

# Stability in Weak Memory Models With Proofs

Jade Alglave<sup>1,2</sup> and Luc Maranget<sup>2</sup>

<sup>1</sup> University of Oxford <sup>2</sup> INRIA

**Abstract.** Concurrent programs running on weak memory models exhibit relaxed behaviours, making them hard to understand and to debug. We examine how to constrain the behaviour of such programs via synchronisation to ensure what we call their stability, i.e. that they behave as if they were running on a stronger model than the actual one, for example Sequential Consistency (SC). First, we define sufficient conditions ensuring stability to a program, and show that Power’s locks and read-modify-write primitives meet them. Second, we minimise the amount of required synchronisation by characterising which parts of a given execution should be synchronised. Third, we characterise the programs stable from a weak architecture to SC. Finally, we present the *offence* tool implementing our approach, by placing either lock-based or lock-free synchronisation in a x86 or Power program to ensure its stability.

Concurrent programs running on modern multiprocessors exhibit subtle behaviours, making them hard to understand and to debug: modern architectures (*e.g.* x86 or Power) provide *weak memory models*, allowing optimisations such as *instruction reordering*, *store buffering* or *write atomicity relaxation* [2]. Thus an execution of a program may not be an interleaving of its instructions, as it would be on a Sequentially Consistent (SC) architecture [17]. Hence standard analyses for concurrent programs might be unsound, as noted by M. Rinard in [21]. There exist a few memory model aware verification tools [11, 15, 20, 26], but they often focus on one model at a time, or cannot handle the write atomicity relaxation exhibited for example by Power: generality remains a challenge.

Fortunately, we can force a program running on a weak architecture to behave as if it were running on a stronger one (*e.g.* SC) by using *synchronisation primitives*. Hence, as observed by S. Burckhart and M. Musuvathi in [12], “*we can sensibly verify the relaxed executions [...] by solving the following two verification problems separately: 1. Use standard verification methodology for concurrent programs to show that the [SC] executions [...] are correct. 2. Use specialized methodology for memory model safety verification [...]*”. Here, *memory model safety* means checking that the executions of a program, although running on a weak architecture, are actually SC. To apply standard verification techniques to concurrent programs running on weak memory models, we thus first need to ensure that our programs have a SC behaviour. S. Burckhart and M. Musuvathi focus in [12] on memory model safety for TSO [24]. We generalise their idea to a wider class of models (the one defined in [5], and recalled in Sec. 1): we examine how to force a program running on a weak architecture  $A_1$  to behave as if running on a stronger one  $A_2$ , a property that we call *stability from  $A_1$  to  $A_2$* .

To ensure stability to a program, we examine the problem of placing *lock-based* or *lock-free* synchronisation primitives in a program. We call *synchronisation mapping* an insertion of synchronisation primitives (either *barriers* (or *fences*), *read-modify-writes*, or *locks*) in a program. We study whether a given synchronisation mapping ensures stability to a program running on a weak memory model, *e.g.* that we placed enough primitives in the code to ensure that it only has SC executions. D. Shasha and M. Snir proposed in [23] the *delay set analysis* to insert barriers in a program, but their work does not provide any semantics for weak memory models. Hence questions remain *w.r.t.* the adequacy of their method in the context of such models. On the contrary, locks allow the programmer to ignore the details of the memory model, thanks to the *data race free guarantee* (DRF guarantee) proposed in [3] by S. Adve and M. Hill.

Yet, from a compilation point of view, locks are costly. As noted by S. Adve and H.-J. Boehm in [4], “[o]n hardware that relaxes write atomicity [...], *e.g.* Power], it is often unclear that more efficient mappings (than the use of locks) are possible; even the fully fenced implementation may not be sequentially consistent.” Hence not only do we need to examine the *soundness* of our synchronisation mappings (*i.e.* that they actually ensure stability to a given program), but also their cost. We present here several new contributions:

1. We define in Sec. 2 sufficient conditions on synchronisation to ensure stability to a program. As an illustration, we provide in Sec. 3 semantics to the locks and read-modify-writes (rmw) of the Power architecture [1] (*i.e.* to the `lwarx` and `stwcx`. instructions) and show in Coq that they meet these conditions.
2. We propose along the way several synchronisation mappings, which we prove in Coq to enforce a SC behaviour to an x86 or Power program.
3. We optimise these mappings by generalising in Sec. 4 the approach of [23] to weak memory models and both lock-based and lock-free synchronisation, and characterise in Coq the executions stable from a weak architecture to SC.
4. We describe in Sec. 5 our new `offence` tool, which places either lock-based or lock-free synchronisation in a x86 or Power assembly program to ensure its stability, following the aforementioned characterisation. We detail how we used `offence` to test and measure the cost of our synchronisation mappings.

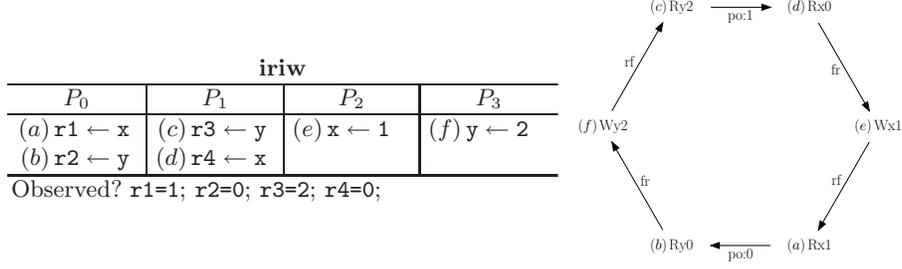
We formalised all our results in Coq. The Coq development<sup>2</sup>, the documentation and sources of `offence` and all the experimental details can be found online<sup>1</sup>.

## 1 Context

We give here the technical background on which we build our results and proofs. This section summarises the generic model [5] on which we build, which embraces SC [17], Sun TSO, PSO and RMO [24], Alpha [7] and a fragment of Power [1].

*Executions* An *event*  $e$  is a read or a write, composed of a direction  $R$  (read) or  $W$  (write), a location  $\text{loc}(e)$ , the instruction from which it comes  $\text{ins}(e)$ , a value

<sup>1</sup> <http://diy.inria.fr/offence> <sup>2</sup> <http://moscova.inria.fr/~alglave/stability>



**Fig. 1.** The **iriw** test and a non-SC execution

$\text{val}(e)$ , an originating processor  $\text{proc}(e)$ , and a unique identifier. We represent each instruction by the events it issues. In Fig. 1, we associate the store  $(e)$   $x \leftarrow 1$  on  $P_2$  with the event  $(e)Wx1$ . We write  $\mathbb{E}$  for the set of events, and  $\mathbb{W}$  (resp.  $\mathbb{R}$ ) for the subset of write (resp. read) events. We write  $w$  (resp.  $r$ ) for a write (resp. read), and  $m$  or  $e$  when the direction is irrelevant.

We associate a program with an *event structure*  $E \triangleq (\mathbb{E}, \overset{\text{po}}{\rightarrow})$ , composed of its events  $\mathbb{E}$  and the *program order*  $\overset{\text{po}}{\rightarrow}$ , a per-processor total order over  $\mathbb{E}$ . In Fig. 1, the read  $(a)$  from  $x$  on  $P_0$  is in program order with the read  $(b)$  from  $y$  on  $P_0$ , *i.e.*  $(a)Rx1 \overset{\text{po}}{\rightarrow} (b)Ry0$ . The  $\overset{\text{dp}}{\rightarrow}$  relation (included in  $\overset{\text{po}}{\rightarrow}$ , the source being a read) models the dependencies between instructions, *e.g.* when we compute the address of a load or store from the value of a preceding load.

Given an event structure  $E$ , we represent an execution  $X \triangleq (\overset{\text{ws}}{\rightarrow}, \overset{\text{rf}}{\rightarrow})$  of the corresponding program by two relations over  $\mathbb{E}$ . The *write serialisation*  $\overset{\text{ws}}{\rightarrow}$  is a per-location total order on writes modeling the *memory coherence* assumed by modern architectures [13], linking a write  $w$  to any write  $w'$  to the same location hitting the memory after  $w$ . The *read-from map*  $\overset{\text{rf}}{\rightarrow}$  links a write  $w$  to a read  $r$  from the same location that reads from  $w$ . We derive the *from-read map*  $\overset{\text{fr}}{\rightarrow}$  from  $\overset{\text{ws}}{\rightarrow}$  and  $\overset{\text{rf}}{\rightarrow}$ . A read  $r$  is in  $\overset{\text{fr}}{\rightarrow}$  with a write  $w$  when the write  $w'$  from which  $r$  reads hit the memory before  $w$  did:  $r \overset{\text{fr}}{\rightarrow} w \triangleq \exists w', w' \overset{\text{rf}}{\rightarrow} r \wedge w' \overset{\text{ws}}{\rightarrow} w$ . We write  $\overset{\text{com}}{\rightarrow} \triangleq \overset{\text{ws}}{\rightarrow} \cup \overset{\text{fr}}{\rightarrow} \cup \overset{\text{rf}}{\rightarrow}$  for the union of our *communication relations*.

In Fig. 1, the specified outcome corresponds to the execution on the right, if each location and register initially holds 0. If  $r1=1$  in the end, the read  $(a)$  read its value from the write  $(e)$  on  $P_2$ , hence  $(e) \overset{\text{rf}}{\rightarrow} (a)$ . If  $r2=0$ , the read  $(b)$  read its value from the initial state, thus before the write  $(f)$  on  $P_3$ , hence  $(b) \overset{\text{fr}}{\rightarrow} (f)$ . Similarly, we have  $(f) \overset{\text{rf}}{\rightarrow} (c)$  from  $r3=2$ , and  $(d) \overset{\text{fr}}{\rightarrow} (e)$  from  $r4=0$ .

*Architectures* In a shared-memory multiprocessor, a write may be committed first into a store buffer, then into a cache, and finally into memory. Hence, while a write transits in store buffers and caches, a processor may read a past value.

We model this by some subrelation of  $\xrightarrow{\text{rf}}$  being *non-global*: they can be ignored by some processors. We write  $\xrightarrow{\text{rfi}}$  (resp.  $\xrightarrow{\text{rfe}}$ ) for the *internal* (resp. *external*) read-from map, *i.e.* a read-from map between two events from the same (resp. distinct) processor(s). Hence we model a read  $r$  by a processor  $P_0$  reading from a write  $w$  in  $P_0$ 's store buffer by  $w \xrightarrow{\text{rfi}} r$  being non-global. When  $r$  reads from a write  $w$  by a distinct processor  $P_1$  into a cache shared by  $P_0$  and  $P_1$  only (a case of *write atomicity* relaxation [2]),  $w \xrightarrow{\text{rfe}} r$  is non-global, and  $w$  is said to be *non-atomic*. TSO authorises *e.g.* store buffering (*i.e.*  $\xrightarrow{\text{rfi}}$  is non-global) but considers stores to be atomic (*i.e.*  $\xrightarrow{\text{rfe}}$  is global). We write  $\xrightarrow{\text{grf}}$  for the global subrelation of  $\xrightarrow{\text{rf}}$ . We consider  $\xrightarrow{\text{ws}}$  and  $\xrightarrow{\text{fr}}$  global, since  $\xrightarrow{\text{ws}}$  is the order in which the writes to a certain location hit the memory.

Moreover, some pairs of events in the program order may be reordered. Thus only a subset of the pairs of events in  $\xrightarrow{\text{po}}$ , gathered in a subrelation  $\xrightarrow{\text{ppo}}$  (*preserved program order*), is guaranteed to occur in this order. TSO for example authorises write-read pairs to be reordered, but nothing else:  $\xrightarrow{\text{ppo}} = \xrightarrow{\text{po}} \setminus (\xrightarrow{\text{po}} \cap (\mathbb{W} \times \mathbb{R}))$ .

