

Fences in Weak Memory Models: Appendix

Jade Alglave¹, Luc Maranget¹, Susmit Sarkar², and Peter Sewell²

¹ INRIA ² University of Cambridge

Abstract. We present an axiomatic framework, implemented in Coq, to define weak memory models *w.r.t.* several parameters: local reorderings of reads and writes, and visibility of inter and intra processor communications through memory, including full store atomicity relaxation. Thereby, we give a formal hierarchy of weak memory models, in which we provide a formal study of what should be the action and placement of fences to restore a given model such as *SC* from a weaker one. Finally, we provide a tool, *diy*, that tests a given machine to determine the architecture it exhibits. We detail the results of our experiments on Power and the model we extract from it. This identified an implementation error in Power 5 memory barriers (for which IBM is providing a software workaround); our results also suggest that Power 6 does not suffer from this problem.

We present here additional data and explanations, as well as the proof sketches of the results that appear in the paper. Moreover, our results are entirely implemented in the Coq proof assistant: the development is available at <http://moscova.inria.fr/~alglave/wmm/>.

1 Tables of notations

Communication relaxations

diy relaxation	Relation	Source	Target	Processor	Location
Rfe	$\xrightarrow{\text{rfe}}$	W	R	Different	Same
Rfi	$\xrightarrow{\text{rfi}}$	W	R	Same	Same
Wse	$\xrightarrow{\text{wse}}$	W	W	Different	Same
Wsi	$\xrightarrow{\text{wsi}}$	W	W	Same	Same
Fre	$\xrightarrow{\text{fre}}$	R	W	Different	Same
Fri	$\xrightarrow{\text{fri}}$	R	W	Same	Same

Program order relaxations

Program order proper

diy relaxation	Relation	Source	Target	Processor	Location
PosRR	$\bigcup_{\ell} \xrightarrow{\text{po}} \cap(\mathbb{R}_{\ell} \times \mathbb{R}_{\ell})$	R	R	Same	Same
PodRR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{po}} \cap(\mathbb{R}_{\ell_1} \times \mathbb{R}_{\ell_2})$	R	R	Same	Diff
PosRW	$\bigcup_{\ell} \xrightarrow{\text{po}} \cap(\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})$	R	W	Same	Same
PodRW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{po}} \cap(\mathbb{R}_{\ell_1} \times \mathbb{W}_{\ell_2})$	R	W	Same	Diff
PosWW	$\bigcup_{\ell} \xrightarrow{\text{po}} \cap(\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})$	W	W	Same	Same
PodWW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{po}} \cap(\mathbb{W}_{\ell_1} \times \mathbb{W}_{\ell_2})$	W	W	Same	Diff
PosWR	$\bigcup_{\ell} \xrightarrow{\text{po}} \cap(\mathbb{W}_{\ell} \times \mathbb{R}_{\ell})$	W	R	Same	Same
PodWR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{po}} \cap(\mathbb{W}_{\ell_1} \times \mathbb{R}_{\ell_2})$	W	R	Same	Diff

Dependencies

diy relaxation	Relation	Source	Target	Processor	Location
DpsR	$\bigcup_{\ell} \xrightarrow{\text{dp}} \cap(\mathbb{R}_{\ell} \times \mathbb{R}_{\ell})$	R	R	Same	Same
DpdR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{dp}} \cap(\mathbb{R}_{\ell_1} \times \mathbb{R}_{\ell_2})$	R	R	Same	Diff
DpsW	$\bigcup_{\ell} \xrightarrow{\text{dp}} \cap(\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})$	R	W	Same	Same
DpdW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{dp}} \cap(\mathbb{R}_{\ell_1} \times \mathbb{W}_{\ell_2})$	R	W	Same	Diff

Barriers relaxations: base cases

Sync

diy relaxation	Relation	Source	Target	Processor	Location
SyncsRR	$\bigcup_{\ell} \xrightarrow{\text{sync}} \cap(\mathbb{R}_{\ell} \times \mathbb{R}_{\ell})$	R	R	Same	Same
SyncdRR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{sync}} \cap(\mathbb{R}_{\ell} \times \mathbb{R}_{\ell})$	R	R	Same	Diff
SyncsRW	$\bigcup_{\ell} \xrightarrow{\text{sync}} \cap(\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})$	R	W	Same	Same
SyncdRW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{sync}} \cap(\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})$	R	W	Same	Diff
SyncsWR	$\bigcup_{\ell} \xrightarrow{\text{sync}} \cap(\mathbb{W}_{\ell} \times \mathbb{R}_{\ell})$	W	R	Same	Same
SyncdWR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{sync}} \cap(\mathbb{W}_{\ell} \times \mathbb{R}_{\ell})$	W	R	Same	Diff
SyncsWW	$\bigcup_{\ell} \xrightarrow{\text{sync}} \cap(\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})$	W	W	Same	Same
SyncdWW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{sync}} \cap(\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})$	W	W	Same	Diff

