Repairing Sequential Consistency in C/C++11

Ori Lahav
MPI-SWS, Germany *
orilahav@mpi-sws.org

Viktor Vafeiadis
MPI-SWS, Germany *
viktor@mpi-sws.org

Jeehoon Kang
Seoul National University, Korea
jeehoon.kang@sf.snu.ac.kr

Chung-Kil Hur
Seoul National University, Korea
gil.hur@sf.snu.ac.kr

Derek Dreyer
MPI-SWS, Germany *
dreyer@mpi-sws.org

Abstract

The C/C++11 memory model defines the semantics of concurrent memory accesses in C/C++, and in particular supports racy "atomic" accesses at a range of different consistency levels, from very weak consistency ("relaxed") to strong, sequential consistency ("SC"). Unfortunately, as we observe in this paper, the semantics of SC atomic accesses in C/C++11, as well as in all proposed strengthenings of the semantics, is flawed, in that (contrary to previously published results) both suggested compilation schemes to the Power architecture are unsound. We propose a model, called RC11 (for Repaired C11), with a better semantics for SC accesses that restores the soundness of the compilation schemes to Power, maintains the DRF-SC guarantee, and provides stronger, more useful, guarantees to SC fences. In addition, we formally prove, for the first time, the correctness of the proposed stronger compilation schemes to Power that preserve load-to-store ordering and avoid "out-of-thin-air" reads.

CCS Concepts • **Theory of computation → Program semantics**; • **Software and its engineering → Concurrent programming languages**

Keywords Weak memory models; C++11; declarative semantics; sequential consistency

1. Introduction

The C/C++11 memory model (C11 for short) [8] defines the semantics of concurrent memory accesses in C/C++, of which there are two general types: *non-atomic* and *atomic*. Non-atomic accesses are intended for normal data: races on such accesses are considered as programming errors and lead to undefined behavior, thus ensuring that they can be compiled to plain machine loads and stores and that it is sound to apply standard sequential optimizations on non-atomic accesses. In contrast, atomic accesses are specifically intended for communication between threads: thus, races on atomics are permitted, but at the cost of introducing hardware fence instructions during compilation and imposing restrictions on how such accesses may be merged or reordered.

The degree to which an atomic access may be reordered with other operations—and more generally, the implementation cost of an atomic access—depends on its *consistency* level, concerning which C11 offers programmers several options according to their needs. Strongest and most expensive are *sequentially consistent* (SC) accesses, whose primary purpose is to restore the simple interleaving semantics of sequential consistency [20] if a program (when executed under SC semantics) only has races on SC accesses. This property is called "DRF-SC" and was a main design goal for C11. To ensure DRF-SC, the standard compilation schemes for modern architectures typically insert hardware "fence" instructions appropriately into the compiled code, with those for weaker architectures (like Power and ARMv7) introducing a full (strong) fence adjacent to each SC access.

Weaker than SC atomics are *release-acquire* accesses, which can be used to perform "message passing" between threads without incurring the implementation cost of a full SC access; and weaker and cheaper still are *relaxed* accesses, which are intended to be compiled down to plain loads and stores at the machine level and which provide only the minimal synchronization guaranteed by the hardware. Finally, the C11 model also supports language-level *fence* instructions, which provide finer-grained control over where hardware fences are to be placed and serve as a barrier to prevent unwanted compiler optimizations.

* Saarland Informatics Campus.

In this paper, we are mainly concerned with the semantics of SC atomics (*i.e.,* SC accesses and SC fences), and their interplay with the rest of the model. Since sequential consistency is such a classical, well-understood notion, one might expect that the semantics of SC atomics should be totally straightforward, but sadly, as we shall see, it is not!

The main problem arises in programs that mix SC and non-SC accesses to the same location. Although not common, such mixing is freely permitted by the C11 standard, and has legitimate uses—*e.g.,* as a way of enabling faster (non-SC) reads from an otherwise quite strongly synchronized data structure. Indeed, we know of several examples of code in the wild that mixes SC accesses together with release/acquire or relaxed accesses to the same location: `seqlocks` [9] and Rust's `crossbeam` library [2]. Now, consider the following program due to Manerkar *et al.* [22]:

$$x :=_{\mathtt{sc}} 1 \;\Big\|\; \begin{array}{l} a := x_{\mathtt{acq}} \,/\!/\, 1 \\ c := y_{\mathtt{sc}} \,/\!/\, 0 \end{array} \;\Big\|\; \begin{array}{l} b := y_{\mathtt{acq}} \,/\!/\, 1 \\ d := x_{\mathtt{sc}} \,/\!/\, 0 \end{array} \;\Big\|\; y :=_{\mathtt{sc}} 1$$

(IRIW-acq-sc)

Here and in all other programs in this paper, we write $a, b, \ldots$ for local variables (registers), and assume that all variables are initialized to 0. The program contains two variables, x and y, which are accessed via SC atomic accesses and also read by acquire atomic accesses. The annotated behavior (reading $a = b = 1$ and $c = d = 0$) corresponds to the two threads observing the writes to x and y as occurring in different orders, and is forbidden by C11. (We defer the explanation of how C11 forbids this behavior to §2.)

Let's now consider how this program is compiled to Power. Two compilation schemes have been proposed [7]. Both use Power's strongest fence instruction, called `sync`, for the compilation of SC atomics. The first scheme, the one implemented in the `GCC` and `LLVM` compilers, inserts a `sync` fence *before* each SC access ("leading sync" convention), whereas the alternative scheme inserts a `sync` fence *after* each SC access ("trailing sync" convention). The intent of both schemes is to have a strong barrier between every pair of SC accesses, enforcing, in particular, sequential consistency on programs containing only SC accesses. Nevertheless, by mixing SC and release-acquire accesses, one can quickly get into trouble, as illustrated by IRIW-acq-sc.

In particular, if one compiles the program into Power using the trailing sync convention, then the behavior is allowed by Power.[1] Since all SC accesses are at the end of the threads, the trailing `sync` fences have no effect, and the example reduces to (the result of compilation of) IRIW with only acquire reads, which is allowed by the Power memory model. In §2.1, we show further examples illustrating that the other, leading sync scheme also leads to behaviors in the target of compilation to Power that are not permitted in the source.

Although the C11 model is known to have multiple problems (*e.g.,* the "out-of-thin-air" problem [31, 11], the lack of monotonicity [30]), none of them until now affected the correctness of compilation to the mainstream architectures. In contrast, the IRIW-acq-sc program from [22] and our examples in §2.1 show that both the suggested compilation schemes to Power are unsound with respect to the C11 model, thereby contradicting the results of [7, 27]. The same problem occurs in some compilation schemes to ARMv7 (see §6), as well as for ARMv8 (see [3, Appendix A.2] for an example).

In the remainder of the paper, we propose a way to repair the semantics of SC accesses that resolves the problems mentioned above. In particular, our corrected semantics restores the soundness of the suggested compilation schemes to Power. Moreover, it still satisfies the standard DRF-SC theorem in the absence of relaxed accesses: if a program's sequentially consistent executions only ever exhibit races on SC atomic accesses, then its semantics under full C11 is also sequentially consistent. It is worth noting that this correction only affects the semantics of programs mixing SC and non-SC accesses to the same location: we show that, without such mixing, it coincides with the strengthened model of Batty *et al.* [5].

We also apply two additional, orthogonal, corrections to the C11 model, which strengthen the semantics of SC fences. The first fix corrects a problem already noted before [27, 21, 17], namely that the current semantics of SC fences does not recover sequential consistency, even when SC fences are placed between every two commands in programs with only release/acquire atomic accesses. The second fix provides stronger "cumulativity" guarantees for programs with SC fences. We justify these strengthenings by proving that the existing compilation schemes for x86-TSO, Power, and ARMv7 remain sound with the stronger semantics.

Finally, we apply another, mostly orthogonal, correction to the C11 model, in order to address the well-known "out-of-thin-air" problem. The problem is that the C11 standard permits certain executions as a result of causality cycles, which break even basic invariant-based reasoning [11] and invalidate DRF-SC in the presence of relaxed accesses. The correction, which is simple to state formally, is to strengthen the model to enforce load-to-store ordering for atomic accesses, thereby ruling out such causality cycles, at the expense of requiring a less efficient compilation scheme for relaxed accesses. The idea of this correction is not novel—it has been extensively discussed in the literature [31, 11, 30]—but the suggested compilation schemes to Power and ARMv7 have not yet been proven sound. Here, we give the first proof that one of these compilation schemes—the one that places a fake control dependency after every relaxed read—is sound. The proof is surprisingly delicate, and involves a novel argument similar to that in DRF-SC proofs.

Putting all these corrections together, we propose a new model called RC11 (for Repaired C11) that supports nearly all features of the C11 model (§3). We prove correctness of compilation to x86-TSO (§4), Power (§5), and ARMv7 (§6),

[1] Formally, we use the recent declarative model of Power by Alglave *et al.* [4].

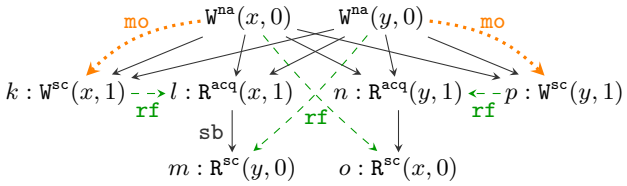

Figure 1. An execution of IRIW-acq-sc yielding the result $a = b = 1 \wedge c = d = 0$.

the soundness of a wide collection of program transformations (§7), and a DRF-SC theorem (§8).

2. The Semantics of SC Atomics in C11: What's Wrong, and How Can We Fix It?

The C11 memory model defines the semantics of a program as a set of consistent executions. Each execution is a graph. Its nodes, $\mathtt{E}$, are called *events* and represent the individual memory accesses and fences of the program, while its edges represent various relations among these events:

- The *sequenced-before* (sb) relation, a.k.a. *program order*, captures the order of events in the program's control flow.
- The *reads-from* (rf) relation associates each write with the set of reads that read from that write. In a consistent execution, the reads-from relation should be functional (and total) in the second argument: a read must read from exactly one write.
- Finally, the *modification order* (mo) is a union of total orders, one for each memory address, totally ordering the writes to that address. Intuitively, it records for each memory address the globally agreed-upon order in which writes to that address happened.

As an example, in Fig. 1, we depict an execution of the IRIW-acq-sc program discussed in the introduction. In addition to the events corresponding to the accesses appearing in the program, the execution contains two events for the implicit non-atomic initialization writes to x and y, which are assumed to be sb-before all other events.

Notation 1. Given a binary relation R, we write $R^?$, R^+, and R^* respectively to denote its reflexive, transitive, and reflexive-transitive closures. The inverse relation is denoted by R^{-1}. We denote by $R_1; R_2$ the left composition of two relations R_1, R_2, and assume that ; binds tighter than $\cup$ and $\setminus$. Finally, we denote by $[A]$ the identity relation on a set A. In particular, $[A]; R; [B] = R \cap (A \times B)$.

Based on these three basic relations, C11 defines some derived relations. First, whenever an acquire or SC read reads from a release or SC write, we say that the write *synchronizes with* (sw) the read.[2] Next, we say that one event *happens before* (hb) another event if they are connected by a sequence of sb or sw edges. Formally, $\mathtt{hb} \triangleq (\mathtt{sb} \cup \mathtt{sw})^+$. For example, in Fig. 1, event k synchronizes with l and therefore k happens-before l and m. Lastly, whenever a read event e reads from a write that is mo-before another write f to the same location, we say that e *reads-before* (rb) f (this relation is also called *"from-read"* [4], but we find *reads-before* more intuitive). Formally, $\mathtt{rb} \triangleq \mathtt{rf}^{-1}; \mathtt{mo} \setminus [\mathtt{E}]$. The "$\setminus [\mathtt{E}]$" part is needed so that RMW events ("read-modify-write", induced by atomic update operations like fetch-and-add and compare-and-swap) do not read-before themselves. For example, in Fig. 1, we have $\langle m, p \rangle \in \mathtt{rb}$ and $\langle o, k \rangle \in \mathtt{rb}$.

Consistent C11 executions require that hb is irreflexive (equivalently, $\mathtt{sb} \cup \mathtt{sw}$ is acyclic), and further guarantee *coherence* (aka *SC-per-location*) and *atomicity* of RMWs. Roughly speaking, coherence ensures that (i) the order of writes to the same location according to mo does not contradict hb (COHERENCE-WW); (ii) reads do not read values written in the future (NO-FUTURE-READ and COHERENCE-RW); (iii) reads do not read overwritten values (COHERENCE-WR); and (iv) two hb-related reads from the same location cannot read from two writes in reversed mo-order (COHERENCE-RR). We refer the reader to Prop. 1 in §3 for a formal definition of coherence.

Now, to give semantics to SC atomics, C11 stipulates that in consistent executions, there should be a strict total order, S, over all SC events, intuitively corresponding to the order in which these events are executed. This order is required to satisfy a number of conditions (but see Remark 1 below), where $\mathtt{E}^{\mathtt{sc}}$ denotes the set of all SC events in E:

(S1) S must include hb restricted to SC events (formally: $[\mathtt{E}^{\mathtt{sc}}]; \mathtt{hb}; [\mathtt{E}^{\mathtt{sc}}] \subseteq \mathtt{S}$);

(S2) S must include mo restricted to SC events (formally: $[\mathtt{E}^{\mathtt{sc}}]; \mathtt{mo}; [\mathtt{E}^{\mathtt{sc}}] \subseteq \mathtt{S}$);

(S3) S must include rb restricted to SC events (formally: $[\mathtt{E}^{\mathtt{sc}}]; \mathtt{rb}; [\mathtt{E}^{\mathtt{sc}}] \subseteq \mathtt{S}$);

(S4-7) S must obey a few more conditions having to do with SC fences.

Remark 1. The S3 condition above, due to Batty *et al.* [5], is slightly simpler and stronger than the one imposed by the official C11. Crucially, however, **all the problems and counterexamples we observe in this section, concerning the C11 semantics of SC atomics, hold for both Batty *et al.*'s model and the original C11**. The reason we use Batty *et al.*'s version here is that it provides a cleaner starting point for our discussion, and our solution to the problems with C11's SC semantics will build on it.

Intuitively, the effect of the above conditions is to enforce that, since S corresponds to the order in which SC events are executed, it should agree with the other global orders of events: hb, mo, and rb. However, as we will see shortly, condition S1 is too strong. Before we get there, let us first

[2] The actual definition of sw contains further cases, which are not relevant for the current discussion. These are included in our formal model in §3.