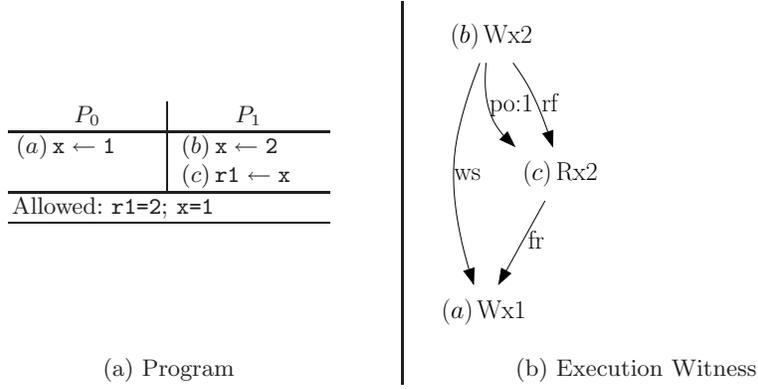
Finally, architectures provide barrier instructions to order certain pairs of events. We gather the orderings induced by barriers in the global relation  $\xrightarrow{\text{ab}}$ . Following [5], the relation  $\xrightarrow{\text{fence}} \subseteq \xrightarrow{\text{po}}$  induced by a barrier **fence** is *non-cumulative* when it orders certain pairs of events surrounding the barrier. For example, the x86 **mfence** barrier is a non-cumulative barrier ordering write-read pairs only: we have  $(w \xrightarrow{\text{mfence}} r) \Rightarrow (w \xrightarrow{\text{ab}} r)$ . Power **bne; isync** sequence is a non-cumulative barrier for read-read and read-write pairs: we have  $(r \xrightarrow{\text{bne;isync}} m) \Rightarrow (r \xrightarrow{\text{ab}} m)$ .

The relation  $\xrightarrow{\text{fence}}$  is *A-cumulative* (resp. *B-cumulative*) when it makes the writes atomic (*e.g.* by flushing the store buffers and caches). For example, Power **sync** barrier is non-, A- and B-cumulative for all pairs: we have  $(m_1 \xrightarrow{\text{sync}} m_2)$  (resp.  $(m_1 \xrightarrow{\text{rf}} w \xrightarrow{\text{sync}} m_2)$ ,  $(m_1 \xrightarrow{\text{sync}} w \xrightarrow{\text{rf}} m_2)$ ) implies  $(m_1 \xrightarrow{\text{ab}} m_2)$ . Power **lwsync** is non-, A- and B-cumulative for all pairs except write-read ones; we have  $(m_1 \xrightarrow{\text{lwsync}} m_2)$  (resp.  $(m_1 \xrightarrow{\text{rf}} w \xrightarrow{\text{lwsync}} m_2)$ ,  $(m_1 \xrightarrow{\text{lwsync}} w \xrightarrow{\text{rf}} m_2)$ ) implies  $(m_1 \xrightarrow{\text{ab}} m_2)$  except when  $(m_1, m_2) \in (\mathbb{W} \times \mathbb{R})$ .

An *architecture*  $A \triangleq (\text{ppo}, \text{grf}, \text{ab})$  specifies the function  $\text{ppo}$  (resp.  $\text{grf}, \text{ab}$ ) returning the relation  $\xrightarrow{\text{ppo}}$  (resp.  $\xrightarrow{\text{grf}}, \xrightarrow{\text{ab}}$ ) when given an execution.

*Validity* The  $\text{uniproc}(E, X) \triangleq \text{acyclic}(\xrightarrow{\text{com}} \cup \xrightarrow{\text{po-loc}})$  condition (where  $\xrightarrow{\text{po-loc}}$  is the program order restricted to events with the same location) forces a processor in a multiprocessor context to respect the memory *coherence* [13]. The  $\text{thin}(E, X) \triangleq \text{acyclic}(\xrightarrow{\text{rf}} \cup \xrightarrow{\text{dr}})$  condition prevents executions where values seem to come *out of thin air* [19]. We define the *global happens-before* relation  $A.\text{ghb}(E, X)$  of an execution  $(E, X)$  on an architecture  $A$  as the union of the relations global on  $A$ :

$$A.\text{ghb}(E, X) \triangleq \xrightarrow{\text{ws}} \cup \xrightarrow{\text{fr}} \cup \xrightarrow{\text{ppo}} \cup \xrightarrow{\text{grf}} \cup \xrightarrow{\text{ab}}$$



**Fig. 2.** A program and an execution witness

An execution  $(E, X)$  is *valid* on an architecture  $A$ , written  $A.\text{valid}(E, X)$ , when the relation  $A.\text{ghb}(E, X)$  is acyclic (together with the two above checks):

$$A.\text{valid}(E, X) \triangleq \text{uniproc}(E, X) \wedge \text{thin}(E, X) \wedge \text{acyclic}(A.\text{ghb}(E, X))$$

*A hierarchy of architectures* We consider an architecture  $A_1$  to be *weaker* than an architecture  $A_2$ , written  $A_1 \leq A_2$ , when  $A_1$  authorises at least all the executions valid on  $A_2$  (writing  $\overset{r_i}{\rightarrow}$  for the relation  $\overset{r}{\rightarrow}$  w.r.t.  $A_i$ ):

$$A_1 \leq A_2 \triangleq \text{pp}^{\text{po}1} \subseteq \text{pp}^{\text{po}2} \wedge \overset{\text{grf}_1}{\rightarrow} \subseteq \overset{\text{grf}_2}{\rightarrow}$$

For example, TSO authorises store buffering, which means the internal  $\overset{\text{rf}}{\rightarrow}$  is not global. Hence for TSO,  $\overset{\text{grf}}{\rightarrow}$  is restricted to the external one, *i.e.*  $\overset{\text{grf}}{\rightarrow} = \overset{\text{rfe}}{\rightarrow}$ . Moreover, TSO authorises the reordering of write-read pairs, hence its  $\text{pp}^{\text{po}}$  is strictly included in  $\text{pp}^{\text{po}}$ . Therefore, TSO is weaker than SC.

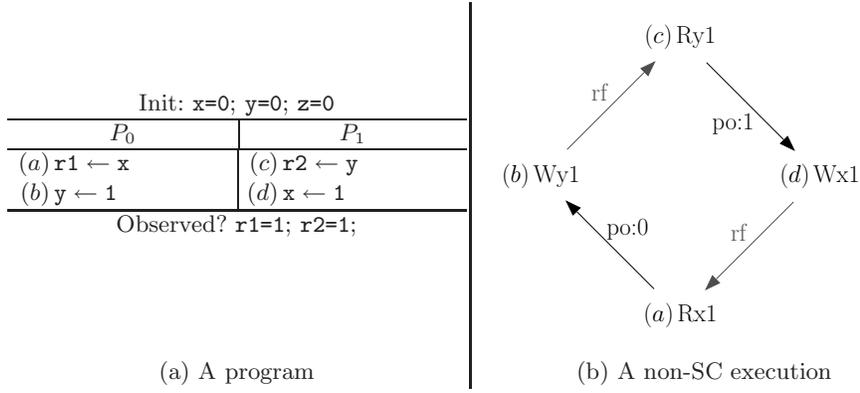
As shown in [5], a weak execution (*i.e.* valid on a weak architecture  $A_1$ ) is valid on a stronger one  $A_2$  if and only if:

**Lem. 1 (Characterisation).**

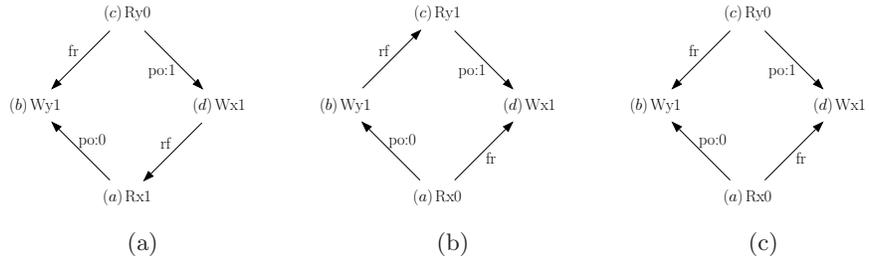
$$\forall A_1 \leq A_2, \forall EX, (A_1.\text{valid}(E, X) \wedge \text{acyclic}(A_2.\text{ghb}(E, X))) \Leftrightarrow A_2.\text{valid}(E, X)$$

For example, the execution  $(E, X)$  of Fig. 2(b) is valid on TSO and  $SC.\text{ghb}(E, X)$  is acyclic, hence  $(E, X)$  is also valid on SC.

*Building a valid execution from an order* We consider two relations to be *compatible* when their union is acyclic. Consider an architecture  $A$  without barriers, *i.e.*  $\overset{\text{ab}^A}{\rightarrow} = \emptyset$ . The characterisation above means in particular that, for any event structure  $E$ , one can build an execution, associated with  $E$  and valid on  $A$ , from a total order  $\overset{\circ}{\rightarrow}$  on  $\text{evts}(E)$  compatible with  $A.\text{ppo}(E)$ .



**Fig. 3.** A program and a non-SC execution



**Fig. 4.** SC executions for the test of Fig. 3(a)

Consider *e.g.* the event structure  $(\{(a), (b), (c), (d)\}, \{(a) \xrightarrow{\text{po}} (b), (c) \xrightarrow{\text{po}} (d)\})$  associated with the program of Fig. 3. On SC we have  $(a) \xrightarrow{\text{ppo}} (b)$  and  $(c) \xrightarrow{\text{ppo}} (d)$ . Hence we can build a valid SC execution from the order  $(a) \xrightarrow{\circ} (b) \xrightarrow{\circ} (c) \xrightarrow{\circ} (d)$ , which is the one we give in Fig. 4(b). The first write in the order  $\xrightarrow{\circ}$  is  $(b)$ , a write to  $y$ , which is immediately followed by the read  $(c)$  to  $y$  with the same value, hence we have  $(b) \xrightarrow{\text{rf}} (c)$ . There is no write preceding the read  $(a)$  from  $x$ , hence  $(a)$  reads from the initial state. Moreover, this initial write to  $x$  precedes the write  $(d)$  in  $\xrightarrow{\text{ws}}$ , hence  $(a) \xrightarrow{\text{fr}} (d)$ .

We write  $\text{rf}(\xrightarrow{\circ})$  (resp.  $\text{ws}(\xrightarrow{\circ})$ ) for the  $\xrightarrow{\text{rf}}$  (resp.  $\xrightarrow{\text{ws}}$ ) extracted from  $\xrightarrow{\circ}$ . We have  $(x, y) \in \text{rf}(\xrightarrow{\circ})$  when  $x$  is a write and  $y$  a read, both to the same location  $\ell$ , with the same value, and  $x$  is the maximal previous write to  $\ell$  before  $y$  in  $\xrightarrow{\circ}$ , *i.e.*  $\neg(\exists z \in \mathbb{W}_{\ell}, x \xrightarrow{\circ} z \xrightarrow{\circ} y)$ . We have  $(x, y) \in \text{ws}(\xrightarrow{\circ})$  when  $x$  and  $y$  are writes to the same location and  $x \xrightarrow{\circ} y$ . Formally, we have:

**Lem. 2 (Extraction of a valid execution from an order).**

$$\forall EX \xrightarrow{\circ}, \text{total-order}(\xrightarrow{\circ}, \text{evts}(E)) \wedge \text{acyclic}(\xrightarrow{\circ} \cup A.\text{ppo}(E)) \wedge \\ X = (\text{ws}(\xrightarrow{\circ}), \text{rf}(\xrightarrow{\circ})) \Rightarrow A.\text{valid}(E, X)$$

**Proof** Since  $\xrightarrow{\circ}$  is a total order,  $\text{ws}(\xrightarrow{\circ})$  is by definition a per-location total order on writes. Moreover,  $\text{rf}(\xrightarrow{\circ})$  trivially satisfies the definition of read-from map.