LwSync

diy relaxation	Relation	Source	Target	Processor	Location
LwSyncsRR	$\bigcup_{\ell} \xrightarrow{\text{lwsync}} \cap(\mathbb{R}_{\ell} \times \mathbb{R}_{\ell})$	R	R	Same	Same
LwSyncdRR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{lwsync}} \cap(\mathbb{R}_{\ell} \times \mathbb{R}_{\ell})$	R	R	Same	Diff
LwSyncsRW	$\bigcup_{\ell} \xrightarrow{\text{lwsync}} \cap(\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})$	R	W	Same	Same
LwSyncdRW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{lwsync}} \cap(\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})$	R	W	Same	Diff
LwSyncsWR	$\bigcup_{\ell} \xrightarrow{\text{lwsync}} \cap(\mathbb{W}_{\ell} \times \mathbb{R}_{\ell})$	W	R	Same	Same
LwSyncdWR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{lwsync}} \cap(\mathbb{W}_{\ell} \times \mathbb{R}_{\ell})$	W	R	Same	Diff
LwSyncsWW	$\bigcup_{\ell} \xrightarrow{\text{lwsync}} \cap(\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})$	W	W	Same	Same
LwSyncdWW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{lwsync}} \cap(\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})$	W	W	Same	Diff

Barriers relaxations: A-cumulativity

Sync

diy relaxation	Relation	Source	Target	Processor	Location
ACSyncsRR	$\bigcup_{\ell} \xrightarrow{\text{rfe}}; (\xrightarrow{\text{sync}} \cap (\mathbb{R}_{\ell} \times \mathbb{R}_{\ell}))$	R	R	Same	Same
ACSyncdRR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{rfe}}; (\xrightarrow{\text{sync}} \cap (\mathbb{R}_{\ell} \times \mathbb{R}_{\ell}))$	R	R	Same	Diff
ACSyncsRW	$\bigcup_{\ell} \xrightarrow{\text{rfe}}; (\xrightarrow{\text{sync}} \cap (\mathbb{R}_{\ell} \times \mathbb{W}_{\ell}))$	R	W	Same	Same
ACSyncdRW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{rfe}}; (\xrightarrow{\text{sync}} \cap (\mathbb{R}_{\ell} \times \mathbb{W}_{\ell}))$	R	W	Same	Diff

LwSync

diy relaxation	Relation	Source	Target	Processor	Location
ACLwSyncsRR	$\bigcup_{\ell} \xrightarrow{\text{rfe}}; (\xrightarrow{\text{lwsync}} \cap (\mathbb{R}_{\ell} \times \mathbb{R}_{\ell}))$	R	R	Same	Same
ACLwSyncdRR	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{rfe}}; (\xrightarrow{\text{lwsync}} \cap (\mathbb{R}_{\ell} \times \mathbb{R}_{\ell}))$	R	R	Same	Diff
ACLwSyncsRW	$\bigcup_{\ell} \xrightarrow{\text{rfe}}; (\xrightarrow{\text{lwsync}} \cap (\mathbb{R}_{\ell} \times \mathbb{W}_{\ell}))$	R	W	Same	Same
ACLwSyncdRW	$\bigcup_{\ell_1 \neq \ell_2} \xrightarrow{\text{rfe}}; (\xrightarrow{\text{lwsync}} \cap (\mathbb{R}_{\ell} \times \mathbb{W}_{\ell}))$	R	W	Same	Diff

Barriers relaxations: B-cumulativity

Sync

diy relaxation	Relation	Source	Target	Processor	Location
BCSyncsRW	$\bigcup_{\ell} (\xrightarrow{\text{sync}} \cap (\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})); \xrightarrow{\text{rfe}}$	R	W	Same	Same
BCSyncdRW	$\bigcup_{\ell_1 \neq \ell_2} (\xrightarrow{\text{sync}} \cap (\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})); \xrightarrow{\text{rfe}}$	R	W	Same	Diff
BCSyncsWW	$\bigcup_{\ell} (\xrightarrow{\text{sync}} \cap (\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})); \text{rfe}$	W	W	Same	Same
BCSyncdWW	$\bigcup_{\ell_1 \neq \ell_2} (\xrightarrow{\text{sync}} \cap (\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})); \xrightarrow{\text{rfe}}$	W	W	Same	Diff

