via transformations.
- ▶ But there is more to weak memory than transformations.

## Technical results

- ▶  $\text{TSO} = \text{SC} + \text{WR-reordering} + \text{RaW-elimination}$
- ▶  $\text{RA} \prec \text{SC} + \text{WR-reordering} + \text{RaW-elimination} + \text{SEQ}$
- ▶  $\text{Power} = \text{StrongPower} + \text{reorderings}$
- ▶  $\text{ARM} \prec \text{StrongARM} + \text{reorderings} + \text{eliminations}$

## Application

- ▶ Simplify compilation correctness proofs

See <http://plv.mpi-sws.org/trns/> for more details and Coq proofs.

*Thank you!*