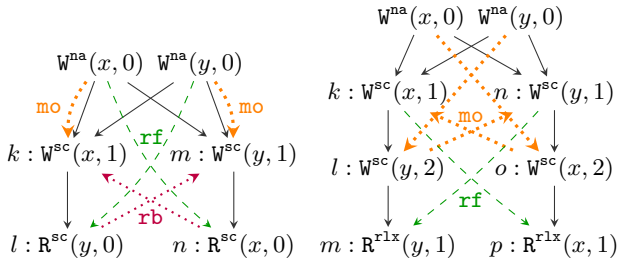

Figure 2. Inconsistent C11 executions of SB and 2+2W.

look at a few examples to illustrate how the conditions on S interact to enforce sequential consistency.

Consider the classic "store buffering" litmus test:

$$\begin{array}{l|\!|l} x :=_{\mathtt{sc}} 1 & y :=_{\mathtt{sc}} 1 \\ a := y_{\mathtt{sc}} \;/\!/ 0 & b := x_{\mathtt{sc}} \;/\!/ 0 \end{array} \qquad \text{(SB)}$$

Here, the annotated behavior is forbidden by C11. To see this, consider the first execution graph in Fig. 2. The rf edges are forced because of the values read, while the mo edges are forced because of COHERENCE-WW. Then, $\mathtt{S}(k,l)$ and $\mathtt{S}(m,n)$ hold because of condition S1; while $\mathtt{S}(l,m)$ and $\mathtt{S}(n,k)$ hold because of condition S3. This entails a cycle in S, which is disallowed.

Similarly, C11's conditions guarantee that the following (variant given in [32] of the) 2+2W litmus test disallows the annotated weak behavior:

$$\begin{array}{l|\!|l} x :=_{\mathtt{sc}} 1 & y :=_{\mathtt{sc}} 1 \\ y :=_{\mathtt{sc}} 2 & x :=_{\mathtt{sc}} 2 \\ a := y_{\mathtt{rlx}} \;/\!/ 1 & b := x_{\mathtt{rlx}} \;/\!/ 1 \end{array} \qquad \text{(2+2W)}$$

To see this, consider the second execution graph in Fig. 2, which has the outcome $a = b = 1$: the rf and mo edges are forced because of the values read and COHERENCE-WR. Now, $\mathtt{S}(k,l)$ and $\mathtt{S}(n,o)$ hold because of condition S1; while $\mathtt{S}(l,n)$ and $\mathtt{S}(o,k)$ hold because of condition S2. Again, this entails a cycle in S.

Let us now move to the IRIW-acq-sc program from the introduction, whose annotated behavior is also forbidden by C11. To see that, suppose without loss of generality that $\mathtt{S}(p,k)$ in Fig. 1. We also know that $\mathtt{S}(k,m)$ because of happens-before via l (S1). Thus, by transitivity, $\mathtt{S}(p,m)$. However, if the second thread reads $y = 0$, then m reads-before p, in which case $\mathtt{S}(m,p)$ (S3), and S has a cycle.

2.1 First Problem: Compilation to Power is Broken

The IRIW-acq-sc example demonstrates that the trailing sync compilation to Power is unsound for the C11 model. We will now see an example showing that the leading sync compilation is also unsound. Consider the following behavior, where all variables are zero-initialized and FAI(y) represents an atomic fetch-and-increment of y returning its value before the increment:

$$\begin{array}{l|\!|l|\!|l} x :=_{\mathtt{sc}} 1 & b := \mathrm{FAI}(y)_{\mathtt{sc}} \;/\!/ 1 & y :=_{\mathtt{sc}} 3 \\ y :=_{\mathtt{rel}} 1 & c := y_{\mathtt{rlx}} \;/\!/ 3 & a := x_{\mathtt{sc}} \;/\!/ 0 \end{array} \qquad \text{(Z6.U)}$$

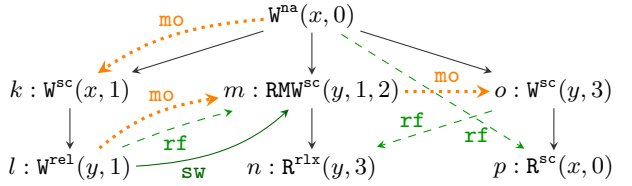

Figure 3. A C11 execution of Z6.U. The initialization of y is omitted as it is not relevant.

We will show that the behavior is disallowed according to C11, but allowed by its compilation to Power.

Fig. 3 depicts the only execution yielding the behavior in question that satisfies the coherence constraints. Again, the rf and mo edges are forced: even if all accesses in the program were relaxed atomic, they would have to go this way. $\mathtt{S}(k,m)$ holds because of condition S1 (k happens-before l, which happens-before m); $\mathtt{S}(m,o)$ holds because of condition S2 (m precedes o in modification order); $\mathtt{S}(o,p)$ holds because of condition S1 (o happens-before p). Finally, since p reads $x = 0$, we have that p reads-before k, so by S3, $\mathtt{S}(p,k)$, thus forming a cycle in S.

Under the leading sync compilation to Power, however, the behavior is allowed. Intuitively, all but one of the sync fences because of the SC accesses are useless because they are at the beginning of a thread. In the absence of other sync fences, the only remaining sync fence, due to the $a := x_{\mathtt{sc}}$ load in the last thread, is equivalent to an lwsync fence (cf. [17, §7]).

In [3, Appendix A.1] we provide a similar example using SC fences instead of RMW instructions, which shows that even placing sync fences both before and after SC accesses is unsound.

What Went Wrong and How to Fix it Generally, in order to provide coherence, hardware memory models provide rather strong ordering guarantees on accesses to the same memory location. Consequently, for conditions S2 and S3, which only enforce orderings between accesses to the same location, ensuring that compilation preserves these conditions is not difficult, even for weaker architectures like Power and ARM.

When, however, it comes to ensuring a strong ordering between accesses of *different* memory locations, as S1 does, compiling to weaker hardware requires the insertion of appropriate memory fence instructions. In particular, for Power, to enforce a strong ordering between two hb-related accesses to different locations, there should be a Power sync fence occurring somewhere in the hb-path (the sequence of sb and sw edges) connecting the two accesses. Unfortunately, in the presence of mixed SC and non-SC accesses, the Power compilation schemes do not always ensure that a sync exists between hb-related SC accesses. Specifically, if we follow the trailing sync convention, the hb-path (in Fig. 1) from k to m starting with an sw edge avoids the sync fence placed after k. Conversely, if we follow the leading sync convention, the hb-

path (in Fig. 3) from k to m ending with an $\mathtt{sw}$ edge avoids the fence placed before m. The result is that S1 enforces more ordering than the hardware provides!

So, if requiring that $\mathtt{hb}$ (on SC events) be included in $\mathtt{S}$ is too strong a condition, what should we require instead? The essential insight is that, according to either compilation scheme, we know that a `sync` fence will necessarily exist between SC accesses a and b if the $\mathtt{hb}$ path from a to b starts and ends with an $\mathtt{sb}$ edge. Second, if a and b access the same location, then the hardware will preserve the ordering anyway. These two observations lead us to replace condition S1 with the following:

(S1fix) $\mathtt{S}$ must relate any two SC events that are related by $\mathtt{hb}$, *provided that* the $\mathtt{hb}$-path between the two events either starts and ends with $\mathtt{sb}$ edges, or starts and ends with accesses to the same location (formally: $[\mathtt{E}^{\mathtt{sc}}];(\mathtt{sb}\cup\mathtt{sb};\mathtt{hb};\mathtt{sb}\cup\mathtt{hb}|_{\mathtt{loc}});[\mathtt{E}^{\mathtt{sc}}]\subseteq\mathtt{S}$, where $\mathtt{hb}|_{\mathtt{loc}}$ denotes $\mathtt{hb}$ edges between accesses to the same location).

We note that condition S1fix, although weaker than S1, suffices to rule out the weak behaviors of the basic litmus tests (*i.e.,* SB and 2+2W). In fact, just to rule out these behaviors, it suffices to require $\mathtt{sb}$ (on SC events) to be included in $\mathtt{S}$.

In essence, according to S1fix, $\mathtt{S}$ must include all the $\mathtt{hb}$-paths between SC accesses to different locations that exist regardless of any synchronization induced by the SC accesses at their endpoints. If a program does not mix SC and non-SC accesses to the same location, then every *minimal* $\mathtt{hb}$-path between two SC accesses to the same location (*i.e.,* one which does not go through another SC access) must start and end with an $\mathtt{sb}$ edge, in which case S1 and S1fix coincide.

Fixing the Model Before formalizing our fix, let us first rephrase conditions S1–S3 in the more concise style suggested by Batty *et al.* [5]. Instead of expressing them as separate conditions on a total order $\mathtt{S}$, they require a single *acylicity* condition, namely that $[\mathtt{E}^{\mathtt{sc}}];(\mathtt{hb}\cup\mathtt{mo}\cup\mathtt{rb});[\mathtt{E}^{\mathtt{sc}}]$ be acyclic. (In general, acyclicity of $\bigcup R_i$ is equivalent to the existence of a total order that contains $R_1, R_2, ...$)

We propose to correct the condition by replacing $\mathtt{hb}$ with $\mathtt{sb}\cup\mathtt{sb};\mathtt{hb};\mathtt{sb}\cup\mathtt{hb}|_{\mathtt{loc}}$. Accordingly, we require that

$$[\mathtt{E}^{\mathtt{sc}}];(\mathtt{sb}\cup\mathtt{sb};\mathtt{hb};\mathtt{sb}\cup\mathtt{hb}|_{\mathtt{loc}}\cup\mathtt{mo}\cup\mathtt{rb});[\mathtt{E}^{\mathtt{sc}}]$$

is acyclic. Note that this condition still ensures SC semantics for programs that have only SC accesses. Indeed, since $[\mathtt{E}^{\mathtt{sc}}];\mathtt{rf};[\mathtt{E}^{\mathtt{sc}}]\subseteq[\mathtt{E}^{\mathtt{sc}}];\mathtt{sw};[\mathtt{E}^{\mathtt{sc}}]\subseteq[\mathtt{E}^{\mathtt{sc}}];\mathtt{hb}|_{\mathtt{loc}};[\mathtt{E}^{\mathtt{sc}}]$, our condition implies acyclicity of $[\mathtt{E}^{\mathtt{sc}}];(\mathtt{sb}\cup\mathtt{rf}\cup\mathtt{mo}\cup\mathtt{rb});[\mathtt{E}^{\mathtt{sc}}]$. The latter suffices for this purpose, as it corresponds exactly to the declarative definition of sequential consistency [28].

2.1.1 Enabling Elimination of SC Accesses

We observe that our condition disallows the elimination of an SC write immediately followed by another SC write to the same location, as well as of an SC read immediately preceded by an SC read from the same location. While neither `GCC` nor `LLVM` performs these eliminations, they are sound under sequential consistency, as well as under C11 (with the fixes of [30]), and one may wish to preserve their soundness.

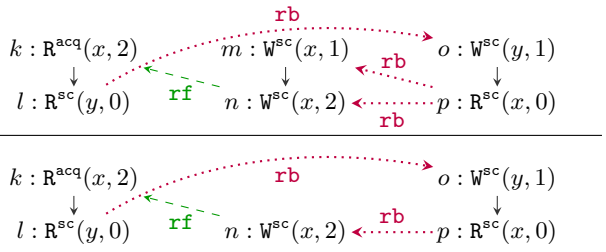

Figure 4. An abbreviated execution of WWmerge (source), and of the resulting program after eliminating the overwritten write m (target). The source execution has a disallowed cycle (m, l, o, p, m), while the target execution does not.

To see the unsoundness of eliminating an overwritten SC write, consider the following program. The annotated behavior is forbidden, but it will become allowed after eliminating $x :=_{\mathtt{sc}} 1$ (see Fig. 4).

$$\begin{array}{l} a := x_{\mathtt{acq}}\ /\!/2 \\ b := y_{\mathtt{sc}}\ /\!/0 \end{array} \;\Big\|\; \begin{array}{l} x :=_{\mathtt{sc}} 1 \\ x :=_{\mathtt{sc}} 2 \end{array} \;\Big\|\; \begin{array}{l} y :=_{\mathtt{sc}} 1 \\ c := x_{\mathtt{sc}}\ /\!/0 \end{array} \qquad \text{(WWmerge)}$$

Similarly, eliminating a repeated SC read is unsound (see example in [3, Appendix A.3]). The problem here is that these transformations remove an $\mathtt{sb}$ edge, and thus remove an $\mathtt{sb};\mathtt{hb};\mathtt{sb}$ path between two SC accesses.

Note that the removed $\mathtt{sb}$ edges are all edges between same-location accesses. Thus, supporting these transformations can be achieved by a slight weakening of our condition: we replace $\mathtt{sb};\mathtt{hb};\mathtt{sb}$ with $\mathtt{sb}|_{\neq\mathtt{loc}};\mathtt{hb};\mathtt{sb}|_{\neq\mathtt{loc}}$, where $\mathtt{sb}|_{\neq\mathtt{loc}}$ denotes $\mathtt{sb}$ edges that are not between accesses to the same location. Thus, we require acyclicity of $[\mathtt{E}^{\mathtt{sc}}];\mathtt{scb};[\mathtt{E}^{\mathtt{sc}}]$, where $\mathtt{scb}$ (*SC-before*) is given by:

$$\mathtt{scb} \triangleq \mathtt{sb}\cup\mathtt{sb}|_{\neq\mathtt{loc}};\mathtt{hb};\mathtt{sb}|_{\neq\mathtt{loc}}\cup\mathtt{hb}|_{\mathtt{loc}}\cup\mathtt{mo}\cup\mathtt{rb}.$$

We note that this change does not affect programs that do not mix SC and non-SC accesses to the same location.

2.2 Second Problem: SC Fences are Too Weak

In this section we extend our model to cover SC fences, which were not considered so far. Denote by $\mathtt{F}^{\mathtt{sc}}$ the set of SC fences in $\mathtt{E}$. The straightforward adaptation of the condition of Batty *et al.* [5] for the full model (obtained by replacing $\mathtt{hb}\cup\mathtt{mo}\cup\mathtt{rb}$ with our $\mathtt{scb}$) is that

$$\mathtt{psc}_1 \triangleq \left([\mathtt{E}^{\mathtt{sc}}]\cup[\mathtt{F}^{\mathtt{sc}}];\mathtt{sb}^{?}\right);\mathtt{scb};\left([\mathtt{E}^{\mathtt{sc}}]\cup\mathtt{sb}^{?};[\mathtt{F}^{\mathtt{sc}}]\right)$$

is acyclic. This condition generalizes the earlier condition by forbidding $\mathtt{scb}$ cycles even between *non-SC* accesses provided they are preceded/followed by an SC fence. This condition rules out weak behaviors of examples such as SB and 2+2W where all accesses are relaxed and SC fences are placed between them in the two threads.