LwSync

diy relaxation	Relation	Source	Target	Processor	Location
BCLwSyncsRW	$\bigcup_{\ell} (\xrightarrow{\text{lwsync}} \cap (\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})); \xrightarrow{\text{rfe}}$	R	W	Same	Same
BCLwSyncdRW	$\bigcup_{\ell_1 \neq \ell_2} (\xrightarrow{\text{lwsync}} \cap (\mathbb{R}_{\ell} \times \mathbb{W}_{\ell})); \xrightarrow{\text{rfe}}$	R	W	Same	Diff
BCLwSyncsWW	$\bigcup_{\ell} (\xrightarrow{\text{lwsync}} \cap (\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})); \xrightarrow{\text{rfe}}$	W	W	Same	Same
BCLwSyncdWW	$\bigcup_{\ell_1 \neq \ell_2} (\xrightarrow{\text{lwsync}} \cap (\mathbb{W}_{\ell} \times \mathbb{W}_{\ell})); \xrightarrow{\text{rfe}}$	W	W	Same	Diff

2 Properties of validity

From our definition of validity arises a simple notion of comparison among architectures. $A_1 \leq A_2$ means that A_1 is *weaker* than A_2 :

Definition 1 (Weaker).

$$A_1 \leq A_2 \triangleq (\overset{pp}{\rightarrow} \subseteq \overset{pp}{\rightarrow}) \wedge (\overset{grf}{\rightarrow} \subseteq \overset{grf}{\rightarrow})$$

2.1 Validity is decreasing

We prove validity of an execution to be decreasing w.r.t the weaker predicate; thus, a weak architecture exhibits at least all the behaviours authorised by a stronger one:

Theorem 1 (Validity is decreasing).

$$\forall A_1 A_2, (A_1 \leq A_2) \Rightarrow (\forall X, A_2^\epsilon . \text{valid}(X) \Rightarrow A_1^\epsilon . \text{valid}(X))$$

Proof (in Coq). From $A_1 \leq A_2$, we immediately have $A_1.ghb \subseteq A_2.ghb$, thus if $A_2.ghb$ is acyclic, so is $A_1.ghb$. \square

2.2 Monotonicity of validity

Programs running on an architecture A_1 exhibit executions that would be valid on a stronger architecture A_2 ; we characterise all such executions as follows:

$$A_1.\text{check}_{A_2}(X) \triangleq \text{acyclic}(\overset{\text{grf}_2}{\rightarrow} \cup \overset{\text{ws}}{\rightarrow} \cup \overset{\text{fr}}{\rightarrow} \cup \overset{\text{ppo}_2}{\rightarrow})$$

Theorem 2 (Characterisation).

$$\forall A_1 A_2, (A_1 \leq A_2) \Rightarrow (\forall X, (A_1^\epsilon . \text{valid}(X) \wedge A_1 . \text{check}_{A_2}(X)) \Leftrightarrow A_2^\epsilon . \text{valid}(X))$$

Proof (in Coq).

- \Rightarrow X being valid on A_1 , we have all requirements – well formedness and uniproc – to guarantee it is valid on A_2 , except the last predicate, which holds by the hypothesis check_{A_2} .
- \Leftarrow X being valid on A_2 gives us all requirements – well formedness and uniproc – to guarantee its validity on A_1 except the last one. As $A_1 \leq A_2$, we know that $A_1.ghb \subseteq A_2.ghb$ (lemma ghb.incl), thus the acyclicity requirement for $A_1.ghb$ holds if $A_2.ghb$ is acyclic. \square

3 Examples of models in our framework

3.1 SC

Definitions

$$\begin{aligned} Sc &\triangleq (MM, \overset{\text{rf}}{\rightarrow}, \emptyset) \\ \overset{\text{hb-seq}}{\rightarrow} &\triangleq \overset{\text{ws}}{\rightarrow} \cup \overset{\text{fr}}{\rightarrow} \cup \overset{\text{rf}}{\rightarrow} \end{aligned}$$

Sc characterisation

Corollary 1 (Sc characterisation).

$$\forall A, (A \leq Sc) \Rightarrow (\forall X, \text{valid } AX \wedge A . \text{check}_{Sc} X \Leftrightarrow Sc . \text{valid } X)$$

Proof (in Coq).

- \Rightarrow As $\overset{\text{po}}{\rightarrow} \cup \overset{\text{hb-seq}}{\rightarrow} = Sc.ghb$, this is a direct consequence of thm. 2.
- \Leftarrow as $A \leq Sc$, this is a direct consequence of thm. 1. \square

Sc is SC SC has been defined in [1] as follows:

[...] the result of any execution is the same as if the operations of all the processors were executed in some sequential order, and the operations of each individual processor appear in this sequence in the order specified by its program.