We define the extracted from-read map  $\text{fr}(\xrightarrow{\circ})$  as:

$$(r, w) \in \text{fr}(\xrightarrow{\circ}) \triangleq \exists w_r, (w_r, r) \in \text{rf}(\xrightarrow{\circ}) \wedge (w_r, w) \in \text{ws}(\xrightarrow{\circ})$$

We define the extracted communication  $\text{com}(\xrightarrow{\circ})$  as:

$$\text{com}(\xrightarrow{\circ}) \triangleq \text{ws}(\xrightarrow{\circ}) \cup \text{rf}(\xrightarrow{\circ}) \cup \text{fr}(\xrightarrow{\circ})$$

Since  $\text{ab}_A = \emptyset$ , we know that  $A.\text{ghb}(E, X)$  is included in  $\text{com}(\xrightarrow{\circ}) \cup A.\text{ppo}(E)$ .

By definition,  $\text{ws}(\xrightarrow{\circ})$  and  $\text{rf}(\xrightarrow{\circ})$  are included in  $\xrightarrow{\circ}$ . Let us show that  $\text{fr}(\xrightarrow{\circ})$  is included in  $\xrightarrow{\circ}$ : consider  $(r, w) \in \text{fr}(\xrightarrow{\circ})$ ; since  $\xrightarrow{\circ}$  is total, we have either  $r \xrightarrow{\circ} w$  or  $w \xrightarrow{\circ} r$ . Suppose  $w \xrightarrow{\circ} r$ ; since  $(r, w) \in \text{fr}(\xrightarrow{\circ})$ , there exists  $w_r$  such that  $(w_r, r) \in \text{rf}(\xrightarrow{\circ})$ , hence  $w_r \xrightarrow{\circ} r$ , and  $(w_r, w) \in \text{ws}(\xrightarrow{\circ})$ , hence  $w_r \xrightarrow{\circ} w$ . Thus we have  $w_r \xrightarrow{\circ} w \xrightarrow{\circ} r$ . By definition of  $\text{rf}(\xrightarrow{\circ})$ ,  $w_r$  is the maximal previous write to  $\text{loc}(r)$  before  $r$  in  $\xrightarrow{\circ}$ , hence a contradiction.

Hence  $A.\text{ghb}(E, X)$  is included in  $\xrightarrow{\circ} \cup A.\text{ppo}(E)$ . Moreover,  $\xrightarrow{\circ}$  is compatible with  $A.\text{ppo}(E)$  by hypothesis (*i.e.* their union is acyclic). Hence there cannot be any cycle in  $A.\text{ghb}(E, X)$ .  $\square$

In the following, we consider  $A_2$  to be without barriers, *i.e.*  $\text{ab}_2 = \emptyset$ . We write  $\xrightarrow{\text{ghb}_2}$  for  $A_2.\text{ghb}(E, X)$ .

## 2 Covering relations

We examine now how to force the executions of a program running on a weak architecture  $A_1$  to be valid on a stronger one  $A_2$ , which we call *stability from  $A_1$  to  $A_2$* , *i.e.* we examine when the following property holds for all  $(E, X)$ :

$$\text{stable}_{A_1, A_2}(E, X) \triangleq A_1.\text{valid}(E, X) \Rightarrow A_2.\text{valid}(E, X)$$

The execution of **iriw** in Fig. 1 is not stable from Power to SC, for it is valid on Power yet not on SC. We can stabilise an execution by using *synchronisation idioms*, *e.g.* barriers or locks. Synchronisation idioms *arbitrate conflicts* between accesses, *i.e.* ensure that one out of two conflicting accesses occurs before the other. We formalise this with an irreflexive *conflict* relation  $\overset{c}{\rightarrow}$  over events, such that  $\forall xy, x \overset{c}{\rightarrow} y \Rightarrow \neg(y \overset{po}{\rightarrow} x)$  and a *synchronisation* relation  $\overset{s}{\rightarrow}$  over events. An execution  $(E, X)$  is *covered* when  $\overset{c}{\rightarrow}$  is *arbitrated* by  $\overset{s}{\rightarrow}$ :

$$\text{covered}_{c,s}(E, X) \triangleq \forall xy, x \overset{c}{\rightarrow} y \Rightarrow x \overset{s}{\rightarrow} y \vee y \overset{s}{\rightarrow} x$$

We consider a relation  $\overset{s}{\rightarrow}$  to be *covering* when ordering by  $\overset{s}{\rightarrow}$  the conflicting accesses of an execution  $(E, X)$  valid on  $A_1$  guarantees its validity on  $A_2$ , *i.e.* the synchronisation  $\overset{s}{\rightarrow}$  arbitrates enough conflicts to enforce a strong behaviour:

$$\text{covering}(\overset{c}{\rightarrow}, \overset{s}{\rightarrow}) \triangleq \forall EX, (A_1.\text{valid}(E, X) \wedge \text{covered}_{c,s}(E, X)) \Rightarrow A_2.\text{valid}(E, X)$$

*Lock-based synchronisation* For example, the DRF guarantee [3] ensures that if the *competing accesses*, defined below, of an execution are ordered by locks, then this execution is SC. Hence locks are covering *w.r.t.* the competing accesses. Following [3], two events are *competing* if they are from two distinct processors, relative to the same location, and one of them at least is a write (*e.g.* in Fig. 1, the read  $(a)$  from  $x$  on  $P_0$  and the write  $(e)$  to  $x$  on  $P_2$ ):

$$m_1 \overset{\text{cmp}}{\leftrightarrow} m_2 \triangleq \text{proc}(m_1) \neq \text{proc}(m_2) \wedge \text{loc}(m_1) = \text{loc}(m_2) \wedge (m_1 \in \mathbb{W} \vee m_2 \in \mathbb{W})$$

We describe the ordering induced by locks by a relation  $\overset{\text{lock}}{\rightarrow}$  (instantiated in Sec. 3.1) over  $\mathbb{E}$  compatible with  $\overset{\text{com}}{\rightarrow}$  (*i.e.*  $\text{acyclic}(\overset{\text{lock}}{\rightarrow} \cup \overset{\text{com}}{\rightarrow})$ ), corresponding in Fig. 1 to placing locks to a variable  $\ell_1$  on the accesses  $(a)$ ,  $(d)$  and  $(e)$  relative to  $x$ , and locks to a different variable  $\ell_2$  on the accesses  $(b)$ ,  $(c)$  and  $(f)$  relative to  $y$ . Thus we have a cycle in  $\overset{\text{lock}}{\rightarrow} \cup \overset{po}{\rightarrow}$ :  $(a) \overset{po}{\rightarrow} (b) \overset{\text{lock}}{\rightarrow} (f) \overset{\text{lock}}{\rightarrow} (c) \overset{po}{\rightarrow} (d) \overset{\text{lock}}{\rightarrow} (e) \overset{\text{lock}}{\rightarrow} (a)$ . If  $\overset{\text{lock}}{\rightarrow} \cup \overset{po}{\rightarrow}$  is acyclic, then the execution of Fig. 1 is forbidden. Formally, we have:

**Lem. 3.**  $\text{acyclic}(\overset{\text{lock}}{\rightarrow} \cup \overset{po}{\rightarrow}) \Rightarrow \text{covering}(\overset{\text{cmp}}{\leftrightarrow}, (\overset{\text{lock}}{\rightarrow} \cup \overset{po}{\rightarrow})^+)$

**Proof** Let  $(E, X)$  be valid on  $A_1$  and covered. Suppose by contradiction that there is a cycle in  $\overset{\text{ghb}_2}{\rightarrow}$ , which is by definition a cycle in  $\overset{\text{ws}}{\rightarrow} \cup \overset{\text{fr}}{\rightarrow} \cup \overset{\text{grf}_2}{\rightarrow} \cup \overset{\text{ppo}_2}{\rightarrow}$ . The events in  $\overset{\text{ws}}{\rightarrow}$ ,  $\overset{\text{fr}}{\rightarrow}$  or  $\overset{\text{rf}}{\rightarrow}$  and from distinct processors are competing. Since  $(E, X)$  is covered,  $\overset{\text{sync}}{\rightarrow}$  orders the competing accesses according to  $\overset{\text{com}}{\rightarrow}$ . The remaining events in  $\overset{\text{com}}{\rightarrow}$  belong to the same processor, hence are in  $\overset{po}{\rightarrow}$  by uniproc. Moreover, we know  $\overset{\text{ppo}_2}{\rightarrow} \subseteq \overset{po}{\rightarrow}$ . Hence a cycle in  $\overset{\text{ws}}{\rightarrow} \cup \overset{\text{fr}}{\rightarrow} \cup \overset{\text{grf}_2}{\rightarrow} \cup \overset{\text{ppo}_2}{\rightarrow}$  is a cycle in  $\overset{\text{sync}}{\rightarrow} \cup \overset{po}{\rightarrow}$ , which contradicts the irreflexivity of  $\overset{\text{sync}}{\rightarrow}$ .  $\square$

This lemma leads to a mapping which we call L (for locks), which simply places a lock by the same lock variable on each side of a given conflict edge; following Lem. 3, it gives a SC behaviour to a program.

*Lock-free synchronisation* We give here an example of a covering lock-free synchronisation relation. A program can distinguish between two architectures  $A_1 \leq A_2$  for one of two reasons. First, if the program involves a pair  $(x, y)$  maintained in program order on  $A_2$  (i.e.  $x \xrightarrow{\text{pp}o_2} y$ ) but not on  $A_1$  (i.e.  $\neg(x \xrightarrow{\text{pp}o_1} y)$ ). In Fig. 1, we have  $(a) \xrightarrow{\text{po}} (b)$ . Hence on a strong architecture  $A_2$  such as SC where  $\text{pp}o_2 = \text{po}$ , we have  $(a) \xrightarrow{\text{pp}o_2} (b)$ . On a weak architecture  $A_1$  such as Power, where the read-read pairs in program order are not maintained, we have  $\neg((a) \xrightarrow{\text{pp}o_1} (b))$ .