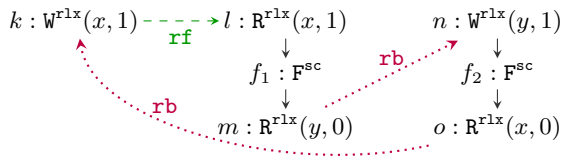

Figure 5. An execution of RWC+syncs yielding the annotated result. The rb edges are due to the reading from the omitted initialization events and the mo edges from those.

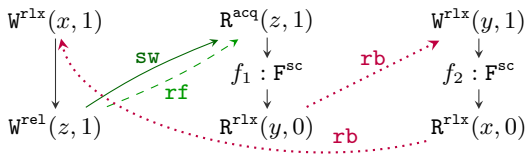

Figure 6. An abbreviated execution of W+RWC.

In general, one might expect that inserting an SC fence between every two instructions restores sequential consistency. This holds for hardware memory models, such as x86-TSO, Power, and ARM, for programs with aligned word-sized accesses (for their analogue of SC fences), but holds neither in the original C11 model nor in its strengthening [5] for two reasons. The first reason is that C11 declares that programs with racy non-atomic accesses have undefined behavior, and even if fences are placed everywhere such races may exist. There is, however, another way in which putting fences everywhere in C11 does not restore sequential consistency, even if all the accesses are atomic. Consider the following program:

$$x :=_{\mathtt{rlx}} 1 \;\Big\|\; \begin{array}{l} a := x_{\mathtt{rlx}} \mathbin{/\!/} 1 \\ \mathtt{fence}_{\mathtt{sc}} \\ b := y_{\mathtt{rlx}} \mathbin{/\!/} 0 \end{array} \;\Big\|\; \begin{array}{l} y :=_{\mathtt{rlx}} 1 \\ \mathtt{fence}_{\mathtt{sc}} \\ c := x_{\mathtt{rlx}} \mathbin{/\!/} 0 \end{array} \qquad \text{(RWC+syncs)}$$

The annotated behavior is allowed according to the model of Batty *et al.* [5] (and so, also by our weaker condition above). Fig. 5 depicts a consistent execution yielding this behavior, as the only $\mathtt{psc}_1$ edge is from f_1 to f_2. Yet, this behavior is disallowed by all implementations of C11. We believe that this is a serious omission of the standard rendering the SC fences too weak, as they cannot be used to enforce sequential consistency. This weakness has also been observed in a C11 implementation of the Chase-Lev deque by Lê *et al.* [21], who report that the weak semantics of SC fences in C11 requires them to unnecessarily strengthen the access modes of certain relaxed writes to SC. (In the context of the RWC+syncs, it would amount to making the write to x in the first thread into an SC write.)

Remark 2 (Itanium)**.** This particular weakness of the standard is attributed to Itanium, whose fences do not guarantee sequential consistency when inserted everywhere. While this would be a problem if C11 relaxed accesses were compiled to plain Itanium accesses, they actually have to be compiled to release/acquire Itanium accesses to guarantee read-read coherence. In this case, Itanium fences guarantee ordering. In fact, Itanium implementations provide multi-copy atomicity for release stores, and thus cannot yield the weak outcome of IRIW even without fences [14, §3.3.7.1].

Fixing the Semantics of SC Fences Analyzing the execution of RWC+syncs, we note that there is a $\mathtt{sb};\mathtt{rb};\mathtt{rf};\mathtt{sb}$ path from f_2 to f_1, but this path does not contribute to $\mathtt{psc}_1$. Although both rb and rf edges contribute to $\mathtt{psc}_1$, their composition $\mathtt{rb};\mathtt{rf}$ does not.

To repair the model, we define the *extended coherence order*, $\mathtt{eco} \triangleq (\mathtt{rf} \cup \mathtt{mo} \cup \mathtt{rb})^+$. This order includes the reads-from relation, rf, the modification order, mo, the reads-before relation, rb, and also all the compositions of these relations with one another—namely, all orders forced because of the coherence axioms. Then, we require that $\mathtt{psc}_1 \cup [\mathtt{F}^{\mathtt{sc}}];\mathtt{sb};\mathtt{eco};\mathtt{sb};[\mathtt{F}^{\mathtt{sc}}]$ is acyclic.

This stronger condition rules out the weak behavior of RWC+syncs because there are $\mathtt{sb};\mathtt{eco};\mathtt{sb}$ paths from one fence to another and vice versa (in one direction via the x accesses and in the other direction via the y accesses). Intuitively speaking, compilation remains correct with this stronger model since eco exists only between accesses to the same location, on which the hardware provides strong ordering guarantees.

Now it is easy to see that, given a program without non-atomic accesses, placing an SC fence between every two accesses guarantees SC. Indeed, by the definition of SC, it suffices to show that $\mathtt{eco} \cup \mathtt{sb}$ is acyclic. Consider a $\mathtt{eco} \cup \mathtt{sb}$ cycle. Since eco and sb are irreflexive and transitive, the cycle necessarily has the form $(\mathtt{eco};\mathtt{sb})^+$. Thus, between every two eco steps, there must be an SC fence. So in effect, we have a cycle in $\mathtt{eco};\mathtt{sb};[\mathtt{F}^{\mathtt{sc}}];\mathtt{sb}$, which can be regrouped to a cycle in $[\mathtt{F}^{\mathtt{sc}}];\mathtt{sb};\mathtt{eco};\mathtt{sb};[\mathtt{F}^{\mathtt{sc}}]$, which is forbidden by our model.

Finally, one might further consider strengthening the model by including eco in scb (which is used to define $\mathtt{psc}_1$), thereby ruling out the weak behavior of a variant of RWC+syncs using SC accesses instead of SC fences in threads 2 and 3. We note, however, that this strengthening is unsound for the default compilation scheme to x86-TSO (see Remark 4 in §4).

2.2.1 Restoring Fence Cumulativity

Consider the following variant of the store buffering program, where the write of $x := 1$ has been moved to another thread with a release-acquire synchronization.

$$\begin{array}{l} x :=_{\mathtt{rlx}} 1 \\ z :=_{\mathtt{rel}} 1 \end{array} \;\Big\|\; \begin{array}{l} a := z_{\mathtt{acq}} \mathbin{/\!/} 1 \\ \mathtt{fence}_{\mathtt{sc}} \\ b := y_{\mathtt{rlx}} \mathbin{/\!/} 0 \end{array} \;\Big\|\; \begin{array}{l} y :=_{\mathtt{rlx}} 1 \\ \mathtt{fence}_{\mathtt{sc}} \\ c := x_{\mathtt{rlx}} \mathbin{/\!/} 0 \end{array} \qquad \text{(W+RWC)}$$

The annotated behavior corresponds to the writes of x and y being observed in different orders by the reads, although SC fences have been used in the observer threads. This behavior is disallowed on x86, Power, and ARM because their fences

are cumulative: the fences order not only the writes performed by the thread with the fence instruction, but also the writes of other threads that are observed by the thread in question [23].

In contrast, the behavior is allowed by the model described thus far. Consider the execution shown in Fig. 6. While there is a $\mathtt{sb};\mathtt{rb};\mathtt{sb}$ path from f_1 to f_2, the only path from f_2 back to f_1 is $\mathtt{sb};\mathtt{rb};\mathtt{sb};\mathtt{sw};\mathtt{sb}$ (or, more generally, $\mathtt{hb};\mathtt{rb};\mathtt{hb}$), and so the execution is allowed.

To disallow such behaviors, we can replace $[\mathtt{F^{sc}}];\mathtt{sb}$ and $\mathtt{sb};[\mathtt{F^{sc}}]$ in the definitions above by $[\mathtt{F^{sc}}];\mathtt{hb}$ and $\mathtt{hb};[\mathtt{F^{sc}}]$.[3] This leads us to our final condition that requires that $\mathtt{psc}_{\text{base}} \cup \mathtt{psc}_{\mathtt{F}}$ is acyclic, where:

$$\mathtt{psc}_{\text{base}} \triangleq \left([\mathtt{E^{sc}}] \cup [\mathtt{F^{sc}}];\mathtt{hb}^{?}\right);\mathtt{scb};\left([\mathtt{E^{sc}}] \cup \mathtt{hb}^{?};[\mathtt{F^{sc}}]\right)$$
$$\mathtt{psc}_{\mathtt{F}} \triangleq [\mathtt{F^{sc}}];(\mathtt{hb} \cup \mathtt{hb};\mathtt{eco};\mathtt{hb});[\mathtt{F^{sc}}]$$

We note that $[\mathtt{F^{sc}}];\mathtt{psc}_{\text{base}};[\mathtt{F^{sc}}] \subseteq \mathtt{psc}_{\mathtt{F}}$. Hence, in programs without SC accesses (but with SC fences) it suffices to require that $\mathtt{psc}_{\mathtt{F}}$ is acyclic.

2.3 A Final Problem: Out-of-Thin-Air Reads

The C11 memory model suffers from a major problem, known as the "out-of-thin-air problem" [31, 11]. Designed to allow efficient compilation and many optimization opportunities for relaxed accesses, the model happened to be too weak, admitting "out-of-thin-air" behaviors, which no implementation exhibits. The standard example is load buffering with some form of dependencies in both threads:

$$\begin{array}{l} a := x_{\mathtt{rlx}} \mathbin{/\!\!/} 1 \\ \textbf{if}\,(a)\ y :=_{\mathtt{rlx}} a \end{array} \;\Big\|\; \begin{array}{l} b := y_{\mathtt{rlx}} \mathbin{/\!\!/} 1 \\ \textbf{if}\,(b)\ x :=_{\mathtt{rlx}} b \end{array} \qquad \text{(LB+deps)}$$

In this program, the formalized C11 model by Batty *et al.* [8] allows reading $a = b = 1$ even though the value 1 does not appear in the program. The reason is that the execution where both threads read and write the value 1 is consistent: each read reads from the write of the other thread. As one might expect, such behaviors are very problematic because they invalidate almost all forms of formal reasoning about programs. In particular, the example above demonstrates a violation of DRF-SC, the most basic guarantee that users of C11 were intended to assume: LB+deps has no races under sequential consistency, and yet has some non-SC behavior.

Fixing the model in a way that forbids all "out-of-thin-air" behaviors and still allows the most efficient compilation is beyond the scope of the current paper (see [16] for a possible solution). In this paper, we settle for a simpler solution of requiring $\mathtt{sb} \cup \mathtt{rf}$ to be acyclic. This is a relatively straightforward way to avoid the problem, although it carries some performance cost. Clearly, it rules out the weak behavior of LB+deps, but also of the following load-buffering program, which is nevertheless permitted by the Power and ARM architectures.

$$\begin{array}{l} a := x_{\mathtt{rlx}} \mathbin{/\!\!/} 1 \\ y :=_{\mathtt{rlx}} 1 \end{array} \;\Big\|\; \begin{array}{l} b := y_{\mathtt{rlx}} \mathbin{/\!\!/} 1 \\ x :=_{\mathtt{rlx}} 1 \end{array} \qquad \text{(LB)}$$

To correctly compile the stronger model to Power and ARM, one has to either introduce a fence between a relaxed atomic read and a subsequent relaxed atomic write or a forced dependency between every such pair of accesses [11]. The latter can be achieved by inserting a dummy control-dependent branch after every relaxed atomic read.

While the idea of strengthening C11 to require acyclicity of $\mathtt{sb} \cup \mathtt{rf}$ is well known [31, 11], we are not aware of any proof showing that the proposed compilation schemes of Boehm and Demsky [11] are correct, nor that DRF-SC holds under this assumption. The latter is essential for assessing our corrected model, as it is a key piece of evidence showing that our semantics for SC accesses is not overly weak.

Importantly, even in this stronger model, non-atomic accesses are compiled to plain machine loads and stores. This is what makes the compilation correctness proof highly non-trivial, as the hardware models allow certain $\mathtt{sb} \cup \mathtt{rf}$ cycles involving plain loads and stores. As a result, one has to rely on the *"catch-fire"* semantics (races on non-atomic accesses result in undefined behavior) for explaining behaviors that involve such cycles. A similar argument is needed for proving the correctness of non-atomic read-write reordering.

3. The Proposed Memory Model

In this section, we formally define our proposed corrected version of the C11 model, which we call RC11. Similar to C11, the RC11 model is given in a "declarative" style in three steps: we associate a set of graphs (called *executions*) to every program (§3.1), filter this set by imposing a consistency predicate (§3.2), and finally define the outcomes of a program based on the set of its consistent executions (§3.3). At the end of the section, we compare our model with C11 (§3.4).

Before we start, we introduce some further notation. Given a binary relation R, $\mathit{dom}(R)$ and $\mathit{codom}(R)$ denote its domain and codomain. Given a function f, $=_f$ denotes the set of f-equivalent pairs ($=_f \triangleq \{\langle a,b\rangle \mid f(a) = f(b)\}$), and $R|_f$ denotes the restriction of R to f-equivalent pairs ($R|_f \triangleq R \cap =_f$). When R is a strict partial order, $R|_{\text{imm}}$ denotes the set of all *immediate R edges*, *i.e.,* pairs $\langle a,b\rangle \in R$ such that for every c, $\langle c,b\rangle \in R$ implies $\langle c,a\rangle \in R^{?}$, and $\langle a,c\rangle \in R$ implies $\langle b,c\rangle \in R^{?}$.

We assume finite sets Loc and Val of locations and values. We use x, y, z as metavariables for locations and v for values. The model supports several modes for accesses and fences, partially ordered by $\sqsubset$ as follows:

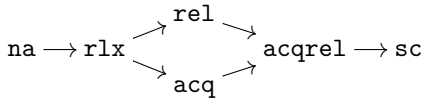

[3] To rule out only the cycle shown in Fig. 6, it would suffice to have replaced only the sb *to* a fence by an hb. We can, however, also construct examples, where it is useful for the sb *from* a fence to be replaced by hb.