Thus an SC execution is a total order $\xrightarrow{\text{ex}}$ consistent with the program order:

$$\text{seq}(\xrightarrow{\text{ex}}) \triangleq \text{total-order}(\xrightarrow{\text{ex}}, \mathbb{E}) \wedge \text{po} \subseteq \xrightarrow{\text{ex}}$$

The implicit execution model of [1] states that a read r takes its value from the most recent write that precedes it in $\xrightarrow{\text{ex}}$. Writing $\text{pw}(o, r)$ for the set of writes to the same location that precede r according to partial order o , we extract our $\xrightarrow{\text{rf}}$ relation from $\xrightarrow{\text{ex}}$ as follows:

$$\begin{aligned} \text{pw}(\xrightarrow{\text{ex}}, r) &\triangleq \{w \mid \text{loc}(w) = \text{loc}(r) \wedge w \xrightarrow{\text{ex}} r\} \\ SC.\text{rf}(\xrightarrow{\text{ex}}) &\triangleq \{(w, r) \mid w = \max_{\xrightarrow{\text{ex}}}(\text{pw}(\xrightarrow{\text{ex}}, r))\} \end{aligned}$$

We extract write serialisation as well, and one of our execution witnesses from $\xrightarrow{\text{ex}}$, which we use to show that the SC definition from [1] is equivalent to ours:

$$\begin{aligned} SC.\text{ws}(\xrightarrow{\text{ex}}) &\triangleq \bigcup_{\ell} (\mathbb{W}_{\ell} \times \mathbb{W}_{\ell}) \cap \xrightarrow{\text{ex}} \\ SC.\text{wit}(\xrightarrow{\text{ex}}) &\triangleq (\mathbb{E}, \text{po}, SC.\text{rf}(\xrightarrow{\text{ex}}), SC.\text{ws}(\xrightarrow{\text{ex}})) \end{aligned}$$

Theorem 3 (Sc is SC).

$$\forall X, Sc.\text{valid}(X) \Leftrightarrow \exists \xrightarrow{\text{ex}}, \text{seq}(\xrightarrow{\text{ex}}) \wedge SC.\text{wit}(\xrightarrow{\text{ex}}) = X$$

Proof (in Coq).

\Rightarrow from X being valid on Sc , we have $\text{acyclic}(\xrightarrow{\text{ghb}})$, that is $\text{acyclic}(\xrightarrow{\text{hb-seq}} \cup \text{po})$ on Sc , which gives us an equivalent SC execution by Cor. 1.
 \Leftarrow from $\xrightarrow{\text{ex}}$, we produce a $SC.\text{wit}$ which is valid by Thm. 1.

□

3.2 TSO

Definitions

$$\text{ppo-tso} \triangleq RM \cup WW$$

$$Tso^{\epsilon} \triangleq (\text{ppo-tso}, \xrightarrow{\text{rfe}}, \emptyset)$$

$$\text{hb-tso} \triangleq \text{ws} \cup \text{fr} \cup \xrightarrow{\text{rfe}}$$

Tso characterisation By Thm. 2 we show the following criterion – where $\xrightarrow{\text{hb-tso}}$ is $\xrightarrow{\text{ws}} \cup \xrightarrow{\text{fr}} \cup \xrightarrow{\text{rfe}}$ – characterises valid executions that are Tso^* on any $A \leq Tso$:

$$A. \text{check}_{Tso}(X) \triangleq \text{acyclic}(\xrightarrow{\text{ppo-tso}} \cup \xrightarrow{\text{hb-tso}})$$

Corollary 2 (*Tso characterisation*).

$$\forall A, (A \leq Tso) \Rightarrow (\forall X, \text{valid } AX \wedge A. \text{check}_{Tso} X \Leftrightarrow \text{valid } Tso(X))$$

Proof (in Coq).