Second, if the program reads from a write atomic on  $A_2$  but not on  $A_1$ . In Fig. 1, we have  $(e) \xrightarrow{\text{rfe}} (a)$ . On a strong architecture  $A_2$  such as SC where the writes are atomic, i.e.  $\text{grf} = \text{rf}$ , we have  $(e) \xrightarrow{\text{grf}} (a)$ . On a weak architecture  $A_1$  such as Power, which relaxes write atomicity, we have  $\neg((e) \xrightarrow{\text{grf}} (a))$ . We call such reads *fragile reads* and define them as  $(\xrightarrow{\text{r}_2 \setminus \text{r}_1} \triangleq \xrightarrow{\text{r}_2} \setminus \xrightarrow{\text{r}_1}$  being the set difference):

$$\text{fragile}(r) \triangleq \exists w, w \xrightarrow{\text{grf}_2 \setminus \text{r}_1} r$$

We consider such differences between architectures as conflicts, and formalise this notion as follows. We consider that two events form a *fragile pair* (written  $\xrightarrow{\text{frag}}$ ) if they are maintained in the program order on  $A_2$ , and either they are not maintained in the program order on  $A_1$ , or the first event is a fragile read:

$$m_1 \xrightarrow{\text{frag}} m_2 \triangleq m_1 \xrightarrow{\text{pp}o_2} m_2 \wedge (\neg(m_1 \xrightarrow{\text{pp}o_1} m_2) \vee \text{fragile}(m_1))$$

An execution is covered if the barrier relation  $\xrightarrow{\text{ab}_1}$  arbitrates the fragile pairs. In Fig. 1, this corresponds to placing a barrier between  $(c)$  and  $(d)$  on  $P_1$ , i.e.  $(c) \xrightarrow{\text{ab}_1} (d)$ , and another barrier between  $(a)$  and  $(b)$  on  $P_0$ , i.e.  $(a) \xrightarrow{\text{ab}_1} (b)$ . Hence we have a cycle in  $\xrightarrow{\text{ab}_1} \cup \xrightarrow{\text{rf}}$ :  $(d) \xrightarrow{\text{rfe}} (a) \xrightarrow{\text{ab}_1} (b) \xrightarrow{\text{rfe}} (c) \xrightarrow{\text{ab}_1} (d)$ . If  $\xrightarrow{\text{ab}_1}$  is *A-cumulative w.r.t.  $\text{grf}_2 \setminus \text{r}_1$*  (i.e.  $\forall xyz, (x \xrightarrow{\text{grf}_2 \setminus \text{r}_1} y \wedge y \xrightarrow{\text{ab}_1} z) \Rightarrow x \xrightarrow{\text{ghb}_1} z$ ), we create a cycle in  $\xrightarrow{\text{ghb}_1}$ , which forbids the execution:  $(d) \xrightarrow{\text{ghb}_1} (b) \xrightarrow{\text{ghb}_1} (d)$ . Indeed, we show that if  $\xrightarrow{\text{ab}_1}$  is *A-cumulative w.r.t.  $\text{grf}_2 \setminus \text{r}_1$*  then  $\xrightarrow{\text{ab}_1}$  is covering:

**Lem. 4.**  $(\forall xyz, (x \xrightarrow{\text{grf}_2 \setminus \text{r}_1} y \wedge y \xrightarrow{\text{ab}_1} z) \Rightarrow x \xrightarrow{\text{ghb}_1} z) \Rightarrow \text{covering}(\xrightarrow{\text{frag}}, \xrightarrow{\text{ab}_1})$

**Proof** Consider an execution  $(E, X)$  valid on  $A_1$  and covered. Suppose by contradiction that there is a cycle in  $\xrightarrow{\text{ghb}_2}$ , which is by definition a cycle in  $\xrightarrow{\text{ws}} \cup \xrightarrow{\text{fr}} \cup \xrightarrow{\text{grf}_2} \cup \xrightarrow{\text{pp}o_2}$ . Since  $\xrightarrow{\text{ws}}$ ,  $\xrightarrow{\text{fr}}$  and  $\xrightarrow{\text{pp}o_1}$  are included in  $\xrightarrow{\text{ghb}_1}$ , this cycle is a cycle in  $\xrightarrow{\text{ghb}_1} \cup \xrightarrow{\text{pp}o_2 \setminus \text{r}_1} \cup (\xrightarrow{\text{grf}_2 \setminus \text{r}_1}; \xrightarrow{\text{pp}o_2})$ . Since  $\xrightarrow{\text{sif}}$  orders all fragile pairs, and is compatible with  $\xrightarrow{\text{pp}o_2}$ , we know that  $\xrightarrow{\text{pp}o_2 \setminus \text{r}_1}$  is included in  $\xrightarrow{\text{sif}}$  and  $(\xrightarrow{\text{grf}_2 \setminus \text{r}_1}; \xrightarrow{\text{pp}o_2}) \subseteq (\xrightarrow{\text{grf}_2 \setminus \text{r}_1}; \xrightarrow{\text{sif}})$ . Since  $\xrightarrow{\text{sif}}$  is *A-cumulative*,  $(\xrightarrow{\text{grf}_2}; \xrightarrow{\text{sif}})$  is in  $\xrightarrow{\text{ab}_1}$ , hence in  $\xrightarrow{\text{ghb}_1}$ . Thus there is a cycle in  $\xrightarrow{\text{sif}} \cup \xrightarrow{\text{ghb}_1}$ , which contradicts their compatibility.  $\square$

Arch.	Fragile pair	Barriers (mapping F)
Power	$r \xrightarrow{\text{po}} r$	$r \xrightarrow{\text{sync}} r$ (need A-cumulativity)
	$r \xrightarrow{\text{po}} w$	$r \xrightarrow{\text{lwsync}} w$ (A-cumulativity OK)
	$w \xrightarrow{\text{po}} w$	$w \xrightarrow{\text{lwsync}} w$ (no need for A-cumulativity)
	$w \xrightarrow{\text{po}} r$	$w \xrightarrow{\text{sync}} r$ (need for write-read non-cumulativity)
x86	$w \xrightarrow{\text{po}} r$	$w \xrightarrow{\text{mfence}} r$ (need for write-read non-cumulativity)

Fig. 5. Mapping F: barriers

This lemma leads to a mapping which we call F (for fences), given in Fig. 5. This mapping places a barrier between each fragile pair of a program; following Lem. 4, it gives a SC behaviour to this program. Recall that we give the semantics of the barriers that we use in the mapping F in Sec. 1, § *Architectures*, on p. 4.

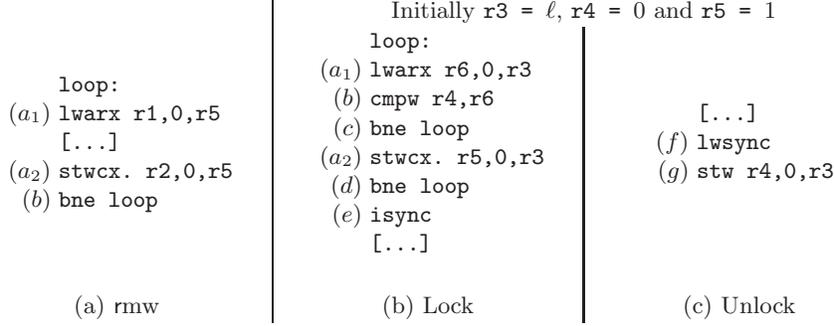
In x86, stores are atomic, and only the write-read pairs in program order are not preserved, *i.e.* the fragile pairs are the write-read pairs  $w \xrightarrow{\text{po}} r$ . We do not need cumulativity in x86, *i.e.* we only need a non-cumulative write-read barrier, the **mfence** barrier:  $w \xrightarrow{\text{mfence}} r$ .

In Power, no pair is preserved in program order except the read-read and read-write pairs with a dependency between the accesses [5]. But since stores are not atomic, even the dependent read-read and read-write pairs are fragile. For a read-read pair  $r_1 \xrightarrow{\text{po}} r_2$ , since  $r_1$  can read from a non-atomic write  $w$ , we need a cumulative barrier between  $r_1$  and  $r_2$ . But **lwsync** does not order write to read chains, *i.e.* **lwsync** between  $r_1$  and  $r_2$  will not order  $w$  and  $r_2$ . Therefore we need a **sync**:  $r_1 \xrightarrow{\text{sync}} r_2$ . For a read-write pair  $r \xrightarrow{\text{po}} w$ , we need a cumulative barrier as well, but **lwsync** is sufficient here, for it will order the read  $w'$  from which  $r$  may read and  $w$ . In the write-write and write-read cases, there is no need for cumulativity. In the write-write case, a **lwsync** is enough, for it orders write-write pairs; but in the write-read case, we need a **sync**.

### 3 Synchronisation idioms

As an illustration to Sec. 2, we now study the semantics of Power’s locks and `rmw` [1]. As noted by S. Adve and H.-J. Boehm in [4] “[o]n hardware that relaxes write atomicity [...], such as Power] even the fully fenced implementation may not be sequentially consistent.” Thus it is unclear whether the synchronisation primitives provided by the architecture actually restore SC: it could perfectly be the architect’s intent (*e.g.* **lwsync** is not strong enough to restore SC, but is faster than **sync**, as we show in Sec. 5), or a bug in the implementation [5]. Hence we need to define the semantics of the synchronisation primitives given in the documentation, and study whether they allow us to restore SC, *i.e.* that we can use them to build covering relations, as defined in Sec. 2.

We first define *atomic pairs*, which are the stepping stone to build locks, studied in Sec. 3.1 and `rmw`, studied in Sec. 3.2. We show how to use these primitives



**Fig. 6.** Read-modify-write, lock and unlock in Power

to build covering relations. Second, because cumulativeness might be too costly in practice, or its implementation challenging, we propose in Sec. 3.2 two lock-free mappings restoring a strong architecture from Power without using cumulativeness, as an alternative to the mapping F (see Sec. 2) which uses cumulativeness.

*Atomicity* Fig. 6(a) gives a generic Power rmw. The `lwarx` ( $a_1$ ) loads from its source address (in register `r5`) and *reserves* it. Any subsequent store to the reserved address from another processor and any subsequent `lwarx` from the same processor invalidates the reservation. The `stwcx.` ( $a_2$ ) checks if the reservation is valid; if so, it is said to be *successful*: it stores into the reserved address and the code exits the loop. Otherwise, the `stwcx.` does not store and the code loops. Thus these instructions ensure *atomicity* to the code they surround (provided this code does not contain any `lwarx` nor `stwcx.`), as no other processor can write to the reserved location between the `lwarx` and the successful `stwcx.`

We distinguish the read and write events issued by such instructions from the plain ones: we write  $\mathbb{R}^*$  (resp.  $\mathbb{W}^*$ ) for the subset of  $\mathbb{R}$  (resp.  $\mathbb{W}$ ) issued by a `lwarx` (resp. a successful `stwcx.`), and define two events  $r$  and  $w$  to form an atomic pair *w.r.t.* a location  $\ell$  if (a)  $w$  was issued by a successful `stwcx.` to  $\ell$ , (b)  $r$  was issued by the last `lwarx` (in  $\xrightarrow{\text{po}}$ ) from  $\ell$  before the `stwcx.` that issued  $w$ , and (c) no other processor wrote to  $\ell$  between  $r$  and  $w$ :

$$\text{atom}(r, w, \ell) \triangleq r \in \mathbb{R}^* \wedge w \in \mathbb{W}^* \wedge \text{loc}(r) = \text{loc}(w) = \ell \wedge \quad (a)$$

$$r = \max_{\text{po}}(\{m \mid m \in (\mathbb{R}^* \cup \mathbb{W}^*) \wedge m \xrightarrow{\text{po}} w\}) \wedge \quad (b)$$

$$\neg(\exists w' \in \mathbb{W}, \text{proc}(w') \neq \text{proc}(r) \wedge \text{loc}(w') = \ell \wedge r \xrightarrow{\text{fr}} w' \xrightarrow{\text{ws}} w) \quad (c)$$

### 3.1 Locks

Atomic pairs are used *e.g.* in *lock* and *unlock* primitives [1, App. B]. The idiomatic Power lock (resp. unlock) is shown in Fig. 6(b) (resp. Fig. 6(c)).



float above a read issued by a `lwarx`: in Fig. 7, the event  $m_2$  in  $cs_2$  is in  $\xrightarrow{\text{ghb}_1}$  with the read ( $a_1$ ) from its Lock's `lwarx`. Hence the read  $r$  of a lock's `lwarx` satisfies the import predicate when no access  $m$  after  $r$  can be speculated before  $r$ :

$$\text{import}(r) \triangleq \forall rm, (r \in \mathbb{R}^* \wedge r \xrightarrow{\text{po}} m) \Rightarrow (r \xrightarrow{\text{ab}_1} m)$$

Fig. 6(c) shows Power's unlock. It starts (line ( $f$ )) with an *export barrier* [1, p. 722] (here a `lwsync`). The export barrier forces the accesses before the write  $w$  of the unlock to be committed to memory before the next lock primitive takes the lock: in Fig. 7, the event  $m_1$  in  $cs_1$  is in  $\xrightarrow{\text{ghb}_1}$  with the read ( $a_1$ ) of  $cs_2$ 's Lock. This means that we define an export barrier to be B-cumulative, but only *w.r.t.* reads issued by the `lwarx` of an atomic pair:

$$\text{export}(w) \triangleq \forall rm, (r \in \mathbb{R}^* \wedge (m \xrightarrow{\text{po}} w \xrightarrow{\text{rf}} r)) \Rightarrow (m \xrightarrow{\text{ab}_1} r)$$