3.1 From Programs to Executions

First, the program is translated into a set of executions. An *execution* G consists of:

1. a finite set of events $\mathtt{E} \subseteq \mathbb{N}$ containing a distinguished set $\mathtt{E}_0 = \{a_0^x \mid x \in \mathtt{Loc}\}$ of initialization events. We use $a, b, \ldots$ as metavariables for events.
2. a function $\mathtt{lab}$ assigning a *label* to every event in $\mathtt{E}$. Labels are of one of the following forms:
   - $\mathtt{R}^o(x, v)$ where $o \in \{\mathtt{na}, \mathtt{rlx}, \mathtt{acq}, \mathtt{sc}\}$.
   - $\mathtt{W}^o(x, v)$ where $o \in \{\mathtt{na}, \mathtt{rlx}, \mathtt{rel}, \mathtt{sc}\}$.
   - $\mathtt{F}^o$ where $o \in \{\mathtt{acq}, \mathtt{rel}, \mathtt{acqrel}, \mathtt{sc}\}$.

   We assume that $\mathtt{lab}(a_0^x) = \mathtt{W}^{\mathtt{na}}(x, 0)$ for every $a_0^x \in \mathtt{E}_0$.

   $\mathtt{lab}$ naturally induces the functions $\mathtt{typ}$, $\mathtt{mod}$, $\mathtt{loc}$, $\mathtt{val_r}$, and $\mathtt{val_w}$ that return (when applicable) the type ($\mathtt{R}$, $\mathtt{W}$ or $\mathtt{F}$), mode, location, and read/written value of an event.

   For $\mathtt{T} \in \{\mathtt{R}, \mathtt{W}, \mathtt{F}\}$, $\mathtt{T}$ denotes the set $\{e \in \mathtt{E} \mid \mathtt{typ}(e) = \mathtt{T}\}$. We also concatenate the event sets notations, use subscripts to denote the accessed location, and superscripts for modes (*e.g.,* $\mathtt{RW} = \mathtt{R} \cup \mathtt{W}$ and $\mathtt{W}_x^{\sqsupseteq \mathtt{rel}}$ denotes all events $a \in \mathtt{W}$ with $\mathtt{loc}(a) = x$ and $\mathtt{mod}(a) \sqsupseteq \mathtt{rel}$).
3. a strict partial order $\mathtt{sb} \subseteq \mathtt{E} \times \mathtt{E}$, called *sequenced-before,* which orders the initialization events before all other events, *i.e.,* $\mathtt{E}_0 \times (\mathtt{E} \setminus \mathtt{E}_0) \subseteq \mathtt{sb}$.
4. a binary relation $\mathtt{rmw} \subseteq [\mathtt{R}]; (\mathtt{sb}|_{\mathsf{imm}} \cap =_{\mathtt{loc}}); [\mathtt{W}]$, called *read-modify-write pairs*, such that for every $\langle a, b\rangle \in \mathtt{rmw}$, $\langle \mathtt{mod}(a), \mathtt{mod}(b)\rangle$ is one of the following:
   - $\langle \mathtt{rlx}, \mathtt{rlx}\rangle$ ($\mathtt{RMW}^{\mathtt{rlx}}$)
   - $\langle \mathtt{acq}, \mathtt{rlx}\rangle$ ($\mathtt{RMW}^{\mathtt{acq}}$)
   - $\langle \mathtt{rlx}, \mathtt{rel}\rangle$ ($\mathtt{RMW}^{\mathtt{rel}}$)
   - $\langle \mathtt{acq}, \mathtt{rel}\rangle$ ($\mathtt{RMW}^{\mathtt{acqrel}}$)
   - $\langle \mathtt{sc}, \mathtt{sc}\rangle$ ($\mathtt{RMW}^{\mathtt{sc}}$)

   We denote by $\mathtt{At}$ the set of all events in $\mathtt{E}$ that are a part of an $\mathtt{rmw}$ edge (that is, $\mathtt{At} = \mathit{dom}(\mathtt{rmw}) \cup \mathit{codom}(\mathtt{rmw})$).

   Note that our executions represent RMWs differently from C11 executions. Here each RMW is represented as two events, a read and a write, related by the $\mathtt{rmw}$ relation, whereas in C11 they are represented by single $\mathtt{RMW}$ events, which act as both the read and the write of the RMW. Our choice is in line with the Power and ARM memory models, and simplifies the formal development (*e.g.,* the definition of receptiveness).
5. a binary relation $\mathtt{rf} \subseteq [\mathtt{W}]; =_{\mathtt{loc}}; [\mathtt{R}]$, called *reads-from*, satisfying (i) $\mathtt{val_w}(a) = \mathtt{val_r}(b)$ for every $\langle a, b\rangle \in \mathtt{rf}$; and (ii) $a_1 = a_2$ whenever $\langle a_1, b\rangle, \langle a_2, b\rangle \in \mathtt{rf}$.
6. a strict partial order $\mathtt{mo}$ on $\mathtt{W}$, called *modification order*, which is a disjoint union of relations $\{\mathtt{mo}_x\}_{x \in \mathtt{Loc}}$, such that each $\mathtt{mo}_x$ is a strict total order on $\mathtt{W}_x$.

In what follows, to resolve ambiguities, we may include a prefix "G." to refer to the components of an execution G.

Executions of a given program represent prefixes of traces of shared memory accesses and fences that are generated by the program. In this paper, we only consider "partitioned" programs of the form $\|_{i \in \mathsf{Tid}}\ c_i$, where Tid is a finite set of thread identifiers, $\|$ denotes parallel composition, and each c_i is a sequential program. Then, the set of executions associated with a given program is defined by induction over the structure of sequential programs. We do not define formally this construction (it depends on the particular syntax and features of the source programming language). In this initial stage the read values are not restricted whatsoever (and $\mathtt{rf}$ and $\mathtt{mo}$ are arbitrary). Note that the set of executions of a program P is taken to be *prefix-closed*: an $\mathtt{sb}$-prefix of an execution of P (which includes at least the initialization events) is also considered to be an execution of P. By *full* executions of P, we refer to executions that represent traces generated by the whole program P.

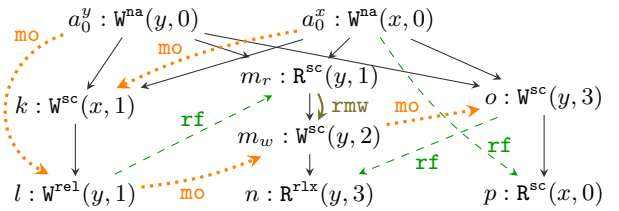

Figure 7. An execution of Z6.U.

We show an example of an execution in Fig. 7. This is a full execution of the Z6.U program, and is essentially the same as the C11 execution shown in Fig. 3, except for the representation of RMWs (see Item 4 above).

3.2 Consistent Executions

The main part of the memory model is filtering the consistent executions among all executions of the program. The first obvious restriction is that every read should read some written value (formally, $\mathtt{R} \subseteq \mathit{codom}(\mathtt{rf})$). We refer to such executions as *complete*.

To state the other constraints we use a number of derived relations:

$$\mathtt{rb} \triangleq \mathtt{rf}^{-1}; \mathtt{mo} \qquad (\textit{reads-before})$$

$$\mathtt{eco} \triangleq (\mathtt{rf} \cup \mathtt{mo} \cup \mathtt{rb})^+ \qquad (\textit{extended coherence order})$$

$$\mathtt{rs} \triangleq [\mathtt{W}]; \mathtt{sb}|_{\mathtt{loc}}^?; [\mathtt{W}^{\sqsupseteq \mathtt{rlx}}]; (\mathtt{rf}; \mathtt{rmw})^* \qquad (\textit{release sequence})$$

$$\mathtt{sw} \triangleq [\mathtt{E}^{\sqsupseteq \mathtt{rel}}]; ([\mathtt{F}]; \mathtt{sb})^?; \mathtt{rs}; \mathtt{rf}; [\mathtt{R}^{\sqsupseteq \mathtt{rlx}}]; (\mathtt{sb}; [\mathtt{F}])^?; [\mathtt{E}^{\sqsupseteq \mathtt{acq}}] \qquad (\textit{synchronizes with})$$

$$\mathtt{hb} \triangleq (\mathtt{sb} \cup \mathtt{sw})^+ \qquad (\textit{happens-before})$$

The first two, $\mathtt{rb}$ and $\mathtt{eco}$, are as described previously. Note that since the modification order, $\mathtt{mo}$, is transitive, we have $\mathtt{eco} = \mathtt{rf} \cup (\mathtt{mo} \cup \mathtt{rb}); \mathtt{rf}^?$ in every execution.

The other three relations, $\mathtt{rs}$, $\mathtt{sw}$ and $\mathtt{hb}$, are taken from [30]. Intuitively, $\mathtt{hb}$ records when an event is globally perceived as occurring before another one. It is defined in terms of two more basic relations. First, the *release sequence* ($\mathtt{rs}$) of a write contains the write itself and all later atomic writes

to the same location in the same thread, as well as all RMWs that recursively read from such writes. Next, a release event a *synchronizes with* ($\mathtt{sw}$) an acquire event b, whenever b (or, in case b is a fence, some $\mathtt{sb}$-prior read) reads from the release sequence of a (or in case a is a fence, of some $\mathtt{sb}$-later write). Then, we say that an event a *happens-before* ($\mathtt{hb}$) an event b if there is a path from a to b consisting of $\mathtt{sb}$ and $\mathtt{sw}$ edges.

Finally, we define the *SC-before* relation, $\mathtt{scb}$, and the *partial SC* relations, $\mathtt{psc}_{\text{base}}$ and $\mathtt{psc}_{\mathtt{F}}$, as follows:

$$\begin{aligned}
\mathtt{sb}|_{\neq \mathtt{loc}} &\triangleq \mathtt{sb} \setminus \mathtt{sb}|_{\mathtt{loc}} \\
\mathtt{scb} &\triangleq \mathtt{sb} \cup \mathtt{sb}|_{\neq \mathtt{loc}}; \mathtt{hb}; \mathtt{sb}|_{\neq \mathtt{loc}} \cup \mathtt{hb}|_{\mathtt{loc}} \cup \mathtt{mo} \cup \mathtt{rb} \\
\mathtt{psc}_{\text{base}} &\triangleq \left([\mathtt{E}^{\mathtt{sc}}] \cup [\mathtt{F}^{\mathtt{sc}}]; \mathtt{hb}^{?}\right); \mathtt{scb}; \left([\mathtt{E}^{\mathtt{sc}}] \cup \mathtt{hb}^{?}; [\mathtt{F}^{\mathtt{sc}}]\right) \\
\mathtt{psc}_{\mathtt{F}} &\triangleq [\mathtt{F}^{\mathtt{sc}}]; (\mathtt{hb} \cup \mathtt{hb}; \mathtt{eco}; \mathtt{hb}); [\mathtt{F}^{\mathtt{sc}}] \\
\mathtt{psc} &\triangleq \mathtt{psc}_{\text{base}} \cup \mathtt{psc}_{\mathtt{F}}
\end{aligned}$$

Using these derived relations, RC11 imposes four constraints on executions:

Definition 1. An execution G is called RC11-*consistent* if it is complete and the following hold:

- $\mathtt{hb}; \mathtt{eco}^{?}$ is irreflexive. (COHERENCE)
- $\mathtt{rmw} \cap (\mathtt{rb}; \mathtt{mo}) = \emptyset$. (ATOMICITY)
- $\mathtt{psc}$ is acyclic. (SC)
- $\mathtt{sb} \cup \mathtt{rf}$ is acyclic. (NO-THIN-AIR)

COHERENCE ensures that programs with only one shared location are sequentially consistent, as at least two locations are needed for a cycle in $\mathtt{sb} \cup \mathtt{eco}$. ATOMICITY ensures that the read and the write comprising a RMW are adjacent in $\mathtt{eco}$: there is no write event in between. The SC condition is the main novelty of RC11 and is used to give semantics to SC accesses and fences. Finally, NO-THIN-AIR rules out thin-air behaviors, albeit at a performance cost, as we will see in §5.

3.3 Program Outcomes

Finally, in order to allow the compilation of non-atomic read and writes to plain machine load and store instructions (as well as the compiler to reorder such accesses), RC11 follows the *"catch-fire"* approach: races on non-atomic accesses result in undefined behavior, that is, any outcome is allowed. Formally, it is defined as follows.

Definition 2. Two events a and b are called *conflicting* in an execution G if $a, b \in \mathtt{E}$, $\mathtt{W} \in \{\mathtt{typ}(a), \mathtt{typ}(b)\}$, $a \neq b$, and $\mathtt{loc}(a) = \mathtt{loc}(b)$. A pair $\langle a, b\rangle$ is called a *race* in G (denoted $\langle a, b\rangle \in \mathtt{race}$) if a and b are conflicting events in G, and $\langle a, b\rangle \notin \mathtt{hb} \cup \mathtt{hb}^{-1}$.

Definition 3. An execution G is called *racy* if there is some $\langle a, b\rangle \in \mathtt{race}$ with $\mathtt{na} \in \{\mathtt{mod}(a), \mathtt{mod}(b)\}$. A program P has *undefined behavior* under RC11 if it has some racy RC11-consistent execution.

Definition 4. The *outcome* of an execution G is the function assigning to every location x the value written by the $\mathtt{mo}$-maximal event in $\mathtt{W}_x$. We say that $O : \mathsf{Loc} \to \mathsf{Val}$ is an *outcome of a program* P *under* RC11 if either O is an outcome of some RC11-consistent full execution of P, or P has undefined behavior under RC11.

3.4 Comparison with C11

Besides the new SC and NO-THIN-AIR conditions, RC11 differs in a few other ways from C11.

- It does not support *consume* accesses, a premature feature of C11 that is not implemented by major compilers, nor *locks*, as they can be straightforwardly implemented with release-acquire accesses.
- For simplicity, it assumes all locations are initialized.
- It incorporates the fixes proposed by Vafeiadis *et al.* [30], namely (i) the strengthening of the release sequences definition, (ii) the removal of restrictions about different threads in the definition of synchronization, and (iii) the lack of distinction between atomic and non-atomic locations (and accordingly omitting the problematic $\mathtt{rf} \subseteq \mathtt{hb}$ condition for non-atomic locations). The third fix avoids "out-of-thin-air" problems that arise when performing non-atomic accesses to atomic location [6, §5].
- It does not consider "unsequenced races" between atomic accesses to have undefined behavior. Our results are not affected by such undefined behavior.

We have also made three presentational changes: (1) we have a much more concise axiomatization of coherence; (2) we model RMWs using two events; and (3) we do not have a total order over SC atomics.