\Rightarrow As $\xrightarrow{\text{ppo-tso}} \cup \xrightarrow{\text{hb-tso}} = Tso. \xrightarrow{\text{ghb}}$, this is a direct consequence of thm. 2.
 \Leftarrow as $A \leq Tso$, this is a direct consequence of thm. 1. \square

Tso is TSO

Equivalence with the native model Sparc [3, V. 8, Appendix K] formally defines a *TSO* execution as a partial order $\xrightarrow{\text{ex}}$ on memory events constrained by some axioms that constrain $\xrightarrow{\text{ex}}$. We formulate¹ those as follows:

$$\begin{aligned} \text{ptso}(\xrightarrow{\text{ex}}) &\triangleq \text{partial-order}(\xrightarrow{\text{ex}}, \mathbb{E}) \wedge RM \subseteq \xrightarrow{\text{ex}} \wedge WW \subseteq \xrightarrow{\text{ex}} \wedge \\ &\exists \xrightarrow{\text{tso}}, \xrightarrow{\text{tso}} \subseteq \xrightarrow{\text{ex}} \wedge \text{total-order}(\xrightarrow{\text{tso}}, \mathbb{W}) \end{aligned}$$

Note that $\xrightarrow{\text{tso}}$ above is the postulated total order on stores that justify the name *TSO*. An additional axiom, *Value*, where L_a (resp. S_a) is the notation of [3] for a load (resp. a store) with location a :

$$Val(L_a) = Val(\max_{\xrightarrow{\text{ex}}} \{S_a \xrightarrow{\text{ex}} L_a \vee S_a \xrightarrow{\text{po}} L_a\})$$

[...] states that the value of a data load is the value written by the most recent store to that location. Two terms combine to define the most recent store. The first corresponds to stores by other processors, while the second corresponds to stores by the processor that issued the load.

We relate Sparc *TSO* and our *Tso* by interpreting the *Value* axiom as specifying $\xrightarrow{\text{rf}}$, and extract $\xrightarrow{\text{ws}}$ from $\xrightarrow{\text{ex}}$, as we did in the *Sc* case:

$$\begin{aligned} TSO. \text{rf}(\xrightarrow{\text{ex}}) &\triangleq \{(w, r) \mid w = \max_{\xrightarrow{\text{ex}}} (\text{pw}(\xrightarrow{\text{ex}} \cup \xrightarrow{\text{po}}, r))\} \\ TSO. \text{ws}(\xrightarrow{\text{ex}}) &\triangleq \{(w_1, w_2) \mid \exists l, (w_1, w_2) \in (\mathbb{W}_l \times \mathbb{W}_l) \wedge w_1 \xrightarrow{\text{ex}} w_2\} \\ TSO. \text{wit}(\xrightarrow{\text{ex}}) &\triangleq (\mathbb{E}, \xrightarrow{\text{po}}, TSO. \text{rf}(\xrightarrow{\text{ex}}), SC. \text{ws}(\xrightarrow{\text{ex}})) \end{aligned}$$

Theorem 4 (*Tso is TSO*).

$$\forall X, Tso^*. \text{valid}(X) \Leftrightarrow \exists \xrightarrow{\text{ex}}, \text{ptso}(\xrightarrow{\text{ex}}) \wedge TSO. \text{wit}(\xrightarrow{\text{ex}}) = X$$

¹ We omit the axioms *Atomicity* and *Termination*.

Proof (in Coq).

\Rightarrow from X being valid on Tso , we have $\text{acyclic}^{\text{ghb}}(\xrightarrow{\text{}})$, that is $\text{acyclic}^{\text{ghb}}(\xrightarrow{\text{hb-tso}} \cup \xrightarrow{\text{po-tso}})$ on Tso , which gives us an equivalent TSO execution by Cor. 2.
 \Leftarrow from $\xrightarrow{\text{ex}}$, we produce a TSO . wit which is valid by Thm. 1. □

4 Barriers

From now on, we note $A. \xrightarrow{\text{hb-seq}}$ for $\xrightarrow{\text{rfe}} \cup (\xrightarrow{\text{ws}}; \xrightarrow{\text{?rfe}}) \cup (\xrightarrow{\text{fr}}; \xrightarrow{\text{?rfe}} \cup \xrightarrow{\text{ws}} \cup \xrightarrow{\text{fr}})$. We trivially have $A. \xrightarrow{\text{hb-seq}} \subseteq A. \xrightarrow{\text{hb-seq}}$, and that $\xrightarrow{\text{hb-seq}}$ is transitive.

A key lemma is that forall relation $\xrightarrow{\text{r}}$ over events, such that there is a cycle in $\xrightarrow{\text{hb-seq}} \cup \xrightarrow{\text{r}}$, then there is a cycle in $\xrightarrow{\text{hb-seq}}; (\xrightarrow{\text{r}})$.