Then a store to the lock variable (line ( $g$ )), or more precisely the next write event to  $\ell$  in program order after a lock acquisition, frees the lock:

$$\begin{aligned} \text{free}(\ell, r, w) \triangleq & w \in \mathbb{W} \wedge \text{loc}(w) = \ell \wedge r \xrightarrow{\text{po}} w \wedge \text{taken}(\ell, r) \wedge \\ & \neg(\exists w' \in \mathbb{W}, \text{loc}(w') = \ell \wedge r \xrightarrow{\text{po}} w' \xrightarrow{\text{po}} w) \end{aligned}$$

A lock primitive thus consists of a taken operation (see Fig. 6(b), lines ( $a_1$ ) to ( $a_2$ )) followed by an import barrier. An unlock consists of an export barrier (line ( $f$ )) followed by a write freeing the lock (line ( $g$ )):

$$\text{Lock}(\ell, r) \triangleq \text{taken}(\ell, r) \wedge \text{import}(r)$$

$$\text{Unlock}(\ell, r, w) \triangleq \text{free}(\ell, r, w) \wedge \text{export}(w)$$

We show that this semantics ensures the acyclicity of  $\xrightarrow{\text{lock}} \cup \xrightarrow{\text{po}}$ , *i.e.* following Lem. 3,  $(\xrightarrow{\text{lock}} \cup \xrightarrow{\text{po}})^+$  is covering for the competing accesses. Hence locks on the competing accesses ensures a SC behaviour to Power programs:

**Lem. 5.**  $\forall EX, A_1. \text{valid}(E, X) \Rightarrow \text{acyclic}(\xrightarrow{\text{lock}} \cup \xrightarrow{\text{po}})$

**Proof** Suppose by contradiction a cycle in  $\xrightarrow{\text{lock}} \cup \xrightarrow{\text{po}}$ . This cycle is a cycle in  $(\xrightarrow{\text{lock}})^+; \xrightarrow{\text{po}}$  since  $\xrightarrow{\text{po}}$  is transitive. Let us show by induction that any path of  $(\xrightarrow{\text{lock}})^+; \xrightarrow{\text{po}}$  from an event  $x$  to an event  $y$  is a path in  $(\xrightarrow{\text{ghb}_1})^+$ . Hence, a cycle is in  $(\xrightarrow{\text{lock}})^+; \xrightarrow{\text{po}}$  is a cycle in  $(\xrightarrow{\text{ghb}_1})^+$ , which contradicts the validity of  $(E, X)$  on  $A_1$ .

Consider the base case with three events  $m_1 \xrightarrow{\text{lock}} m_2 \xrightarrow{\text{po}} m_3$ . Let us do an induction over  $m_1 \xrightarrow{\text{lock}} m_2$ . Consider the base case where  $m_1$  and  $m_2$  belong respectively to the critical sections  $cs_1$  and  $cs_2$ , such that  $cs_1 \xrightarrow{\text{css}} cs_2$ .

In this case,  $m_3$  is in  $\xrightarrow{\text{po}}$  after  $cs_2$ 's import barrier, which prevents any event to float above the read issued by  $cs_2$ 's `lwarx`. Thus we have  $R(cs_2) \xrightarrow{\text{ab}_1} m_3$ , hence  $R(cs_2) \xrightarrow{\text{ghb}_1} m_3$ . Moreover,  $m_1$  is in  $\xrightarrow{\text{po}}$  before  $cs_1$ 's export barrier (*i.e.*  $m_1 \xrightarrow{\text{ab}_1} W(cs_1)$ ) hence  $m_1 \xrightarrow{\text{ghb}_1} W(cs_1)$ . Since  $W(cs_1) \xrightarrow{\text{rfe}} R(cs_2)$ , by B-cumulativity of  $cs_1$ 's export barrier, we have  $m_1 \xrightarrow{\text{ab}_1} R(cs_2)$ , hence  $m_1 \xrightarrow{\text{ghb}_1} R(cs_2)$ . Thus  $m_1 (\xrightarrow{\text{ghb}_1})^+ m_3$ .

The transitive cases follows by induction.  $\square$

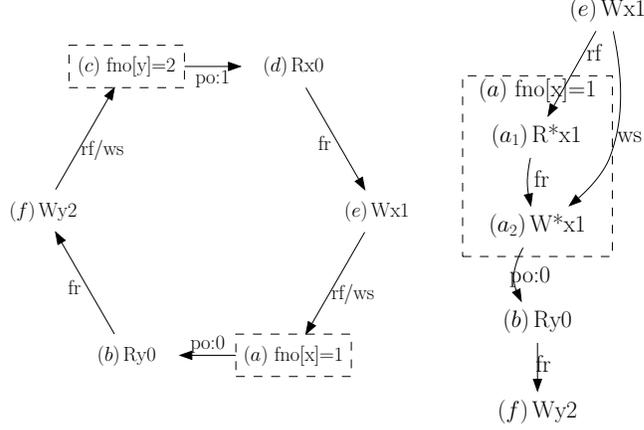


Fig. 8. (a) **iriw** after P mapping      5.(b) Opening fno on  $P_0$

Our import barrier allows events to be delayed so that they are performed inside the critical section. Our export barrier allows the events after the unlock to be speculated before the lock is released. Such relaxed semantics already exist for high-level lock and unlock primitives [8, 22].

In the documentation [1, p. 721], the import barrier is a sequence **bne**; **isync** (*i.e.* a read-read, read-write non-cumulative barrier) or a **lwsync**, *i.e.* cumulative [1, p.721]. Lem. 5 shows that the first one is enough, for our import barrier does not need cumulativity. The export barrier is a **sync** (*i.e.* cumulative for all pairs) or a **lwsync** [1, p. 722]. Lem. 5 shows that we only need a B-cumulative barrier towards reads issued by a **lwarx**, *i.e.* a **sync** is unnecessarily costly. Moreover, although a **lwsync** is not B-cumulative towards plain reads, its implementations appear experimentally to treat the reads issued by the **lwarx** of an atomic pair specially. We tested this semantics of **lwsync** with our **diy** tool [5], and ran our automatically generated tests up to  $10^9$  times each (see the logs online<sup>1</sup>). Hence our semantics for the export barrier is experimentally sound.

### 3.2 Read-modify-write primitives

By Lem. 4, we can restore SC in the **iriw** test of Fig. 1 using A-cumulative barriers between the fragile pairs (a) and (b) on  $P_0$ , and (c) and (d) on  $P_1$ . Yet, cumulativity may be challenging to implement or too costly in practice [5]. We propose a mapping of certain reads to rmw (as in Fig. 6(a)), and show that this restores a strong architecture from a weaker one without using cumulativity.

In Fig. 8(a), we replaced the fragile reads (a) and (c) of **iriw** by rmw: we say these fragile reads are *protected* (a notion defined below). In the example we use *fetch and no-op* (fno) primitives [1, p.719] to implement atomic reads. Yet, our results hold for any kind of rmw. We show that when the fragile reads are

protected, we do not need cumulative barriers, but just non-cumulative ones. If a read is protected by a rmw, then the rmw compensates the need for cumulativity by enforcing enough order to the write from which the protected read reads.

*Protecting the fragile reads with rmw* We consider that two events  $r$  and  $w$  form a rmw *w.r.t.* a location  $\ell$  if they form an atomic pair *w.r.t.*  $\ell$  (*i.e.* the code in Fig. 6(a) does not loop), or there is a read  $r'$  after  $r$  in the program order forming an atomic pair *w.r.t.*  $\ell$  with  $w$ , such that  $r'$  is the last read issued by the loop before the `stwcx`. succeeds (*i.e.* the code in Fig. 6(a) loops). We do not consider the case where the loop never terminates:

$$\text{rmw}(r, w, \ell) \triangleq \text{atom}(r, w, \ell) \vee (\exists r', r \xrightarrow{\text{po}} r' \wedge \text{loc}(r) = \text{loc}(r') \wedge \text{atom}(r', w, \ell))$$

In Fig. 8(b), we open up the fno box protecting the read ( $a$ ) from  $x$  on  $P_0$ . We suppose that the fno is immediately successful, *i.e.* the code in Fig. 6(a) does not loop. Hence we expand the fno event ( $a$ ) on  $P_0$  to the  $r^*$  ( $a_1$ ) (from the `lwarx`) in program order with the  $w^*$  ( $a_2$ ) (from the successful `stwcx`).

We define a read to be *protected* when it is issued by the `lwarx` of a rmw immediately followed in program order by a non-cumulative barrier; an execution  $(E, X)$  is protected when its fragile reads are:

$$\begin{aligned} \text{protected}(r) &\triangleq \exists w, \text{rmw}(r, w, \text{loc}(r)) \wedge (\forall m, w \xrightarrow{\text{po}} m \Rightarrow w \xrightarrow{\text{ab}_1} m) \\ \text{protected}(E, X) &\triangleq \forall r, \text{fragile}(r) \Rightarrow \text{protected}(r) \end{aligned}$$

In Fig. 8(b), the write ( $e$ ) from which ( $a_1$ ) reads hits the memory before ( $a_2$ ), *i.e.*  $(e) \xrightarrow{\text{ws}} (a_2)$ . Hence there are two paths from ( $e$ ) to ( $b$ ):  $(e) \xrightarrow{\text{rf}} (a_1) \xrightarrow{\text{po}} (b)$  and  $(e) \xrightarrow{\text{ws}} (a_2) \xrightarrow{\text{po}} (b)$ . Thus we can trade the fragile pair  $(a_1), (b)$  for the pair  $(a_2), (b)$ : we compensate the lack of write atomicity of ( $e$ ) (*i.e.*  $(e) \xrightarrow{\text{rfe}} (a)$  not global) by using the write serialisation between ( $e$ ) and ( $a_2$ ) (thanks to the rmw) instead of cumulativity before. Formally, we prove that a sequence  $w \xrightarrow{\text{grf}_2^1} r \xrightarrow{\text{ppo}_2} m$  with  $r$  protected is in  $\xrightarrow{\text{ws}; \text{ghb}_1}$ , *i.e.* globally ordered on  $A_1$ :

**Lem. 6.**  $\forall wrm, (\text{protected}(r) \wedge w \xrightarrow{\text{grf}_2^1} r \xrightarrow{\text{ppo}_2} m) \Rightarrow w \xrightarrow{\text{ws}; \text{ghb}_1} m$

**Proof** Since  $r$  is protected, there are  $r'$  and  $w'$  such that  $\text{rmw}(r', w', \text{loc}(r))$  where  $r'$  is  $r$  or a subsequent read in  $\xrightarrow{\text{po}}$ . In both cases, we have  $w \xrightarrow{\text{ws}} w'$ . Moreover, since there is a barrier between  $w'$  and  $m$  (*i.e.*  $w' \xrightarrow{\text{ab}_1} m$ ), we know that  $w' \xrightarrow{\text{ghb}_1} m$ .  $\square$