Proposition 1. RC11*'s* COHERENCE *condition is equivalent to the conjunction of the following constraints of C11:*

- $\mathtt{hb}$ *is irreflexive.* (IRREFLEXIVE-HB)
- $\mathtt{rf}; \mathtt{hb}$ *is irreflexive.* (NO-FUTURE-READ)
- $\mathtt{mo}; \mathtt{rf}; \mathtt{hb}$ *is irreflexive.* (COHERENCE-RW)
- $\mathtt{mo}; \mathtt{hb}$ *is irreflexive.* (COHERENCE-WW)
- $\mathtt{mo}; \mathtt{hb}; \mathtt{rf}^{-1}$ *is irreflexive.* (COHERENCE-WR)
- $\mathtt{mo}; \mathtt{rf}; \mathtt{hb}; \mathtt{rf}^{-1}$ *is irreflexive.* (COHERENCE-RR)

Proposition 2. *The* SC *condition is equivalent to requiring the existence of a total strict order* $\mathtt{S}$ *on* $\mathtt{E}^{\mathtt{sc}}$ *such that* $\mathtt{S}; \mathtt{psc}$ *is irreflexive.*

Finally, the next proposition ensures that without mixing SC and non-SC accesses to the same location, RC11 supplies the stronger guarantee of C11. As a consequence, programmers that never mix such accesses may completely ignore the difference between RC11 and C11 regarding SC accesses.

Proposition 3. *If SC accesses are to distinguished locations (for every* $a, b \in \mathtt{E} \setminus \mathtt{E}_0$*, if* $\mathtt{mod}(a) = \mathtt{sc}$ *and* $\mathtt{loc}(a) = \mathtt{loc}(b)$ *then* $\mathtt{mod}(b) = \mathtt{sc}$*) then* $[\mathtt{E}^{\mathtt{sc}}]; \mathtt{hb}; [\mathtt{E}^{\mathtt{sc}}] \subseteq \mathtt{psc}^{+}$.

4. Compilation to x86-TSO

In this section, we present the x86-TSO memory model, and show that its intended compilation scheme is sound. We use a

$(\!\|\mathtt{R}\|\!)$	$\triangleq$ MOV (from memory)	$(\!\|\mathtt{W}^{\sqsubseteq\mathtt{rel}}\|\!)$	$\triangleq$ MOV (to memory)
$(\!\|\mathtt{W}^{\mathtt{sc}}\|\!)$	$\triangleq$ MOV;MFENCE	$(\!\|\mathtt{RMW}\|\!)$	$\triangleq$ CMPXCHG
$(\!\|\mathtt{F}^{\neq\mathtt{sc}}\|\!)$	$\triangleq$ No operation	$(\!\|\mathtt{F}^{\mathtt{sc}}\|\!)$	$\triangleq$ MFENCE

Figure 8. Compilation to TSO.

declarative model of x86-TSO from [17], that we denote by TSO. By [25, Theorem 3] and [17, Theorem 5], this definition is equivalent to the better known operational one. TSO executions are similar to the ones defined above, with the following exceptions:

- Read/write/fence labels have the form $\mathtt{R}(x,v)$, $\mathtt{W}(x,v)$, and F (they do not include a "mode"). In addition, labels may also be $\mathtt{RMW}(x,v_r,v_w)$, and executions do not include an rmw component (*i.e.,* RMWs are represented with a single event). We use RMW to denote the set of all events $a \in \mathtt{E}$ with $\mathtt{typ}(a) = \mathtt{RMW}$.
- The modification order, mo, is a strict total order on $\mathtt{W} \cup \mathtt{RMW} \cup \mathtt{F}$ (rather than a union of total order on writes to the same location).
- Happens-before is given by $\mathtt{hb} \triangleq (\mathtt{sb} \cup \mathtt{rf})^+$.
- Reads-before is given by $\mathtt{rb} \triangleq \mathtt{rf}^{-1}; \mathtt{mo}|_{\mathtt{loc}} \setminus [\mathtt{E}]$.

Remark 3. Lahav *et al.* [17] treat fence instructions as syntactic sugar for RMWs of a distinguished location. Here, we have fences as primitive instructions that induce fence events in TSO executions.

Definition 5. A TSO execution G is TSO-*consistent* if it is complete and the following hold:

1. hb is irreflexive.
2. mo;hb is irreflexive.
3. rb;hb is irreflexive.
4. rb;mo is irreflexive.
5. rb;mo;rfe;sb is irreflexive (where $\mathtt{rfe} = \mathtt{rf} \setminus \mathtt{sb}$).
6. rb;mo;[RMW ∪ F];sb is irreflexive.

Unlike RC11, well-formed TSO programs do not have undefined behavior. Thus, a function $O : \mathsf{Loc} \to \mathsf{Val}$ is an *outcome of a* TSO *program* P if it is an outcome of some TSO-consistent full execution of P (see Def. 4).

Fig. 8 presents the compilation scheme from C11 to x86-TSO that is implemented in the GCC and the LLVM compilers. Since TSO provides strong consistency guarantees, it allows most language primitives to be compiled to plain loads and stores. Barriers are only needed for the compilation of SC writes. Our next theorem says that this compilation scheme is also correct for RC11.

Theorem 1. *For a program P, denote by $(\!|P|\!)$ the* TSO *program obtained by compiling P using the scheme in Fig. 8. Then, given a program P, every outcome of $(\!|P|\!)$ under* TSO *is an outcome of P under* RC11.

Proof (Outline). We consider the compilation as if it happens in three steps, and prove the soundness of each step:

1. All non-atomic/relaxed accesses are strengthened to release/acquire ones, and all relaxed/release/acquire RMWs are strengthened to acquire-release ones. It is easy to see that this step does not introduce new outcomes (see §7).
2. All non-SC fences are removed. Due to the previous step, it is easy to show that non-SC fences have no effect.
3. The mappings in Fig. 8 are applied. The correctness of this step, given in [3, Appendix K], is established by showing that given a TSO-consistent TSO execution G_t of $(\!|P|\!)$ (where P has no non-SC fences), there exists an RC11-consistent execution G of P that has the same outcome as G_t. □

In fact, the proof of Thm. 1 establishes the correctness of compilation even for a strengthening of RC11 obtained by replacing the scb relation by $\mathtt{scb}' \triangleq \mathtt{hb} \cup \mathtt{mo} \cup \mathtt{rb}$. This entails that the original C11 model, as well as Batty *et al.*'s strengthening [5], are correctly compiled to x86-TSO. Additionally, the proof only assumes the existence of an MFENCE between every store originated from an SC write and load originated from an SC read. The compilation scheme in Fig. 8 achieves this by placing an MFENCE after each store that originated from an SC write. An alternative correct compilation scheme may place MFENCE before SC reads, rather than after SC writes [1]. (Since there are typically more SC reads than SC writes in programs, the latter scheme is less preferred.)

Remark 4. The compilation scheme that places MFENCE before SC reads can be shown to be sound even for a very strong SC condition that requires acyclicity of

$$\mathtt{psc}_{\mathrm{strong}} = ([\mathtt{E}^{\mathtt{sc}}] \cup [\mathtt{F}^{\mathtt{sc}}]; \mathtt{hb}^?); (\mathtt{hb} \cup \mathtt{eco}); ([\mathtt{E}^{\mathtt{sc}}] \cup \mathtt{hb}^?; [\mathtt{F}^{\mathtt{sc}}]).$$

To prove this (see [3, Appendix K.1]), we are able to follow a simpler approach utilizing the recent result of Lahav and Vafeiadis [19] that provides a characterization of TSO in terms of program transformations (or "compiler optimizations"). This allows one to reduce compilation correctness to soundness of certain transformations. The preferred compilation scheme to x86-TSO, which uses barriers after SC writes (see Fig. 8), is unsound if one requires acyclicity of $\mathtt{psc}_{\mathrm{strong}}$, or even if one requires acyclicity of $[\mathtt{E}^{\mathtt{sc}}]; (\mathtt{sb} \cup \mathtt{eco}); [\mathtt{E}^{\mathtt{sc}}]$. To see this, consider the following variant of SB:

$$\begin{array}{l|l} x :=_{\mathtt{rel}} 1 & y :=_{\mathtt{rel}} 1 \\ a := x_{\mathtt{sc}} \mathbin{/\!/} 1 & c := y_{\mathtt{sc}} \mathbin{/\!/} 1 \\ b := y_{\mathtt{sc}} \mathbin{/\!/} 0 & d := x_{\mathtt{sc}} \mathbin{/\!/} 0 \end{array} \qquad \text{(SB+rfis)}$$

Any execution of this program that yields the annotated behavior has a cycle in $[\mathtt{E}^{\mathtt{sc}}]; (\mathtt{sb} \cup \mathtt{eco}); [\mathtt{E}^{\mathtt{sc}}]$ (we have rb;rf both from $\mathtt{R}^{\mathtt{sc}}(x,0)$ to $\mathtt{R}^{\mathtt{sc}}(x,1)$, and from $\mathtt{R}^{\mathtt{sc}}(y,0)$ to $\mathtt{R}^{\mathtt{sc}}(y,1)$). However, since the program has no SC writes, following Fig. 8, all accesses are compiled to plain accesses, and x86-TSO clearly allows this behavior.

5. Compilation to Power

In this section, we present the Power model and the mappings of language operations to Power instructions. We then prove the correctness of compilation from RC11 to Power.

As a model of the Power architecture, we use the recent declarative model by Alglave *et al.* [4], which we denote by Power. Its executions are similar to the RC11's execution, with the following exceptions:

- Power executions track syntactic dependencies between events in the same thread, and derive a relation called *preserved program order*, denoted ppo, which is a subset of sb guaranteed to be preserved. The exact definition of ppo is quite intricate, and is included in [3, Appendix F].
- Read/write labels have the form $\mathtt{R}(x, v)$ and $\mathtt{W}(x, v)$ (they do not include a "mode"). Power has two types of fence events: a "lightweight fence" and a "full fence". We denote by $\mathtt{F}^{\mathtt{lwsync}}$ and $\mathtt{F}^{\mathtt{sync}}$ the set of all lightweight fence and full fence events in a Power execution. Power's "instruction fence" (isync) is used to derive ppo but is not recorded in executions.

In addition to ppo, the following additional derived relations are needed to define Power-consistency (see [4] for further explanations and details).

- $\mathtt{sync} \triangleq [\mathtt{RW}]; \mathtt{sb}; [\mathtt{F}^{\mathtt{sync}}]; \mathtt{sb}; [\mathtt{RW}]$
- $\mathtt{lwsync} \triangleq [\mathtt{RW}]; \mathtt{sb}; [\mathtt{F}^{\mathtt{lwsync}}]; \mathtt{sb}; [\mathtt{RW}] \setminus (\mathtt{W} \times \mathtt{R})$
- $\mathtt{fence} \triangleq \mathtt{sync} \cup \mathtt{lwsync}$ (*fence order*)
- $\mathtt{hb} \triangleq \mathtt{ppo} \cup \mathtt{fence} \cup \mathtt{rfe}$ (*Power's happens-before*)
- $\mathtt{prop}_1 \triangleq [\mathtt{W}]; \mathtt{rfe}^?; \mathtt{fence}; \mathtt{hb}^*; [\mathtt{W}]$
- $\mathtt{prop}_2 \triangleq (\mathtt{moe} \cup \mathtt{rbe})^?; \mathtt{rfe}^?; (\mathtt{fence}; \mathtt{hb}^*)^?; \mathtt{sync}; \mathtt{hb}^*$
- $\mathtt{prop} \triangleq \mathtt{prop}_1 \cup \mathtt{prop}_2$ (*propagation relation*)

where for every relation c (e.g., rf, mo, etc.), we denote by ce its thread-external restriction. Formally, $\mathtt{ce} = \mathtt{c} \setminus \mathtt{sb}$.

Definition 6. A Power execution G is Power-*consistent* if it is complete and the following hold:

1. $\mathtt{sb}|_{\mathtt{loc}} \cup \mathtt{rf} \cup \mathtt{rb} \cup \mathtt{mo}$ is acyclic. (SC-PER-LOC)
2. $\mathtt{rbe}; \mathtt{prop}; \mathtt{hb}^*$ is irreflexive. (OBSERVATION)
3. $\mathtt{mo} \cup \mathtt{prop}$ is acyclic. (PROPAGATION)
4. $\mathtt{rmw} \cap (\mathtt{rbe}; \mathtt{moe}) = \emptyset$. (POWER-ATOMICITY)
5. hb is acyclic. (POWER-NO-THIN-AIR)

Remark 5. The model in [4] contains an additional constraint: $\mathtt{mo} \cup [\mathtt{At}]; \mathtt{sb}; [\mathtt{At}]$ should be acyclic (recall that $\mathtt{At} = \mathit{dom}(\mathtt{rmw}) \cup \mathit{codom}(\mathtt{rmw})$). Since none of our proofs requires this property, we excluded it from Def. 6.

Like in the case of TSO, we say that a function $O : \mathsf{Loc} \to \mathsf{Val}$ is an *outcome of a* Power *program* P if it is an outcome of some Power-consistent full execution of P (see Def. 4).

As already mentioned, the two compilation schemes from C11 to Power that have been proposed in the literature [1] differ only in the mappings used for SC accesses (see Fig. 10). The first scheme follows the *leading sync* convention, and places a sync fence *before* each SC access. The alternative scheme follows the *trailing sync* convention, and places a sync fence *after* each SC access. Importantly, the same scheme should be used for all SC accesses in the program, since mixing the schemes is unsound. The mappings for the non-SC accesses and fences are common to both schemes and are shown in Fig. 9. Note that our compilation of relaxed reads is stronger than the one proposed for C11 (see §2.3).

Our main theorem says that the compilation schemes are correct.