4.1 Barriers guarantee

Consider two architectures $A_1 \leq A_2$. We define the predicate fb (*fully barriered*) on $A_1 \leq A_2$, where $\xrightarrow{\text{r}_2 \setminus \text{r}_1} \triangleq \xrightarrow{\text{r}_2} \setminus \xrightarrow{\text{r}_1}$ is the set difference, and $x \xrightarrow{\text{r}_1}; \xrightarrow{\text{r}_2} y \triangleq \exists z, x \xrightarrow{\text{r}_1} z \wedge z \xrightarrow{\text{r}_2} y$ stands for the sequence:

$$A_1. \text{fb}_{A_2}(X) \triangleq ((\text{ppo}_2 \setminus \text{r}_1) \cup (\text{grf}_2 \setminus \text{r}_1; \text{ppo}_2)) \subseteq \xrightarrow{\text{ab}_1}$$

We prove that the above condition on $\xrightarrow{\text{ab}_1}$ suffices to restore A_2 from A_1 :

Theorem 5 (Barriers guarantee).

$$\forall A_1 A_2, (A_1 \leq A_2) \Rightarrow (\forall X, A_1. \text{valid}(X) \wedge A_1. \text{fb}_{A_2}(X) \Rightarrow A_2. \text{valid}(X))$$

Proof (in Coq).

Let relation $\xrightarrow{\text{hb}'_2}$ be $\xrightarrow{\text{?rfe}_2} \cup (\xrightarrow{\text{ws}}; \xrightarrow{\text{?rfe}_2}) \cup (\xrightarrow{\text{fr}}; \xrightarrow{\text{?rfe}_2}) \cup \xrightarrow{\text{ws}} \cup \xrightarrow{\text{fr}}$, and relation $\xrightarrow{\text{hb}''_2}$ be $(\xrightarrow{\text{?rfi}_2} \cup \text{ppo}_2)^+$. We prove that $\xrightarrow{\text{?rf}_2} \cup \xrightarrow{\text{ws}} \cup \xrightarrow{\text{fr}} \cup \text{ppo}_2$ is acyclic if and only if $R_2 = \xrightarrow{\text{hb}'_2}; \xrightarrow{\text{hb}''_2}$ is. Then, to prove the theorem it suffices to prove that $\xrightarrow{\text{ghb}_1}$ is cyclic whenever R_2 is.

Let us consider a cycle in R_2 . Because $m_1 \xrightarrow{\text{rf}} m_2 \xrightarrow{\text{hb}'_2} m_3$ implies $m_1 \xrightarrow{\text{hb}'_2} m_3$, we can safely assume that no $\xrightarrow{\text{hb}''_2}$ step ends by a $\xrightarrow{\text{?rfi}_1}$ step. Then, as no two $\xrightarrow{\text{rfi}}$ steps can follow one another, condition $A_1. \text{fb}_{A_2}(X)$ suffices to guarantee the inclusion $R_2 \subseteq (\xrightarrow{\text{ghb}_1})^+$ and conclude. □

4.2 Considering a weaker barriers guarantee

When two architectures A_2 and A_1 have the same policy *w.r.t.* the store atomicity and store buffer relaxations, which we model by $\xrightarrow{\text{grf}_1} = \xrightarrow{\text{grf}_2}$, there is no need for a barrier as powerful as above to restore A_2 from A_1 : a barrier that only orders the events that surround it statically—that is, a non cumulative barrier, which action we model by $\text{non-cumul}(X, \xrightarrow{\text{fenced}}) \triangleq \xrightarrow{\text{fenced}}$ —is enough. Consider

the wfb predicate, which states that the barriers provided by A_1 maintain the pairs that are preserved in the program order on A_2 but not on A_1 :

$$A_1 . \text{wfb}_{A_2}(X) \triangleq \text{ppo}_2 \setminus \text{ab}_1 \subseteq \text{ab}_1$$

The same guarantee applies if A_2 hinders the store buffer relaxation by its preserved program order, *i.e.* when $\xrightarrow{\text{rfi}} \subseteq \text{ppo}_2$ —which is particular to Sc :

Theorem 6 (Non cumulative barriers guarantee).

$$\begin{aligned} & \forall A_1 A_2, ((A_1 \leq A_2) \wedge (\overset{\text{grf}_1}{\xrightarrow{\text{rfi}}} = \overset{\text{grf}_2}{\xrightarrow{\text{rfi}}} \vee (\overset{\text{rfi}}{\xrightarrow{\text{rfi}}} \subseteq \text{ppo}_2))) \Rightarrow \\ & (\forall X, A_1 . \text{valid}(X) \wedge A_1 . \text{wfb}_{A_2}(X) \Rightarrow A_2^\epsilon . \text{valid}(X)) \end{aligned}$$