Thus, if we protect the fragile reads, the only remaining fragile pairs are the ones in  $\text{ppo}_2^1$ . In Fig. 8(a), we have  $(e) \xrightarrow{\text{ws}} (a_2) \xrightarrow{\text{po}} (b) \xrightarrow{\text{fr}} (f)$  and  $(f) \xrightarrow{\text{ws}} (c_2) \xrightarrow{\text{po}} (d) \xrightarrow{\text{fr}} (e)$ , hence a cycle in  $\xrightarrow{\text{ws}} \cup \xrightarrow{\text{fr}} \cup \xrightarrow{\text{po}}$ . Since  $\xrightarrow{\text{ws}}$  and  $\xrightarrow{\text{fr}}$  are global, to invalidate this cycle, we need to order globally (*e.g.* by a barrier) the accesses ( $a_2$ ) and ( $b$ ) on  $P_0$  and ( $c_2$ ) and ( $d$ ) on  $P_1$ . Indeed, if an execution is protected, non-cumulative barriers placed between the remaining fragile pairs in  $\text{ppo}_2^1$  ensure stability:

**Lem. 7.**  $A_1.\text{valid}(E, X) \wedge \text{protected}(E, X) \wedge (\text{ppo}_2^1 \subseteq \xrightarrow{\text{ab}_1}) \Rightarrow A_2.\text{valid}(E, X)$

Arch.	Fragile pair	rmw (mapping A)	rmw (mapping P)
Power	$r \xrightarrow{po} r$	$\mathbf{fno} \xrightarrow{po} \mathbf{fno}$	$\mathbf{fno} \xrightarrow{\mathbf{sync}} r$
	$r \xrightarrow{po} w$	$\mathbf{fno} \xrightarrow{po} \mathbf{sta}$	$\mathbf{fno} \xrightarrow{\mathbf{lwsync}} w$
	$w \xrightarrow{po} w$	$\mathbf{sta} \xrightarrow{po} \mathbf{sta}$	$w \xrightarrow{\mathbf{lwsync}} w$
	$w \xrightarrow{po} r$	$\mathbf{sta} \xrightarrow{po} \mathbf{fno}$	$w \xrightarrow{\mathbf{sync}} r$
x86	$w \xrightarrow{po} r$	$\mathbf{xchg} \xrightarrow{po} r$	na

**Fig. 9.** Mappings A and P: rmw

**Proof** Barriers are by definition compatible with  $\xrightarrow{\text{ghb}_1}$ . The A-cumulativity is handled by the protection of the fragile reads as shown in Lem. 6. Finally, the barriers order globally the remaining fragile pairs. Hence the barriers induce a covering synchronisation relation by Lem. 4.  $\square$

This lemma leads to a mapping which we call P (for protected reads), given in Fig. 9. This mapping places a **fno** on the first read of a fragile pair, and a barrier between this **fno** and the second access of the fragile pair of a given piece of code. If the first access of the fragile pair is a write, it remains unchanged and we only place a barrier between the two accesses, following the mapping F. For the read-read (resp. read-write) case, since replacing a read by a **fno** amounts to replacing the read by a sequence of events ending with a write, we choose a barrier ordering write-read (resp. write-write) pairs, *i.e.* Power **sync** (resp. **lwsync**). Following Lem. 7, the mapping P gives a SC behaviour to a program.

H.-J. Boehm and S. Adve propose in [10] a mapping of all stores into rmw (*i.e.* **xchg**) on x86 (which has no fragile reads), to provide a SC semantics to C++ atomics. We call this mapping A-x86 (for atomics), and give it in Fig. 9. For models with fragile reads, *e.g.* Power, they question in [4] the existence of “*more efficient mappings (than the use of locks)*”. The mapping P could be more efficient, since it removes the need for cumulativity. Yet, mapping reads to rmw introduces additional stores (issued by **stwcx.**), which may impair the performance. Moreover, we have to use cumulative barriers in the mapping P, for Power does not provide non-cumulative barriers. Yet, we show in Sec. 5 that the mapping P is more efficient than locks on Power machines.

We propose another mapping, given in Fig. 9, which we call A-Power. All reads and writes are mapped into rmw (using **fno** for reads and fetch-and-store (**sta**) [1, p. 719] for writes). The documentation stipulates indeed that “*a processor has at most one reservation at any time*” [1, p. 663]. Hence two rmw on the same processor in program order may be preserved in this order, because the writes issued by their **stwcx.**, though to different locations, would be ordered by a dependency over the reservation. Although the documentation does not state whether this dependency even exists, we show (see Sec. 5) that the mapping A-Power restores SC experimentally and is more efficient than locks as well.

## 4 Stability from a weak architecture to SC

We now want to minimise the synchronisation that we use, *i.e.* we would like to synchronise only the conflicting accesses (either competing accesses or fragile pairs) that are actually necessary. For example, if in the **iriw** test of Fig. 1, we add a write ( $g$ ) to a fresh variable  $z$  after (in program order) the write ( $e$ ) to  $x$  on  $P_2$ , ( $e$ ) and ( $g$ ) may not be preserved in program order, *i.e.* ( $e$ ) and ( $g$ ) may form a fragile pair. Yet, there is no need to maintain them, since they do not contribute to the cycle we want to forbid.

*Critical cycles* D. Shasha and M. Snir provide in [23] an analysis to place barriers in a program, in order to enforce a SC behaviour. They examine in [23, Thm. 3.9 p. 297] the *critical cycles* of an execution, and show that placing a barrier along each program order arrow of such a cycle (each *delay* arrow) is enough to restore SC. Yet, this work does not provide any semantics of weak memory models. We show in Coq that their technique applies to the models embraced by our framework, *e.g.* models with store buffering, like TSO or relaxing store atomicity, like Power.

Given an event structure  $E$ , a cycle  $\sigma \subseteq (\overset{\text{cmp}}{\leftrightarrow} \cup \overset{\text{po}}{\rightarrow})^+$  (where  $\overset{\text{cmp}}{\leftrightarrow}$  is the competing relation of Sec. 2) is critical, written  $\text{critical}(E, \sigma)$ , when it satisfies the two following properties: **(i)** Per processor, there are at most two memory accesses  $(x, y)$  on this processor, such that  $x \overset{\text{po}}{\rightarrow} y$  and  $\text{loc}(x) \neq \text{loc}(y)$ . **(ii)** For a given memory location  $x$ , there are at most three accesses relative to  $x$ , and these accesses are from distinct processors ( $w \overset{\text{cmp}}{\leftrightarrow} w$ ,  $w \overset{\text{cmp}}{\leftrightarrow} r$ ,  $r \overset{\text{cmp}}{\leftrightarrow} w$  or  $r \overset{\text{cmp}}{\leftrightarrow} w \overset{\text{cmp}}{\leftrightarrow} r$ ). For example, the execution of **iriw** in Fig. 1 has a critical cycle.

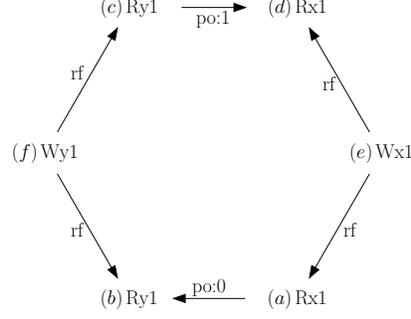
Note that a critical cycle does not forbid an execution on SC: the execution of **iriw** in Fig. 10(a) exhibits a critical cycle, *i.e.*  $(a) \overset{\text{po}}{\rightarrow} (b) \overset{\text{cmp}}{\leftrightarrow} (f) \overset{\text{cmp}}{\leftrightarrow} (c) \overset{\text{po}}{\rightarrow} (d) \overset{\text{cmp}}{\leftrightarrow} (e) \overset{\text{cmp}}{\leftrightarrow} (a)$ , but is authorised on SC (since there is no cycle in  $\overset{\text{com}}{\rightarrow} \cup \overset{\text{po}}{\rightarrow}$ ). Yet, a critical cycle does not require any information about the execution witness. A critical cycle becomes a *violation* of SC *w.r.t.*  $A$  (*w.r.t.* an execution witness  $X$ , written  $\text{viol}_{A,SC}(E, X, \sigma)$ ) when it is *oriented*, *i.e.* when it is included in  $\overset{\text{com}}{\rightarrow} \cup \overset{\text{po}}{\rightarrow}$  (but not in  $\overset{\text{ghb}^A}{\rightarrow}$ , otherwise it would violate  $A$ ). For example, the execution of **iriw** in Fig. 1 exhibits a violation of SC *w.r.t.* Power:  $(a) \overset{\text{po}}{\rightarrow} (b) \overset{\text{fre}}{\rightarrow} (f) \overset{\text{rfe}}{\rightarrow} (c) \overset{\text{po}}{\rightarrow} (d) \overset{\text{fre}}{\rightarrow} (e) \overset{\text{po}}{\rightarrow} (a)$ .

*Covering the critical cycles* We show in the following that, given a conflict relation  $\overset{c}{\rightarrow}$  and a synchronisation relation  $\overset{s}{\rightarrow}$ , it is enough to synchronise by  $\overset{s}{\rightarrow}$  the conflicting accesses of  $\overset{c}{\rightarrow}$  that belong to a critical cycle.

We call these conflicts *critical conflicts*. Given an event structure  $E$ , we write  $\overset{\text{cmin}}{\rightarrow}$  for the critical conflicts of  $E$ , *i.e.*  $\overset{c}{\rightarrow}$  restricted to the critical cycles of  $E$ :

**Definition 1 (Critical conflicts).**

$$m_1 \overset{\text{cmin}}{\rightarrow} m_2 \triangleq m_1 \overset{c}{\rightarrow} m_2 \wedge (\exists \sigma, \text{mv}(E, \sigma) \wedge m_1 \overset{\sigma}{\rightarrow} m_2)$$



**Fig. 10.** A SC execution of `iriw`

We show that if all the critical conflicts of an execution  $(E, X)$  valid on  $A$  are covered, there cannot be any violations in  $(E, X)$ . A violation is what makes an execution invalid on SC. Hence forbidding them ensures that  $(E, X)$  is valid on SC as well:

**Lem. 8 (Covering critical conflicts).**

$$\forall \vec{c} \xrightarrow{s}, \left( (\vec{c}, \vec{s}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow}) \vee (\vec{c}, \vec{s}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow}) \right) \Rightarrow \text{covering}(\overset{c}{\leftarrow}, \overset{s}{\rightarrow})$$

**Proof** Let  $(E, X)$  be an execution valid on  $A$  where  $\overset{c}{\leftarrow}$  is arbitrated by  $\overset{s}{\rightarrow}$ . Consider by contradiction a cycle in  $\text{SC.ghb}(E, X)$ . Hence, there exists a violation  $\vec{\sigma}$  in  $(E, X)$ .

We know that  $\vec{\sigma}$  is a cycle in  $(\overset{oc}{\leftarrow} \cup \overset{po}{\rightarrow})^+$ . Therefore, we know that  $\vec{\sigma}$  is equal to  $(\vec{\sigma} \cap \overset{oc}{\leftarrow}) \cup (\vec{\sigma} \cap \overset{po}{\rightarrow})$ . By hypothesis, since  $\vec{\sigma}$  is a violation, all the conflicts in  $\vec{\sigma}$  are arbitrated by  $\overset{s}{\rightarrow}$ .