$(\!\lvert \mathtt{R}^{\mathtt{na}} \rvert\!)$	$\triangleq$ ld	$(\!\lvert \mathtt{W}^{\mathtt{na}} \rvert\!)$	$\triangleq$ st
$(\!\lvert \mathtt{R}^{\mathtt{rlx}} \rvert\!)$	$\triangleq$ ld;cmp;bc	$(\!\lvert \mathtt{W}^{\mathtt{rlx}} \rvert\!)$	$\triangleq$ st
$(\!\lvert \mathtt{R}^{\mathtt{acq}} \rvert\!)$	$\triangleq$ ld;cmp;bc;isync	$(\!\lvert \mathtt{W}^{\mathtt{rel}} \rvert\!)$	$\triangleq$ lwsync;st
$(\!\lvert \mathtt{F}^{\neq \mathtt{sc}} \rvert\!)$	$\triangleq$ lwsync	$(\!\lvert \mathtt{F}^{\mathtt{sc}} \rvert\!)$	$\triangleq$ sync
$(\!\lvert \mathtt{RMW}^{\mathtt{rlx}} \rvert\!)$	$\triangleq$ L:lwarx;cmp;bc Le;stwcx.;bc L;Le:		
$(\!\lvert \mathtt{RMW}^{\mathtt{acq}} \rvert\!)$	$\triangleq (\!\lvert \mathtt{RMW}^{\mathtt{rlx}} \rvert\!)$;isync		
$(\!\lvert \mathtt{RMW}^{\mathtt{rel}} \rvert\!)$	$\triangleq$ lwsync;$(\!\lvert \mathtt{RMW}^{\mathtt{rlx}} \rvert\!)$		
$(\!\lvert \mathtt{RMW}^{\mathtt{acqrel}} \rvert\!)$	$\triangleq$ lwsync;$(\!\lvert \mathtt{RMW}^{\mathtt{rlx}} \rvert\!)$;isync		

Figure 9. Compilation of non-SC primitives to Power.

Leading sync		Trailing sync	
$(\!\lvert \mathtt{R}^{\mathtt{sc}} \rvert\!)$	$\triangleq$ sync;$(\!\lvert \mathtt{R}^{\mathtt{acq}} \rvert\!)$	$(\!\lvert \mathtt{R}^{\mathtt{sc}} \rvert\!)$	$\triangleq$ ld;sync
$(\!\lvert \mathtt{W}^{\mathtt{sc}} \rvert\!)$	$\triangleq$ sync;st	$(\!\lvert \mathtt{W}^{\mathtt{sc}} \rvert\!)$	$\triangleq (\!\lvert \mathtt{W}^{\mathtt{rel}} \rvert\!)$;sync
$(\!\lvert \mathtt{RMW}^{\mathtt{sc}} \rvert\!)$	$\triangleq$ sync;$(\!\lvert \mathtt{RMW}^{\mathtt{acq}} \rvert\!)$	$(\!\lvert \mathtt{RMW}^{\mathtt{sc}} \rvert\!)$	$\triangleq (\!\lvert \mathtt{RMW}^{\mathtt{rel}} \rvert\!)$;sync

Figure 10. Compilations of SC accesses to Power.

Theorem 2. *For a program P, denote by $(\!\lvert P \rvert\!)$ the* Power *program obtained by compiling P using the scheme in Fig. 9 and either of the schemes in Fig. 10 for SC accesses. Then, given a program P, every outcome of $(\!\lvert P \rvert\!)$ under* Power *is an outcome of P under* RC11.

Proof (Outline). The main idea is to consider the compilation as if it happens in three steps, and prove the soundness of each step:

1. Leading sync: Each $\mathtt{R}^{\mathtt{sc}}/\mathtt{W}^{\mathtt{sc}}/\mathtt{RMW}^{\mathtt{sc}}$ in P is replaced by $\mathtt{F}^{\mathtt{sc}}$ followed by $\mathtt{R}^{\mathtt{acq}}/\mathtt{W}^{\mathtt{rel}}/\mathtt{RMW}^{\mathtt{acqrel}}$.
   Trailing sync: Each $\mathtt{R}^{\mathtt{sc}}/\mathtt{W}^{\mathtt{sc}}/\mathtt{RMW}^{\mathtt{sc}}$ in P is replaced by $\mathtt{R}^{\mathtt{acq}}/\mathtt{W}^{\mathtt{rel}}/\mathtt{RMW}^{\mathtt{acqrel}}$ followed by $\mathtt{F}^{\mathtt{sc}}$.
2. The mappings in Fig. 9 are applied.
3. Leading sync: Pairs of the form sync;lwsync that originated from $\mathtt{W}^{\mathtt{sc}}/\mathtt{RMW}^{\mathtt{sc}}$ are reduced to sync (eliminating the redundant lwsync).
   Trailing sync: Any cmp;bc;isync;sync sequences originated from $\mathtt{R}^{\mathtt{sc}}/\mathtt{RMW}^{\mathtt{sc}}$ are reduced to sync (eliminating the redundant cmp;bc;isync).

The resulting Power program is clearly identical to the one obtained by applying the mappings in Figures 9 and 10.

X \ Y	$\mathtt{R}_y^{o_2}$	$\mathtt{W}_y^{o_2}$	$\mathtt{RMW}_y^{o_2}$	$\mathtt{F}^{o_2}$
$\mathtt{R}_x^{o_1}$	$o_1 \sqsubseteq \mathtt{rlx}$	$o_1, o_2 \sqsubseteq \mathtt{rlx} \wedge (o_1 = \mathtt{na} \vee o_2 = \mathtt{na})$	$o_1 = \mathtt{na} \wedge o_2 \sqsubseteq \mathtt{acq}$	$o_1 \neq \mathtt{rlx} \wedge o_2 = \mathtt{acq}$
$\mathtt{W}_x^{o_1}$	$o_1 \neq \mathtt{sc} \vee o_2 \neq \mathtt{sc}$	$o_2 \sqsubseteq \mathtt{rlx}$	$o_2 \sqsubseteq \mathtt{acq}$	$o_2 = \mathtt{acq}$
$\mathtt{RMW}_x^{o_1}$	$o_1 \sqsubseteq \mathtt{rel}$	$o_1 \sqsubseteq \mathtt{rel} \wedge o_2 = \mathtt{na}$	–	$o_1 \sqsupseteq \mathtt{acq} \wedge o_2 = \mathtt{acq}$
$\mathtt{F}^{o_1}$	$o_1 = \mathtt{rel}$	$o_1 = \mathtt{rel} \wedge o_2 \neq \mathtt{rlx}$	$o_1 = \mathtt{rel} \wedge o_2 \sqsupseteq \mathtt{rel}$	$o_1 = \mathtt{rel} \wedge o_2 = \mathtt{acq}$

Table 1. Deorderable pairs of accesses/fences (x and y are distinct locations).

The soundness for each step (that is, none of them introduces additional outcomes) is established in [3, Lemmas E.3 and H.3 and Appendix F.3]. □

The main difficulty (and novelty of our proof) lies in proving soundness of the second step, and more specifically in establishing the NO-THIN-AIR condition. Since Power, unlike RC11, does not generally forbid $\mathtt{sb} \cup \mathtt{rf}$ cycles, we have to show that such cycles can be untangled to produce a racy RC11-consistent execution, witnessing the undefined behavior. Here, the idea is, similar to DRF-SC proofs, to detect a first rf edge that closes an $\mathtt{sb} \cup \mathtt{rf}$ cycle, and replace it by a different rf edge that avoids the cycle. This is highly non-trivial because it is unclear how to define a "first" rf edge when $\mathtt{sb} \cup \mathtt{rf}$ is cyclic. To solve this problem, we came up with a different ordering of events, which does not include all sb edges, and Power ensures to be acyclic (a relation we call *Power-before* in [3, Appendix G]).

For completeness, we also show that the conditional branch after the relaxed read is only necessary if we care about enforcing the NO-THIN-AIR condition. That is, let weakRC11 be the model obtained from RC11 by omitting the NO-THIN-AIR condition, and denote by $(\!|P|\!)_{\mathsf{weak}}$ the Power program obtained by compiling P as above, except that relaxed reads are compiled to plain loads (again, with either leading or trailing syncs for SC accesses). Then, this scheme is correct with respect to the weakRC11 model.

Theorem 3 (Compilation of weakRC11 to Power)**.** *Given a program P, every outcome of $(\!|P|\!)_{\mathsf{weak}}$ under* Power *is an outcome of P under* weakRC11.

Finally, we note that it is also possible to use a lightweight fence (`lwsync`) instead of a fake control dependency and an instruction fence (`isync`) in the compilation of (all or some) acquire accesses.

6. Compilation to ARMv7

The ARMv7 model [4] is very similar to the Power model just presented in §5. There are only two differences.

First, while ARMv7 has analogues for Power's strong fence and instruction fence (`dmb` for `sync`, and `isb` for `isync`), it lacks an analogue for Power's lightweight fence (`lwsync`). Thus, on ARMv7 we have $\mathtt{F}^{\mathtt{lwsync}} = \emptyset$ and so $\mathtt{fence} = \mathtt{sync}$.

The second difference is that ARMv7 has a somewhat weaker *preserved program order*, ppo, than Power, which in particular does not always include $[\mathtt{R}]; \mathtt{sb}|_{\mathtt{loc}}; [\mathtt{W}]$ (following the model in [4]). In our Power compilation proofs, however, we never rely on this property of Power's ppo (see [3, Appendix F]).

The compilation schemes to ARMv7 are essentially the same as those to Power substituting the corresponding ARMv7 instructions for the Power ones: `dmb` instead of `sync` and `lwsync`, and `isb` instead of `isync`. The soundness of compilation to ARMv7 follows directly from Theorems 2 and 3.

We note that neither `GCC` (version 5.4) nor `LLVM` (version 3.9) map acquire reads into `ld;cmp;bc;isb`. Instead, they emit `ld;dmb` (that corresponds to Power's `ld;sync`). With this stronger compilation scheme, there is no correctness problem in compilation of C11 to ARMv7. Nevertheless, if one intends to use `isb`'s, the same correctness issue arises (*e.g.,* the one in Fig. 1), and RC11 overcomes this issue.

7. Correctness of Program Transformations

In this section, we list program transformations that are sound in RC11, and prove that this is the case. As in [30], to have a simple presentation, all of our arguments are performed at the *semantic* level, as if the transformations were applied to events in an execution. Thus, to prove soundness of a program transformation $P_{\mathsf{src}} \rightsquigarrow P_{\mathsf{tgt}}$, we are given an arbitrary RC11-consistent execution G_{tgt} of P_{tgt}, and construct a RC11-consistent execution G_{src} of P_{src}, such that either G_{src} and G_{tgt} have the same outcome or G_{src} is racy. In the former case, we also show that G_{tgt} is racy only if G_{src} is. Consequently, one obtains that every outcome of P_{tgt} under RC11 is also an outcome of P_{src} under RC11.

The soundness proofs (sketched in [3, Appendix I]) are mostly similar to the proofs in [30], with the main difference concerning the new SC condition.

Strengthening Strengthening transforms the mode o of an event in the source into o' in the target where $o \sqsubseteq o'$. Soundness of this transformation is trivial, because RC11-consistency is monotone with respect to the mode ordering.

Sequentialization Sequentialization merges two program threads into one, by interleaving their events in sb. Essentially sequentialization just adds edges to the sb relation. Its soundness trivially follows from the monotonicity of RC11-consistency with respect to sb.

$\mathtt{R}^o;\mathtt{R}^o \rightsquigarrow \mathtt{R}^o$		$\mathtt{W}^o;\mathtt{W}^o \rightsquigarrow \mathtt{W}^o$	
$\mathtt{W}^{\mathtt{sc}};\mathtt{R}^{\mathtt{sc}} \rightsquigarrow \mathtt{W}^{\mathtt{sc}}$		$\mathtt{W}^o;\mathtt{R}^{\mathtt{acq}} \rightsquigarrow \mathtt{W}^o$	
$\mathtt{RMW}^o;\mathtt{R}^{o_\mathrm{r}} \rightsquigarrow \mathtt{RMW}^o$		$\mathtt{RMW}^o;\mathtt{RMW}^o \rightsquigarrow \mathtt{RMW}^o$	
$\mathtt{W}^{o_\mathrm{w}};\mathtt{RMW}^o \rightsquigarrow \mathtt{W}^{o_\mathrm{w}}$		$\mathtt{F}^o;\mathtt{F}^o \rightsquigarrow \mathtt{F}^o$	

Figure 11. Mergeable pairs (assuming both accesses are to the same location). o_r denotes the maximal mode in $\{\mathtt{rlx},\mathtt{acq},\mathtt{sc}\}$ satisfying $o_\mathrm{r} \sqsubseteq o$; and o_w denotes the maximal mode in $\{\mathtt{rlx},\mathtt{rel},\mathtt{sc}\}$ satisfying $o_\mathrm{w} \sqsubseteq o$.

Deordering Table 1 defines the *deorderable* pairs, for which we proved the soundness of the transformation $\mathtt{X};\mathtt{Y} \rightsquigarrow \mathtt{X} \parallel \mathtt{Y}$ in RC11. (Note that reordering is obtained by applying deordering and sequentialization.) Generally speaking, RC11 supports all reorderings that are intended to be sound in C11 [30], except for load-store reorderings of relaxed accesses, which are unsound in RC11 due to the conservative NO-THIN-AIR condition (if one omits this condition, these reorderings are sound). Importantly, load-store reorderings of *non-atomic* accesses are sound due to the "catch-fire" semantics. The soundness of these reorderings (in the presence of NO-THIN-AIR) was left open in [30], and requires a non-trivial argument of the same nature as the one used to show NO-THIN-AIR in the compilation correctness proof.

Merging Merges are transformations of the form $\mathtt{X};\mathtt{Y} \rightsquigarrow \mathtt{Z}$, eliminating one memory access or fence. Fig. 11 defines the set of *mergeable* pairs. Note that using strengthening, the modes mentioned in Fig. 11 are upper bounds (*e.g.,* $\mathtt{R}_x^{\mathtt{acq}};\mathtt{R}_x^{\mathtt{rlx}}$ can be first strengthened to $\mathtt{R}_x^{\mathtt{acq}};\mathtt{R}_x^{\mathtt{acq}}$ and then merged). Generally speaking, RC11 supports all mergings that are intended to be mergeable in C11 [30].

Remark 6. The elimination of redundant read-after-write allows the write to be non-atomic. Nevertheless, an SC read cannot be eliminated in this case, unless it follows an SC write. Indeed, eliminating an SC read after a non-SC write is unsound in RC11. We note that the effectiveness of this optimization seems to be low, and, in fact, it is already unsound for the model in [5] (see [3, Appendix A.4] for a counterexample). Note also that read-after-RMW elimination does not allow the read to be an acquire read unless the update includes an acquire read (unlike read-after-write). This is due to release sequences: eliminating an acquire read after a relaxed update may remove the synchronization due to a release sequence ending in this update.

Register Promotion Finally, "register promotion" is sound in RC11. This global program transformation replaces all the accesses to a memory location by those to a register, provided that the location is used by only one thread. At the execution level, all accesses to a particular location are removed from the execution, provided that they are all $\mathtt{sb}$-related.

8. Programming Guarantees

In this section, we demonstrate that our semantics for SC atomics (*i.e.,* the SC condition in Def. 1) is not overly weak. We do so by proving theorems stating that programmers who follow certain defensive programming patterns can be assured that their programs exhibit no weak behaviors. The first such theorem is DRF-SC, which says that if a program has no races on non-SC accesses under SC semantics, then its outcomes under RC11 coincide with those under SC.