Proof (in Coq). As for Thm. 5 we need prove $R_2 = (\overset{\text{hb}'_2}{\xrightarrow{\text{rfi}}}; \overset{\text{hb}''_2}{\xrightarrow{\text{rfi}}}) \subseteq (\overset{\text{ghb}_1}{\xrightarrow{\text{rfi}}})^+$. First, by hypothesis $\text{ext}_1 = \text{ext}_2$, we get $\overset{\text{hb}'_2}{\xrightarrow{\text{rfi}}} = \overset{\text{hb}'_1}{\xrightarrow{\text{rfi}}} \subseteq \overset{\text{ghb}_1}{\xrightarrow{\text{rfi}}}$. Second, we prove $\overset{\text{hb}''_2}{\xrightarrow{\text{rfi}}} \subseteq (\overset{\text{ghb}_1}{\xrightarrow{\text{rfi}}})^+$, by considering each step in $\overset{\text{hb}''_2}{\xrightarrow{\text{rfi}}}$. For a particular step $m_1 \xrightarrow{\text{ppo}_2} m_2$, we have either $m_1 \xrightarrow{\text{ppo}_1} m_2 \Rightarrow m_1 \xrightarrow{\text{ghb}_1} m_2$, or $m_1 \xrightarrow{\text{ppo}_2 \setminus \text{ab}_1} m_2 \Rightarrow m_1 \xrightarrow{\text{ab}_1} m_2 \Rightarrow m_1 \xrightarrow{\text{ghb}_1} m_2$. While for all $\xrightarrow{\text{rfi}}$ steps, we conclude either by $\xrightarrow{\text{rfi}} \subseteq \overset{\text{ghb}_1}{\xrightarrow{\text{rfi}}}$ — when $\text{int}_2 = \text{int}_1 = \text{true}$, or by $\xrightarrow{\text{rfi}} \subseteq \text{ppo}_2$. \square

From Tso to Sc As $\xrightarrow{\text{rfe}}$ are global in both Tso and Sc , and Sc hinders the store buffering relaxation by its ppo definition, it suffices by Thm. 6 to fence all pairs in $\text{ppo-sc} \setminus \text{ppo-tso} = WR$ (where $WR \triangleq (\mathbb{W} \times \mathbb{R}) \cap \overset{\text{po}}{\xrightarrow{\text{po}}}$) to restore Sc from Tso :

Corollary 3 (Barriers restoring Sc from Tso).

$$\forall X, (Tso . \text{valid}(X) \wedge \text{non-cumul}(X, WR) \subseteq \overset{\text{ab}_{Tso}}{\xrightarrow{\text{rfe}}}) \Rightarrow Sc . \text{valid}(X)$$

From Pso to Tso $^\epsilon$ We comment here on the two definitions of PSO given in Sparc documentations [3]. We adapt the definition of [3, V8] to our framework:

$$Pso^\epsilon . \text{Arch} \triangleq (\lambda X . RM, \overset{\text{rfe}}{\xrightarrow{\text{rfi}}}, \lambda X . \emptyset)$$

Tso and Pso agree on both the store atomicity and the store buffering relaxations, which allows us to apply Thm. 6: Tso^ϵ is restored from Pso by inserting non cumulative barriers between all $\text{ppo-tso} \setminus \text{ppo-psso} = WW$ pairs. Indeed, Tso is obtained from PSO by adding *StoreStore* barriers after each write [3, V9].

5 diy

5.1 Cycles generation

A type *edge* represents the arrows of an execution witness. It can be in $\overset{\text{rf}}{\xrightarrow{\text{rf}}}$, $\overset{\text{ws}}{\xrightarrow{\text{ws}}}$, $\overset{\text{fr}}{\xrightarrow{\text{fr}}}$: their source and target share the same location and have the appropriate

directions (*e.g.* \xrightarrow{rf} has a write as source and a read as target). Similarly to \xrightarrow{rfi} and \xrightarrow{rfe} , we define \xrightarrow{wsi} and $\xrightarrow{wse, fri}$ and \xrightarrow{fre} . An edge can also be in \xrightarrow{po} , \xrightarrow{dp} or \xrightarrow{fenced} , where we need to specify the source and target directions (except for \xrightarrow{dp} whose source is a read by definition), and if they share the same location.

We specify a concrete syntax for the edges: thus Rfe represents a \xrightarrow{rfe} arrow, Fre a \xrightarrow{fre} arrow, PosRW specifies a \xrightarrow{po} arrow with a read (R) as source, a write (W) as target, knowing these events share the same location (s), and DpdR a \xrightarrow{dp} arrow with a read (R) as target, with different (d) source and target locations.