- Suppose  $(\vec{c}, \vec{s}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow})$ . Since  $\overset{s}{\rightarrow}$  is compatible with  $\overset{com}{\rightarrow}$ , all the pairs in  $\overset{oc}{\leftarrow}$  are in  $\overset{s}{\rightarrow}$ . Hence  $\vec{\sigma} = (\vec{\sigma} \cap \overset{s}{\rightarrow}) \cup (\vec{\sigma} \cap \overset{po}{\rightarrow})$ . Thus, we have  $\vec{\sigma} \subseteq (\overset{s}{\rightarrow} \cup \overset{po}{\rightarrow})^+$ . Since  $\overset{s}{\rightarrow}$  is compatible with  $\overset{po}{\rightarrow}$ , there cannot be any such cycle  $\vec{\sigma}$ .
- Suppose  $(\vec{c}, \vec{s}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow}) = (\overset{c}{\leftarrow}, \overset{s}{\rightarrow})$ . Since  $\vec{\sigma}$  is a cycle in  $(\overset{oc}{\leftarrow} \cup \overset{po}{\rightarrow})^+$ , we know that  $\vec{\sigma}$  is included in  $(\overset{ws}{\rightarrow} \cup \overset{fr}{\rightarrow} \cup \overset{rf}{\rightarrow} \cup \overset{po}{\rightarrow})^+$ . Since  $\overset{ws}{\rightarrow}$  and  $\overset{fr}{\rightarrow}$  are in  $\overset{ghb_1}{\rightarrow}$ ,  $\vec{\sigma}$  is included in  $(\overset{ghb_1}{\rightarrow} \cup \overset{ppo_2 \setminus 1}{\rightarrow} \cup (\overset{grf_2 \setminus 1}{\rightarrow}, \overset{po}{\rightarrow}))^+$ , *i.e.* in  $(\overset{ghb_1}{\rightarrow} \cup \overset{c_{if}}{\rightarrow})^+$ . Since  $\overset{s}{\rightarrow}$  covers  $\overset{c_{if}}{\rightarrow}$  and is compatible with  $\overset{rf}{\rightarrow}$  and  $\overset{po}{\rightarrow}$ ,  $\vec{\sigma}$  is included in  $(\overset{ghb_1}{\rightarrow} \cup \overset{s_{if}}{\rightarrow})^+$ . Since  $\overset{s_{if}}{\rightarrow}$  is compatible with  $\overset{ghb_1}{\rightarrow}$ , there cannot be any such cycle  $\vec{\sigma}$ .  $\square$

Yet, even if a critical cycles is not necessarily an actual violation of SC, we can build an execution  $Y$  (by transforming  $X$ ) associated with the same event structure, which contains an actual violation of SC. In Fig. 10(a), if the  $\overset{rf}{\rightarrow}$

relations  $(f) \xrightarrow{\text{rf}} (b)$  and  $(e) \xrightarrow{\text{rf}} (d)$  become  $\xrightarrow{\text{fr}}$  ones (*i.e.*  $(b) \xrightarrow{\text{fr}} (f)$  and  $(d) \xrightarrow{\text{fr}} (e)$ ), then we build the execution of **iriw** depicted in Fig. 1, which is forbidden on SC. Hence we build an execution violating SC (the one in Fig. 1), from an execution which does not violate SC, but exhibits a critical cycle (the one in Fig. 10(a)). We formalise this idea in the following lemma:

**Lem. 9.**  $\forall E \xrightarrow{\sigma}, \text{critical}(E, \xrightarrow{\sigma}) \Rightarrow (\exists X, A. \text{valid}(E, X) \wedge \text{viol}_{A, \text{SC}}(E, X, \xrightarrow{\sigma}))$

**Proof** Let  $(E, X)$  be an execution valid on  $A$ . Let  $\xrightarrow{\sigma}$  be a critical cycle in  $E$ . By definition of critical cycle, we know that  $\text{acyclic}((\xrightarrow{\sigma} \cap \xrightarrow{\text{cdrf}}) \cup \xrightarrow{\text{ppo}^A})$ . Hence we know by Lem. 2 that we can build an execution  $Y$  associated with  $E$  and valid on  $A$  from any linear extension  $\xrightarrow{\text{le}}$  of the order  $((\xrightarrow{\sigma} \cap \xrightarrow{\text{cdrf}}) \cup \xrightarrow{\text{ppo}^A})^+$ .

Let us show that  $\xrightarrow{\sigma}$  is a violation of SC *w.r.t.*  $A$  in  $(E, Y)$ . Since  $\xrightarrow{\sigma}$  is a critical cycle in  $E$ , we know it is a critical cycle in  $(E, Y)$ . Hence, we need to show that  $\xrightarrow{\sigma}$  is included in  $(\xrightarrow{\text{cdrf}} \cup \xrightarrow{\text{ppo}})^+$  in  $(E, Y)$ . Thus, for all  $x$  and  $y$  such that  $x \xrightarrow{\sigma} y$ , we need to show that  $(x, y) \in (\xrightarrow{\text{r}})^+$ , where:

$$m_1 \xrightarrow{\text{r}} m_2 \triangleq ((m_1, m_2) \in \text{com}(Y) \wedge (\text{proc}(m_1) \neq \text{proc}(m_2))) \vee (m_1 \xrightarrow{\text{ppo}} m_2)$$

Consider two events  $x \xrightarrow{\sigma} y$ . Since  $\xrightarrow{\sigma}$  is a critical cycle, we have  $(x, y) \in (\xrightarrow{\text{cdrf}} \cup \xrightarrow{\text{ppo}})^+$ . Let us reason by induction over this statement. In the base case, let us do a case disjunction.

- When  $(x, y) \in \xrightarrow{\text{cdrf}}$ , we do case disjunction over the directions of  $x$  and  $y$ . Since they are in  $\xrightarrow{\text{cdrf}}$ , we know that they cannot be both reads.
  - If  $x$  and  $y$  are both writes, we know since they are in  $\xrightarrow{\text{cdrf}}$  that they are to the same location and from distinct processors. Moreover, we know that  $x \xrightarrow{\sigma} y$ . Therefore by definition of extracted write serialisation (see Lem. 2), they are in  $\text{ws}(\xrightarrow{\text{le}})$ , *i.e.* in  $\text{ws}(Y)$ .
  - If  $x$  is a read and  $y$  a write, we know since  $(E, Y)$  is valid on  $A$  that there exists a write  $w_x$  such that  $(w_x, x) \in \text{rf}(Y)$ . Hence, by definition of  $\text{rf}(Y)$ , we know that  $w_x \xrightarrow{\text{le}} x$ . Moreover, we know that  $(x, y) \in (\xrightarrow{\sigma} \cap \xrightarrow{\text{cdrf}})$  by hypothesis, hence  $x \xrightarrow{\text{le}} y$ . Thus by transitivity we have  $w_x \xrightarrow{\text{le}} y$ . Since  $w_x$  and  $y$  are both writes, and to the same location, we know by definition that  $(w_x, y) \in \text{ws}(Y)$ . Hence, since  $(w_x, x) \in \text{rf}(Y)$  and  $(w_x, y) \in \text{ws}(Y)$ , we have  $(x, y) \in \text{fr}(Y)$ .
  - If  $x$  is a write and  $y$  a read, we know since  $(E, Y)$  is valid on  $A$  that there exists a write  $w_y$  such that  $(w_y, y) \in \text{rf}(Y)$ . Suppose  $x = w_y$ ; in this case, we have  $(x, y) \in \text{rf}(Y)$ , hence the result. Suppose  $x \neq w_y$ . In this case, since  $x$  and  $w_y$  are both writes to the same location, we have  $(x, w_y) \in \text{ws}(Y) \vee (w_y, x) \in \text{ws}(Y)$ .
    - \* When  $(x, w_y) \in \text{ws}(Y)$ , suppose  $\text{proc}(x) = \text{proc}(w_y)$ . In this case, we know that  $x \xrightarrow{\text{ppo}} w_y$  (since in  $\xrightarrow{\text{ws}}$  and from the same processor). Otherwise (*i.e.* if  $\text{proc}(x) \neq \text{proc}(w_y)$ ),  $(x, w_y)$  is in  $\text{ws}(Y)$  and the two events are from distinct processors. The same reasoning applies

- to  $(w_y, y)$ : we have either  $w_y \xrightarrow{\text{po}} y$  if they are on the same processor, or  $(w_y, y) \in \text{rf}(Y)$  if not.
- \* When  $(w_y, x) \in \text{ws}(Y)$ , since  $(w_y, y) \in \text{rf}(Y)$ , we know that  $w_y$  is the maximal previous write to  $\text{loc}(y)$  before  $y$  in  $\xrightarrow{\text{le}}$ . Since  $(x, y) \in (\xrightarrow{\sigma} \cap \xrightarrow{\text{cdrf}})$  by hypothesis, we have  $x \xrightarrow{\text{le}} y$ . Since  $(w_y, x) \in \text{ws}(Y)$ , we have  $w_y \xrightarrow{\text{le}} x$ . Hence  $x$  occurs in between  $w_y$  and  $y$  in  $\xrightarrow{\text{le}}$ , a contradiction.
- When  $x \xrightarrow{\text{po}} y$ , we trivially have  $(x, y) \in ((\xrightarrow{\sigma} \cap \xrightarrow{\text{ocdrf}}) \cup \xrightarrow{\text{po}})^+$ .

The transitive case follows by immediate induction.  $\square$

In our framework, we show that the execution witnesses  $X$  of an event structure  $E$  are stable from  $A$  to SC if and only if there is no critical cycle in  $E$ , *i.e.* that an execution valid on  $A$  is SC if and only if there is no critical cycle in  $E$ :

**Thm. 1.**  $\forall E, (\forall X, \text{stable}_{A, \text{SC}}(E, X)) \Leftrightarrow \neg(\exists \xrightarrow{\sigma}, \text{critical}(E, \xrightarrow{\sigma}))$

**Proof** Let  $E$  be an event structure.

- $\Rightarrow$  Let  $X$  be an associated execution witness;  $(E, X)$  is stable by hypothesis. Suppose by contradiction the existence of a critical cycle  $\xrightarrow{\sigma}$  in  $E$ . In this case, by Lem. 9, we know that there exists another execution witness  $Y$ , such that  $(E, Y)$  is valid on  $A$ , in which  $\xrightarrow{\sigma}$  is a violation of SC *w.r.t.*  $A$ . Since all the executions associated with  $E$  are stable,  $(E, Y)$  is stable. Since  $(E, Y)$  is valid on  $A$  and stable, we know by definition of stable that  $(E, Y)$  is valid on SC. But since  $(E, Y)$  contains a violation of SC *w.r.t.*  $A$ ,  $(E, Y)$  cannot be valid on SC.
- $\Leftarrow$  Suppose  $(\xrightarrow{c}, \xrightarrow{s}) = (\xrightarrow{\text{cdrf}}, \xrightarrow{\text{sdrf}})$  or  $(\xrightarrow{c}, \xrightarrow{s}) = (\xrightarrow{\text{cif}}, \xrightarrow{\text{sif}})$ . Let  $(E, X)$  be an associated execution valid on  $A$ . We know that there is always such a  $X$  since any execution valid on SC is valid on a weaker architecture by Thm. 1: take for instance any linear extension of  $\xrightarrow{\text{po}}$ . Suppose there is no critical cycle in  $E$ . Thus,  $\xrightarrow{\text{cmin}}$  is empty, hence we know that  $\xrightarrow{\text{cmin}}$  is trivially covered by  $\xrightarrow{s}$  in  $(E, X)$ . Moreover,  $\xrightarrow{s}$  is covering for  $\xrightarrow{\text{cmin}}$  by Lem. 8. Therefore, we know by definition of covering that  $(E, X)$  is valid on SC.  $\square$

This theorem means that it is enough to arbitrate (with a covering relation) the conflicting accesses (either competing accesses or fragile pairs) occurring in the critical cycles of a given event structure to give it a SC behaviour. Hence we do not have to synchronise all the conflicts to ensure stability from a weak architecture to SC, but only those occurring in critical cycles.

## 5 Offence: a synchronisation tool

We implemented our study in our new offence tool, illustrating techniques that can be included in a compiler. Given a program in x86 or Power assembly, offence places either lock-based or lock-free synchronisation along the critical cycles of its input, following the mapping A, P, L or F, to enforce a SC behaviour.

## 5.1 Control flow graphs and critical cycles

Offence builds one control flow graph (cfg) per thread of the input program, containing *static events* (*i.e.* nodes representing memory accesses), and control flow instructions. A static memory event  $f$  has a direction, a location, originating instruction and processor, as events do, but no value component.