In our proofs we use the standard declarative definition of SC: an execution is SC-consistent if it is complete, satisfies ATOMICITY, and $\mathtt{sb} \cup \mathtt{rf} \cup \mathtt{mo} \cup \mathtt{rb}$ is acyclic [28].

Theorem 4. *If in all SC-consistent executions of a program P, every race $\langle a,b\rangle$ has $\mathtt{mod}(a) = \mathtt{mod}(b) = \mathtt{sc}$, then the outcomes of P under* RC11 *coincide with those under SC.*

Note that the NO-THIN-AIR condition is essential for the correctness of Thm. 4 (recall the LB+deps example).

Next, we show that adding a fence instruction between every two accesses to *shared* locations restores SC, or there remains a race in the program, in which case the program has undefined behavior.

Definition 7. A location x is *shared* in an execution G if $\langle a,b\rangle \notin \mathtt{sb} \cup \mathtt{sb}^{-1}$ for some distinct events $a,b \in \mathtt{E}_x$.

Theorem 5. *Let G be an* RC11*-consistent execution. Suppose that for every two distinct shared locations x and y, $[\mathtt{E}_x];\mathtt{sb};[\mathtt{E}_y] \subseteq \mathtt{sb};[\mathtt{F}^{\mathtt{sc}}];\mathtt{sb}$. Then, G is SC-consistent.*

We remark that for the proofs of Theorems 4 and 5, we do not need the full SC condition: for Thm. 4 it suffices for $[\mathtt{E}^{\mathtt{sc}}];(\mathtt{sb}\cup\mathtt{rf}\cup\mathtt{mo}\cup\mathtt{rb});[\mathtt{E}^{\mathtt{sc}}]$ to be acyclic; and for Thm. 5 it suffices for $[\mathtt{F}^{\mathtt{sc}}];\mathtt{sb};\mathtt{eco};\mathtt{sb};[\mathtt{F}^{\mathtt{sc}}]$ to be acyclic.

9. Related Work

The C11 memory model was designed by the C++ standards committee based on a paper by Boehm and Adve [10]. During the standardization process, Batty *et al.* [8] formalized the C11 model and proved soundness of its compilation to x86-TSO. They also proposed a number of key technical improvements to the model (such as some coherence axioms), which were incorporated into the standard.

Since then, however, a number of problems have been found with the C11 model. In 2012, Batty *et al.* [7] and Sarkar *et al.* [27] studied the compilation of C11 to Power, and incorrectly proved the correctness of two compilation schemes. In their proofs, from a consistent Power execution, they constructed a corresponding C11 execution, which they tried to prove consistent, but in doing so they forgot to check the overly strong condition S1. The examples shown in §1 and in §2.1 are counterexamples to their theorems.

Quite early on, a number of papers [12, 31, 24, 11] noticed the disastrous effects of thin-air behaviors allowed by the C11 model, and proposed strengthening the definition of consistency by disallowing $\mathtt{sb} \cup \mathtt{rf}$ cycles. Boehm and

Demsky [11] further discussed how the compilation schemes of relaxed accesses to Power and ARM would be affected by the change, but did not formally prove the correctness of their proposed schemes.

Next, Vafeiadis *et al.* [30] noticed a number of other problems with the C11 memory model, which invalidated a number of source-to-source program transformations that were assumed to hold. They proposed local fixes to those problems, and showed that these fixes enabled proving correctness of a number of local transformations. We have incorporated their fixes in the RC11-consistency definition.

Then, in 2016, Batty *et al.* [5] proposed a more concise semantics for SC atomics, whose presentation we have followed in our proposed RC11 model. As their semantics is stronger than C11, it cannot be compiled efficiently to Power, contradicting the claim of that paper. Moreover, as already discussed, SC fences are still too weak according to their model: in particular, putting them between every two accesses in a program with only atomic accesses does not guarantee SC.

Recently, Manerkar *et al.* [22] discovered the problem with trailing-sync compilation to Power (in particular, they observed the IRIW-acq-sc counterexample), and identified the mistake in the existing proof. Independently, we discovered the same problem, as well as the problem with leading-sync compilation. Moreover, in this paper, we have proposed a fix for both problems, and proven that it works.

A number of previous papers [31, 29, 18, 17] have studied only small fragments of the C11 model—typically the release/acquire fragment. Among these, Lahav *et al.* [17] proposed strengthening the semantics of SC fences in a different way from the way we do here, by treating them as read-modify-writes to a distinguished location. That strengthening, however, was considered in the restricted setting of only release/acquire accesses, and does not directly scale to the full set of C11 access modes. In fact, for the fragment containing only SC fences and release/acquire accesses, RC11-consistency is equivalent to RA-consistency that treats SC fences as RMWs to a distinguished location [17].

Finally, several solutions to the "out-of-thin-air" problem were recently suggested, *e.g.,* [26, 15, 16]. These solutions aim to avoid the performance cost of disallowing $\mathtt{sb} \cup \mathtt{rf}$ cycles, but none of them follows the declarative framework of C11. The conservative approach of disallowing $\mathtt{sb} \cup \mathtt{rf}$ cycles allows us to formulate our model in the style of C11.

10. Conclusion

In this paper, we have introduced the RC11 memory model, which corrects all the known problems of the C11 model (albeit at a performance cost for the "out-of-thin-air" problem). We have further proved (i) the correctness of compilation from RC11 to x86-TSO [25], Power and ARMv7 [4]; (ii) the soundness of various program transformations; (iii) a DRF-SC theorem; and (iv) a theorem showing that for programs without non-atomic accesses, weak behaviors can be always avoided by placing SC fences. It would be very useful to mechanize the proofs of this paper in a theorem prover; we leave this for future work.

A certain degree of freedom exists in the design of the SC condition. A very weak version, which maintains the two formal programming guarantees of this paper, would require acyclicity of $[\mathtt{E}^{\mathtt{sc}}];(\mathtt{sb} \cup \mathtt{rf} \cup \mathtt{mo} \cup \mathtt{rb});[\mathtt{E}^{\mathtt{sc}}] \cup [\mathtt{F}^{\mathtt{sc}}];\mathtt{sb};\mathtt{eco};\mathtt{sb};[\mathtt{F}^{\mathtt{sc}}]$. At the other extreme, one can require the acyclicity of $\mathtt{psc}_{\text{strong}} = ([\mathtt{E}^{\mathtt{sc}}] \cup [\mathtt{F}^{\mathtt{sc}}];\mathtt{hb}^{?});(\mathtt{hb} \cup \mathtt{eco});([\mathtt{E}^{\mathtt{sc}}] \cup \mathtt{hb}^{?};[\mathtt{F}^{\mathtt{sc}}])$, and either disallow mixing SC and non-SC accesses to the same location, or have rather expensive compilation schemes (for Power/ARMv7: compile release-acquire atomics exactly as the SC ones; for TSO: place a barrier before every SC read). Our choice of psc achieves the following: (i) it allows free mixing of different access modes to the same location in the spirit of C11; (ii) it ensures the correctness of the existing compilation schemes; and (iii) it coincides with $\mathtt{psc}_{\text{strong}}$ in the absence of mixing of SC and non-SC accesses to the same location.

Regarding the infamous "out-of-thin-air" problem, we employed in RC11 a conservative solution at the cost of including a fake control dependency after every relaxed read. While this was already considered a valid solution before, we are the first to prove the correctness of this compilation scheme, as well as the soundness of reordering of independent non-atomic accesses under this model. Correctness of an alternative scheme that places a lightweight fence before every relaxed write is left for future work. It would be interesting to evaluate the practical performance costs of each scheme. On the one hand, relaxed writes (which are not followed by a fence) are perhaps rare in real programs, compared to relaxed reads. On the other hand, a control dependency is cheaper than a lightweight fence, and relaxed reads are often anyway followed by a control dependency.

Another important future direction would be to combine our SC constraint with our recent operational model in [16], which prevents "out-of-thin-air" values (and avoids undefined behaviors altogether), while still allowing the compilation of relaxed reads and writes to plain loads and stores. This is, in particular, crucial for adopting a model like RC11 in a type-safe language, like Java, which cannot allow undefined behaviors. Integrating our SC condition in that model, however, is non-trivial because the model is defined in a very different style from C11, and thus we will have to find an equivalent operational way to check our SC condition.

Finally, extending RC11 with additional features of C11 (see §3.4) and establishing the correctness of compilation of RC11 to ARMv8 [13] are important future goals as well.

Acknowledgments

We thank Hans Boehm, Soham Chakraborty, Doug Lea, Peter Sewell and the PLDI'17 reviewers for their helpful feedback. This research was supported in part by Samsung Research Funding Center of Samsung Electronics under Project

Number SRFC-IT1502-07, and in part by an ERC Consolidator Grant for the project "RustBelt" (grant agreement no. 683289). The third author has been supported by a Korea Foundation for Advanced Studies Scholarship.

References

[1] C/C++11 mappings to processors, available at http://www.cl.cam.ac.uk/~pes20/cpp/cpp0xmappings.html. [Online; accessed 27-September-2016].

[2] Crossbeam: support for concurrent and parallel programming, available at https://github.com/aturon/crossbeam. [Online; accessed 24-October-2016].

[3] Supplementary material for this paper, available at http://plv.mpi-sws.org/scfix/.

[4] J. Alglave, L. Maranget, and M. Tautschnig. Herding cats: Modelling, simulation, testing, and data mining for weak memory. *ACM Trans. Program. Lang. Syst.*, 36(2):7:1–7:74, July 2014.

[5] M. Batty, A. F. Donaldson, and J. Wickerson. Overhauling SC atomics in C11 and OpenCL. In *POPL 2016*, pages 634–648. ACM, 2016.

[6] M. Batty, K. Memarian, K. Nienhuis, J. Pichon-Pharabod, and P. Sewell. The problem of programming language concurrency semantics. In *ESOP 2015*, pages 283–307. Springer, 2015.

[7] M. Batty, K. Memarian, S. Owens, S. Sarkar, and P. Sewell. Clarifying and compiling C/C++ concurrency: From C++11 to POWER. In *POPL 2012*, pages 509–520. ACM, 2012.

[8] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL 2011*, pages 55–66. ACM, 2011.

[9] H.-J. Boehm. Can seqlocks get along with programming language memory models? In *MSPC 2012*, pages 12–20. ACM, 2012.

[10] H.-J. Boehm and S. V. Adve. Foundations of the C++ concurrency memory model. In *PLDI 2008*, pages 68–78. ACM, 2008.

[11] H.-J. Boehm and B. Demsky. Outlawing ghosts: Avoiding out-of-thin-air results. In *MSPC 2014*, pages 7:1–7:6. ACM, 2014.

[12] M. Dodds, M. Batty, and A. Gotsman. C/C++ causal cycles confound compositionality. *TinyToCS*, 2, 2013.

[13] S. Flur, K. E. Gray, C. Pulte, S. Sarkar, A. Sezgin, L. Maranget, W. Deacon, and P. Sewell. Modelling the ARMv8 architecture, operationally: Concurrency and ISA. In *POPL 2016*, pages 608–621. ACM, 2016.

[14] Intel. A formal specification of Intel Itanium processor family memory ordering, 2002. http://download.intel.com/design/Itanium/Downloads/25142901.pdf. [Online; accessed 14-November-2016].

[15] A. Jeffrey and J. Riely. On thin air reads: Towards an event structures model of relaxed memory. In *LICS 2016*, pages 759–767. ACM, 2016.

[16] J. Kang, C.-K. Hur, O. Lahav, V. Vafeiadis, and D. Dreyer. A promising semantics for relaxed-memory concurrency. In *POPL 2017*, pages 175–189. ACM, 2017.

[17] O. Lahav, N. Giannarakis, and V. Vafeiadis. Taming release-acquire consistency. In *POPL 2016*, pages 649–662. ACM, 2016.

[18] O. Lahav and V. Vafeiadis. Owicki-Gries reasoning for weak memory models. In *ICALP 2015*, pages 311–323. Springer, 2015.

[19] O. Lahav and V. Vafeiadis. Explaining relaxed memory models with program transformations. In *FM 2016*, pages 479–495. Springer, 2016.

[20] L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691, 1979.

[21] N. M. Lê, A. Pop, A. Cohen, and F. Zappa Nardelli. Correct and efficient work-stealing for weak memory models. In *PPoPP 2013*, pages 69–80. ACM, 2013.

[22] Y. A. Manerkar, C. Trippel, D. Lustig, M. Pellauer, and M. Martonosi. Counterexamples and proof loophole for the C/C++ to POWER and ARMv7 trailing-sync compiler mappings. *arXiv preprint arXiv:1611.01507*, 2016.

[23] L. Maranget, S. Sarkar, and P. Sewell. A tutorial introduction to the ARM and POWER relaxed memory models. http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test7.pdf, 2012.

[24] B. Norris and B. Demsky. CDSchecker: checking concurrent data structures written with C/C++ atomics. In *OOPSLA 2013*, pages 131–150. ACM, 2013.

[25] S. Owens, S. Sarkar, and P. Sewell. A better x86 memory model: x86-TSO. In *TPHOLs 2009*, pages 391–407. Springer-Verlag, 2009.

[26] J. Pichon-Pharabod and P. Sewell. A concurrency semantics for relaxed atomics that permits optimisation and avoids thin-air executions. In *POPL 2016*, pages 622–633. ACM, 2016.

[27] S. Sarkar, K. Memarian, S. Owens, M. Batty, P. Sewell, L. Maranget, J. Alglave, and D. Williams. Synchronising C/C++ and POWER. In *PLDI 2012*, pages 311–322. ACM, 2012.

[28] D. Shasha and M. Snir. Efficient and correct execution of parallel programs that share memory. *ACM Trans. Program. Lang. Syst.*, 10(2):282–312, Apr. 1988.

[29] A. Turon, V. Vafeiadis, and D. Dreyer. GPS: Navigating weak memory with ghosts, protocols, and separation. In *OOPSLA 2014*, pages 691–707. ACM, 2014.

[30] V. Vafeiadis, T. Balabonski, S. Chakraborty, R. Morisset, and F. Zappa Nardelli. Common compiler optimisations are invalid in the C11 memory model and what we can do about it. In *POPL 2015*, pages 209–220. ACM, 2015.

[31] V. Vafeiadis and C. Narayan. Relaxed separation logic: A program logic for C11 concurrency. In *OOPSLA 2013*, pages 867–884. ACM, 2013.