We define two relations \xrightarrow{relax} and \xrightarrow{safe} , whose default values are respectively $\xrightarrow{po} \cup \xrightarrow{rf}$ and $\xrightarrow{ws} \cup \xrightarrow{fr} \cup \xrightarrow{dp} \cup \xrightarrow{fenced}$. They can be modified by an input of the user: *e.g.* on a x86 machine, which has a *Tso* model [2], \xrightarrow{rfe} can be specified as \xrightarrow{safe} .

We show a cycle in $\xrightarrow{hb-seq} \cup \xrightarrow{po}$ is a cycle in $\xrightarrow{(relax)^+} \cup \xrightarrow{(safe)^+} \cup \xrightarrow{(relax)^+}; \xrightarrow{(safe)^+}$. Thus, a test violating *Sc* corresponds to a cycle in this second relation. But we do not generate cycles with more than one relaxation: if a test exhibits a relaxation, we cannot decide which one is responsible. A cycle Rfe; DpsW; Fre determines whether \xrightarrow{rfe} is global or not since \xrightarrow{dp} and \xrightarrow{fr} are safe, whereas Rfe; DpsW; Rfi; Fre is inconclusive as it exercises both \xrightarrow{rfe} and \xrightarrow{rfi} . We also generate the cycles exhibiting no relaxation (in $\xrightarrow{(safe)^+}$), to test our assumption \xrightarrow{safe} is global.

5.2 Code generation

diy interprets a sequence of edges as a cycle from which it computes a litmus test or fails. One of its execution witnesses includes a cycle compliant with the input edges sequence. The final condition is a conjunction of equalities over the final state of registers and memory locations characterising this cycle.

Test generation performs the following successive steps.

1. We map the edges sequence to a circular double-linked list, whose cells represent memory events, with direction, location, and value fields. An additional field records the edge starting from the event. This list represents the *input cycle* and appear in at least one of the execution witnesses of the produced test.
2. A linear scan sets the events directions, by comparing each target direction with the following source direction. When equal, the in-between cell direction is set to the common value; otherwise (*e.g.* Rfe; Rfe), the generation fails.
3. If no event *e* has an incoming edge specifying a location change, generation fails. Otherwise, a linear scan starting from *e* sets the locations. We reject the cycles where *e* and its predecessor have different locations (*e.g.* Rfe; PodRW).
4. We cut the input cycle into maximal sequences of events with the same location, each being scanned *w.r.t.* the cycle order: we give the value 1 to the first write in this sequence, 2 to the second one, *etc.* Thus the values also reflect the write serialisation order for the specified location.

5. *Significant reads* are the sources of $\xrightarrow{\text{fr}}$ edges and the targets of $\xrightarrow{\text{rf}}$ edges. They are associated with the write on the other side of the edge. In the $\xrightarrow{\text{rf}}$ case, the value of the read is the one of its associated write. In the $\xrightarrow{\text{fr}}$ case, the value of the read is the value of the predecessor of its associated write in $\xrightarrow{\text{ws}}$, *i.e.* by construction the value of its associated write minus 1. Non significant reads do not appear in the test condition.
6. We cut the cycle into maximal sequences of events from the same processor, each being scanned, generating load instructions to (resp. stores from) fresh registers for reads (resp. writes). We insert code implementing a dependency in front of events targeting $\xrightarrow{\text{dr}}$ and the appropriate barrier instruction for events targeting $\xrightarrow{\text{fenced}}$ edges. We built initial state at this step: stores and loads take their addresses from fresh registers whose content (a memory location) is defined in the initial state. Part of the final condition is also built: for any significant read with value v resulting in a load instruction to register r , we add the equality $r = v$.
7. We complete the final condition to characterise write serialisations. The write serialisation for a given location x is defined by the sequence of values 0 (initial value of x), \dots , n , where n is the last value allocated for location x at step 4. If n is 0 or 1 then no addition to the final condition needs to be performed, because the write serialisation is either a singleton or a pair. If n is 2, we add the equality $x = 2$. Otherwise ($n > 2$), we add an *observer* to the program, *i.e.* we add a thread performing n loads from x to r_1, \dots, r_n and add the equalities $r_1 = 1 \wedge \dots \wedge r_n = n$ to the final condition.

References

1. L. Lamport. How to Make a Correct Multiprocess Program Execute Correctly on a Multiprocessor. *IEEE Trans. Comput.*, 46(7):779–782, 1979.
2. S. Owens, S. Sarkar, and P. Sewell. A Better x86 Memory Model: x86-TSO. In *TPHOL 2009*.
3. Sparc Architecture Manual Versions 8 and 9, 1992 and 1994.