Given an event structure and two events  $e_1 \xrightarrow{\text{po}} e_2$ , mapping to static events  $f_1$  and  $f_2$ , we compute the *static program order* relation  $\xrightarrow{\text{pos}}$ , such that  $e_1 \xrightarrow{\text{po}} e_2$  entails  $f_1 \xrightarrow{\text{pos}} f_2$ , using a standard forward data flow analysis. If memory locations accessed by a given instruction are constant, we have  $\text{loc}(e_1) = \text{loc}(f_1)$  and  $\text{loc}(e_2) = \text{loc}(f_2)$ . Hence static conflicts computed from the cfg, written  $\xleftrightarrow{\text{cmps}}$ , abstract the conflicts of the event structures. When locations are not constant, we need alias analysis to compute an over-approximation of the locations of each static event. One can safely consider that all pairs of memory accesses by distinct processors conflict, as soon as one of the accesses is a write.

With  $F$  the set of static events, we call the triple  $(F, \xrightarrow{\text{pos}}, \xleftrightarrow{\text{cmps}})$  *static event structure*. Following Sec. 4, we enumerate the cycles of  $F$  that have properties **(i)** and **(ii)**, *i.e.* we build an over-approximation of the runtime critical cycles.

## 5.2 Placing synchronisation primitives

We then collect the fragile pairs (*i.e.* the write-read pairs in x86 and all pairs in Power) occurring in the critical cycles of  $F$ . By Thm. 1 it is necessary and sufficient to maintain these fragile pairs to reach stability, *i.e.* to restore SC.

*Barriers* Then, offence follows the mapping  $F$  on these fragile pairs. Given a pair  $(f_1, f_2)$ , offence issues the barrier request  $(i_1, i_2, b)$  where  $i_1 = \text{ins}(f_1)$ ,  $i_2 = \text{ins}(f_2)$  and  $b$  is the required barrier. Every path from  $i_1$  to  $i_2$  in the cfg should pass through a barrier instruction  $b$ . We use the global barrier placement of [18], which maximises the number of pairs maintained by a given barrier.

*Alternative to barriers* Offence can also follow the mappings A and P. For A-x86, the `xchg` instruction has an implicit write-read barrier semantics [10]. Thus, we use the global barrier placement of [18] for `xchg`. For locks, offence follows the mapping L on the conflict edges of the cfg. Sec. 3.1 describes the lock and unlock idioms that we use for Power. For x86, lock uses the `xchg` instruction to build a compare-and-swap loop, while unlock uses a single store instruction.

## 5.3 Experiments

*Generating tests* We generated two kinds of tests to exercise offence, using our previous diy tool [5], which computes tests in x86 or Power assembly from a given cycle of relations. First, we generate tests built from critical cycles, *e.g.* **iriw** in Fig. 1. Second, using a new tool, we mix such tests: given two tests built from critical cycles, we randomly permute processors of one of the given tests, alpha-convert its memory locations and registers to fresh ones, and interleave the codes of the programs. The process yields two series of tests, written X in the following, each series consisting of 209 tests for Power and 58 tests for x86.

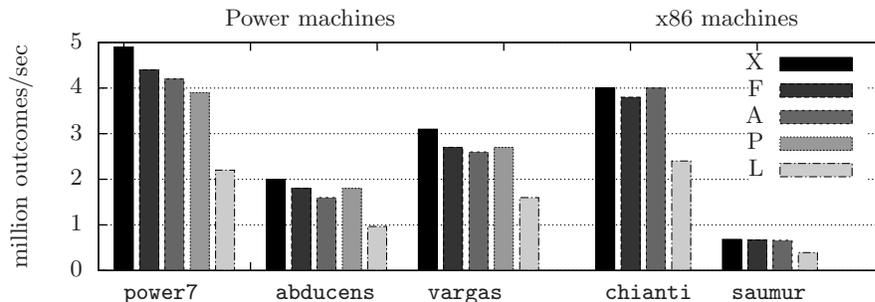


Fig. 11. Productivity observed during soundness experiments.

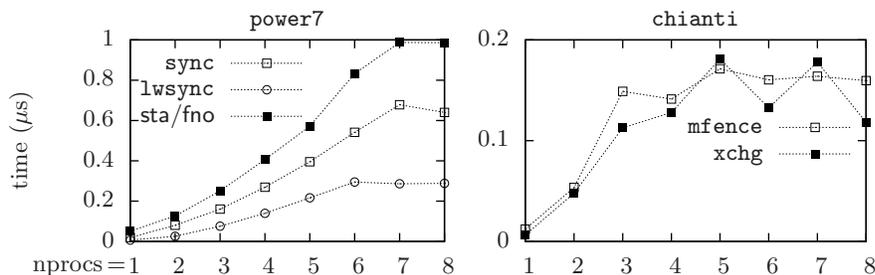
*Experimental soundness* We run these tests against hardware using our litmus tool [6]. We observed that all tests from the initial X series exhibit violations of SC and that the tests transformed by offence (following the mappings F, A, P and L) do *not* exhibit violations of SC, running each test at least  $10^9$  times. Thus we confirmed experimentally that our mappings enforce SC, which we established formally for the mappings F (Lem. 4), P (Lem. 7) and L (Lem. 3 and 5).

*Measures of synchronisation cost* We also estimate the cost of synchronisation constructs. We show in Fig. 11 the *productivity*, *i.e.* the number of outcomes collected per second, for the initial series of tests X, and for the tests transformed by offence following the mappings F, A, P and L. We ran tests on three Power machines: `power7` (Power7, 8 cores 4-ways SMT), `abducens` (Power6, 4 cores 2-ways SMT) and `vargas` (Power6, 32 cores 2-ways SMT); and on two AMD64 machines `chianti` (Intel Xeon, 8 cores, 2-ways HT) and `saumur` (Intel Xeon, 4 cores, 2-ways HT). Our mappings F, P and A outperform the L one, *i.e.* provide “*more efficient mappings (than the use of locks)*”, answering the question of [4].

To compare the barriers and `rmw` more precisely, we consider 8 specific tests from 1 to 8 threads, where we add synchronisation with offence so that there is only one synchronisation primitive per thread, and insert the code for each thread inside a tight loop. We then measure running times on our two 8 core machines, `power7` and `chianti`, subtract the time of the original test from the time of synchronised tests and divide the result by loop size. We give the results in Fig. 12. While fences and `rmw` are fast in isolation (10–20 ns on one thread), their cost raises to hundreds of ns when communication by shared memory occurs.

## 6 Related Work and Conclusion

*Related work* The DRF guarantee [3, 10], the semantics of synchronisation idioms [9, 8], and the insertion of barriers [23, 14, 11, 16] have been extensively studied,



**Fig. 12.** Time of synchronisation constructs, in microseconds.

but most of these works focus on one kind of synchronisation, and none of them addresses Power features such as *cumulativity* or the lack of write atomicity.

S. Burckhardt and M. Musuvathi examine in [12] whether we can simulate a program running on TSO by enumerating only its SC executions. They distinguish a class of such executions, the *TSO-safe* ones. We believe these executions to be an instance of our stable ones, *i.e.* the stable executions from TSO to SC. Yet, our characterisation of stability in the general case is a novel contribution.

J. Lee and D. Padua examine in [18] the issue of restoring SC at the compiler level. We have re-used their global fence placement algorithm. Yet, our work improves on [18] *w.r.t.* semantical foundations: as a result, we use Power `lwsync` when possible and we do not use the x86 barriers `lfence` and `sfence` that are irrelevant to user-level code. We believe that our mappings could be included in their Java compiler [25], *i.e.* using `lwsync` for Power, and `xchg` for x86.

*Conclusion* We propose a formal study of stability in weak memory models. This allows us to define several mappings of Power or x86 assembly code, which, as we prove in Coq, give a SC behaviour to a given program. Along the way, we give a semantics to Power’s `lwarx` and `stwcx`. instructions and show how to use the lightweight Power barrier `lwsync`, which are novel contributions. In addition, we characterise the executions stable from a weak architecture to SC, hence generalise the result of [23] to weak memory models. Finally, we implement our study in our `offence` tool, to measure the cost of these mappings: our lock-free mappings outperform locks. Our work could benefit to compiler writers, for we propose several mappings ensuring that an assembly program is SC: for example, the problem of ensuring a SC semantics to C++ atomics is of great interest. Moreover, our work could benefit to the verification community because it allows to separate the verification of a program running on a weak architecture into two steps: first, check the stability of a program from a weak architecture to SC; second, apply SC-sound verification techniques on this program.

## References

1. *Power ISA Version 2.06*. 2009.
2. S. V. Adve and K. Gharachorloo. Shared Memory Consistency Models: A Tutorial. *IEEE Computer*, 29:66–76, 1995.
3. S. V. Adve and M. D. Hill. Weak Ordering - A New Definition. In *ISCA 1990*.
4. S.V. Adve and H.-J. Boehm. Memory Models: A Case for Rethinking Parallel Languages and Hardware. To appear in CACM.
5. J. Alglave, L. Maranget, S. Sarkar, and P. Sewell. Fences in Weak Memory Models. In *CAV 2010*.
6. J. Alglave, L. Maranget, S. Sarkar, and P. Sewell. Litmus: Running Tests Against Hardware. In *TACAS 2011*.
7. Alpha Architecture Reference Manual, Fourth Edition, 2002.
8. H.-J. Boehm. Reordering Constraints for Pthread-Style Locks. In *PPoPP 2007*.
9. H.-J. Boehm. Threads Cannot Be Implemented As a Library. In *PLDI 2005*.
10. H.-J. Boehm and S. V. Adve. Foundations of the C++ Concurrency Memory Model. In *PLDI 2008*.
11. S. Burckhardt, R. Alur, and M. K. Martin. Checkfence: Checking consistency of concurrent data types on relaxed memory models. In *PLDI 2007*.
12. S. Burckhardt and M. Musuvathi. Effective Program Verification for Relaxed Memory Models. In *CAV 2008*.
13. J. Cantin, M. Lipasti, and J. Smith. The Complexity of Verifying Memory Coherence. In *SPAA 2003*.
14. X. Fang, J. Lee, and S. Midkiff. Automatic fence insertion for shared memory multiprocessing. In *ICS 2003*.
15. T. Huynh and A. Roychoudhury. A memory model sensitive checker for C#. In *FM 2006*.
16. M. Kuperstein, M. Vechev, and E. Yahav. Automatic inference of memory fences. In *FMCAD 2010*.
17. L. Lamport. How to Make a Correct Multiprocess Program Execute Correctly on a Multiprocessor. *IEEE Trans. Comput.*, 46(7):779–782, 1979.
18. J. Lee and D.A. Padua. Hiding relaxed memory consistency with a compiler. *IEEE Transactions on Computers*, 50:824–833, 2001.
19. J. Manson, W. Pugh, and S. V. Adve. The Java Memory Model. In *POPL 2005*.
20. S. Park and D. L. Dill. An executable specification, analyzer and verifier for RMO. In *SPAA 1995*.
21. M. Rinard. Analysis of Multithreaded Programs. In *SAS 2001*.
22. J. Sevcik. *Program Transformations in Weak Memory Models*. PhD thesis, University of Edinburgh, 2008.
23. D. Shasha and M. Snir. Efficient and Correct Execution of Parallel Programs that Share Memory. In *TOPLAS 1988*.
24. Sparc Architecture Manual Version 9, 1994.
25. Z. Sura, X. Fang, C.-L. Wong, S.P. Midkiff, J. Lee, and D.A. Padua. Compiler techniques for high performance SC Java programs. In *PPoPP'05*. ACM, 2005.
26. Y. Yang, G. Gopalakrishnan, and G. Lindstrom. Memory model sensitive data race analysis. In *ICFEM 2004*.