[32] J. Wickerson, M. Batty, T. Sorensen, and G. A. Constantinides. Automatically comparing memory consistency models. In *POPL 2017*, pages 190–204. ACM, 2017.

A. Further Examples

A.1 Failure of leading sync convention with SC fences

The following behavior is disallowed according to C11, but allowed by its compilation to Power.

$$\begin{array}{l} x :=_{\texttt{rlx}} 2 \\ \texttt{fence}_{\texttt{sc}} \\ b := y_{\texttt{rlx}} \; /\!/ 0 \end{array} \;\Big\|\; y :=_{\texttt{sc}} 1 \;\Big\|\; \begin{array}{l} d := y_{\texttt{acq}} \; /\!/ 1 \\ x :=_{\texttt{rel}} 1 \\ e := x_{\texttt{rlx}} \; /\!/ 2 \end{array} \;\Big\|\; f := x_{\texttt{sc}} \; /\!/ 1 \qquad \text{(Rsync+Rsc)}$$

Under C11, this behavior is forbidden. Consider the following execution (the initialization of x is omitted):

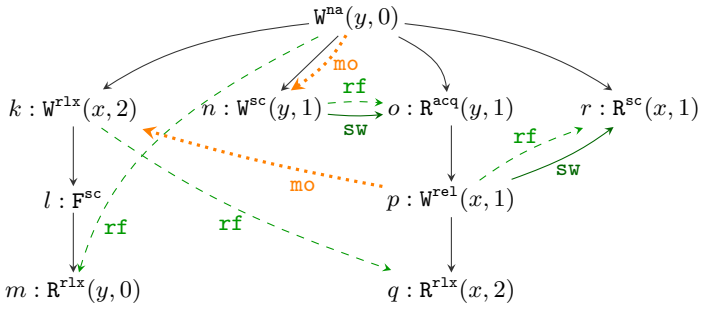

The rf and mo edges are forced because of read values and coherence. Now, the C11 conditions on SC fences require, in particular, that $[\texttt{F}^{\texttt{sc}}]; \texttt{sb}; \texttt{rb}; [\texttt{E}^{\texttt{sc}}] \subseteq \texttt{S}$ and $[\texttt{E}^{\texttt{sc}}]; \texttt{rb}; \texttt{sb}; [\texttt{F}^{\texttt{sc}}] \subseteq \texttt{S}$. Hence, we must have $\texttt{S}(l,n)$ (essentially because if we had $\texttt{S}(n,l)$, then m would have been reading from an overwritten write), as well as $\texttt{S}(r,l)$ (essentially because if we had $\texttt{S}(l,r)$, then r would have been reading from an mo-overwritten write before the fence). By transitivity, we thus have $\texttt{S}(r,n)$ which contradicts condition S1, which requires $\texttt{S}(n,r)$ because of the happens-before path via o and p.

The compilation to Power allows the behavior because the sync fences do not provide sufficient synchronization: again all but one sync fences are useless, as they are placed at the beginning and end of a thread. In fact, this example shows the unsoundness of compiliation of C11 to Power even for a compliation scheme that places a sync fence *both* before and after each SC access.

A.2 Unsoundness of compilation of C11 to ARMv8

$$x :=_{\texttt{sc}} 1 \;\Big\|\; \begin{array}{l} a := x_{\texttt{rlx}} \; /\!/ 1 \\ \texttt{fence}_{\texttt{acq}} \\ b := y_{\texttt{sc}} \; /\!/ 0 \end{array} \;\Big\|\; \begin{array}{l} y :=_{\texttt{sc}} 1 \\ c := x_{\texttt{sc}} \; /\!/ 0 \end{array} \qquad \text{(RWC+acq+sc)}$$

C11 disallows the annotated behavior of this program (we have hb from the write of x to the read of y; rb from the read of y to the write of y; hb from the write of y to the read of x; and finally rb from the read of x to the write of x).

Nevertheless, the compilation to ARMv8 (using its special load and store instructions for SC accesses) allows the behavior following the model in [13]:

$$\texttt{STLR}\ \#1, [x] \;\Big\|\; \begin{array}{l} \texttt{LDR}\ a, [x] \; /\!/ 1 \\ \texttt{DMB\ \ LD} \\ \texttt{LDAR}\ b, [y] \; /\!/ 0 \end{array} \;\Big\|\; \begin{array}{l} \texttt{STLR}\ \#1, [y] \\ \texttt{LDAR}\ c, [x] \; /\!/ 0 \end{array}$$

First, the store to x is committed and propagated to thread 2 (but not to thread 3). Then, the load of x in thread 2 is issued, satisfied and committed, the fence is committed, and the load of y is issued. In the storage subsystem, the issued load of y is propagated to thread 1, reordered with the store to x (as they originate from different threads), propagated to thread 3 and to the main memory, and satisfied with value 0. Then, thread 3 executes: the store to y is propagated to the main memory, and the load of x is issued and propagated to the main memory, satisfied with the value 0. Finally, the store of x propagates to the main memory.

A.3 Failure of read-after-read elimination

Let $\texttt{psc}_{\text{old}} = [\texttt{E}^{\texttt{sc}}]; (\texttt{sb} \cup \texttt{rf} \cup \texttt{mo} \cup \texttt{rb}); [\texttt{E}^{\texttt{sc}}]$. We present the executions showing the failure of read-after-read elimination when requiring acyclicity of $\texttt{psc}_{\text{old}}$ (see §2.1.1).

Consider the following program:

$$\begin{array}{l} y :=_{\texttt{sc}} 1 \\ x :=_{\texttt{rel}} 1 \end{array} \;\Big\|\; \begin{array}{l} a := x_{\texttt{sc}} \;/\!/ 1 \\ b := x_{\texttt{sc}} \;/\!/ 1 \end{array} \;\Big\|\; \begin{array}{l} c := x_{\texttt{rlx}} \;/\!/ 1 \\ x :=_{\texttt{sc}} 2 \\ d := y_{\texttt{sc}} \;/\!/ 0 \end{array} \qquad \text{(RRmerge)}$$

The annotated behavior is forbidden, but it will become allowed after replacing $b := x_{\texttt{sc}}$ by $b := a$. Indeed, the following execution is an execution of RRmerge yielding the result $a = b = c = 1 \wedge d = 0$.

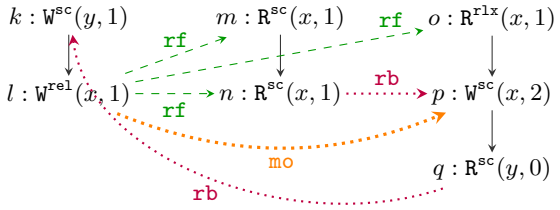

In this execution we have a $\texttt{psc}_{\text{old}}$ cycle (k, n, p, q, k). It is, however, consistent using our final psc relation ($\langle k, n \rangle \notin \texttt{psc}$).

Now, the following execution is an execution of the same RRmerge program, but after replacing $b := x_{\texttt{sc}}$ by $b := a$, again yielding the result $a = b = c = 1 \wedge d = 0$. This execution is consistent when requiring acyclicity of $\texttt{psc}_{\text{old}}$ (as well as with our final psc relation).

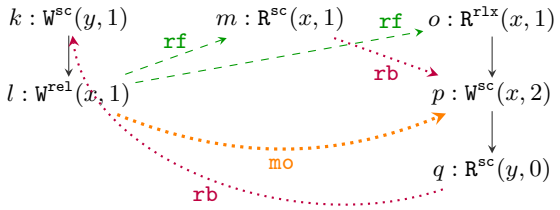

A.4 Failure of SC-read-after-non-SC-write elimination

$$\begin{array}{l} y :=_{\texttt{sc}} 2; \\ x :=_{\texttt{rlx}} 1; \\ a := x_{\texttt{sc}}; \;/\!/ 1 \\ b := x_{\texttt{rlx}}; \;/\!/ 2 \end{array} \;\Big\|\; \begin{array}{l} x :=_{\texttt{sc}} 2; \\ y :=_{\texttt{sc}} 1; \\ c := y_{\texttt{rlx}}; \;/\!/ 2 \end{array} \quad \rightsquigarrow \quad \begin{array}{l} y :=_{\texttt{sc}} 2; \\ x :=_{\texttt{rlx}} 1; \\ a := 1; \\ b := x_{\texttt{rlx}}; \;/\!/ 2 \end{array} \;\Big\|\; \begin{array}{l} x :=_{\texttt{sc}} 2; \\ y :=_{\texttt{sc}} 1; \\ c := y_{\texttt{rlx}}; \;/\!/ 2 \end{array}$$

The annotated behavior is allowed under RC11 for the target, but not for the source. The same applies to the model of Batty *et al.* [5].

B. Programs to Executions: Receptiveness Assumption

To carry out the compilation correctness proof, we need to record syntactic dependencies between instructions, as in the Power model. (This is only needed if one is interested in the NO-THIN-AIR condition; compilation correctness for weakRC11 may completely ignore this extension.) Dependencies are classified into data, address and control dependencies. Accordingly, we extend the definition of an execution (see §3.1), with additional relations data, addr and ctrl. We use deps to denote the union of the three relations. We require data, addr and ctrl to satisfy the following:

1. $\texttt{data} \subseteq \texttt{R} \times \texttt{W}$.
2. $\texttt{addr} \subseteq \texttt{R} \times (\texttt{R} \cup \texttt{W})$.
3. $\texttt{ctrl} \subseteq \texttt{R} \times \texttt{E}$.
4. $\texttt{ctrl}; \texttt{sb} \subseteq \texttt{ctrl}$.
5. $\texttt{rmw} \subseteq \texttt{deps}$.

The dependency relations are calculated from the program syntax, together with the generation of program's execution (like in Power), and the construction ensures that the above properties hold. Moreover, the construction of executions from programs provides us with the following *receptiveness* property:

Definition B.1. A function $lab' : \mathsf{Event} \rightharpoonup \mathsf{Label}$ is called a *reevaluation* of $lab : \mathsf{Event} \rightharpoonup \mathsf{Label}$ if for every event a, the label $lab'(a)$ is identical to $lab(a)$, except possibly for read/written value.

Notation B.1. Given an execution G and a reevaluation lab of $G.\texttt{lab}$, $lab(G)$ denotes the execution G' given by: $G'.\texttt{lab} = lab$, $G'.\texttt{rf} = \emptyset$, and $G'.\texttt{c} = G.\texttt{c}$ for every $\texttt{c} \in \{\texttt{E}, \texttt{sb}, \texttt{rmw}, \texttt{mo}, \texttt{data}, \texttt{addr}, \texttt{ctrl}\}$.

Assumption B.1 (receptiveness)**.** Let G be an execution of a program P. Let $a \in \texttt{R}$, and suppose that $a \notin dom(\texttt{deps}^*; (\texttt{ctrl} \cup \texttt{addr}))$. For every $v \in \mathsf{Val}$, there exists a reevaluation lab of $G.\texttt{lab}$ such that:

- $lab(G)$ is an execution of P.
- $lab(G).\texttt{val}_\texttt{r}(a) = v$.
- $lab(b) = G.\texttt{lab}(b)$ whenever $\langle a, b\rangle \notin G.\texttt{deps}^+$.

Note that a more basic receptiveness property follows from this assumption: if $a \notin dom(\texttt{sb})$ then for every $v \in \mathsf{Val}$, we have that $lab(G)$ is an execution of P, for the reevaluation lab of $G.\texttt{lab}$ that sets the read value of a to v, and otherwise is identical to $G.\texttt{lab}$.

In addition, we assume that the set of executions of a program is *prefix-closed*:

Notation B.2. Given an execution G and a set $E \subseteq \texttt{E}$ that is downwards closed w.r.t. $\texttt{sb}$ (*i.e.,* $a \in E$ whenever $\langle a, b\rangle \in \texttt{sb}$ for some $b \in E$), and contains at least all the initialization events, the *restriction* of G to E, denoted $G|_E$, is the execution G' given by $G'.\texttt{E} = E$, $G'.\texttt{lab} = G.\texttt{lab}|_E$, and $G'.\texttt{c} = [E]; G.\texttt{c}; [E]$ for $\texttt{c} \in \{\texttt{sb}, \texttt{rmw}, \texttt{rf}, \texttt{mo}, \texttt{data}, \texttt{addr}, \texttt{ctrl}\}$.

Assumption B.2 (prefix-closed executions)**.** Let G be an execution of a program P, and let E be a subset of $\texttt{E}$ that is downwards closed w.r.t. $\texttt{sb}$, and contains at least all the initialization events. Then, $G|_E$ is an execution of P.

C. Properties of RC11

In this section, we present some basic properties of the derived relations $\texttt{eco}$, $\texttt{sw}$, $\texttt{hb}$ and of RC11-consistent executions. We omit some of the proofs that straightforwardly follow from our definitions. For the rest of this section, consider an arbitrary execution G.

Proposition C.1. $\texttt{eco}$ *is a strict partial order.*

Proposition C.2. *Suppose that* $[\texttt{W}]; \texttt{sb}|_{\texttt{loc}}; [\texttt{W}] \subseteq \texttt{mo}$ *and* $\texttt{rmw} \subseteq \texttt{rb}$*. Then, the following hold:*

1. $\texttt{rs} \subseteq \texttt{eco}^?$.
2. $[\texttt{W}]; \texttt{sw}; [\texttt{R}] \subseteq \texttt{eco}$.
3. $[\texttt{W}]; \texttt{sw}; [\texttt{F}] \subseteq \texttt{eco}; \texttt{sb}$.
4. $\texttt{eco}; \texttt{hb} \subseteq \texttt{eco} \cup \texttt{eco}; (\texttt{sb} \setminus \texttt{rmw}); \texttt{hb}^?$.

Proof.

1. Let $\langle a, b\rangle \in \texttt{rs}$. Then, by definition, $\langle a, b\rangle \in [\texttt{W}]; \texttt{sb}|^?_{\texttt{loc}}; [\texttt{W}^{\sqsupseteq \texttt{rlx}}]; (\texttt{rf}; \texttt{rmw})^*$. Since $[\texttt{W}]; \texttt{sb}|_{\texttt{loc}}; [\texttt{W}] \subseteq \texttt{mo}$ and $\texttt{rmw} \subseteq \texttt{rb}$, we have $\langle a, b\rangle \in \texttt{eco}^*$. Since $\texttt{eco}$ is transitive, we have $\langle a, b\rangle \in \texttt{eco}^?$.
2. Let $\langle a, b\rangle \in [\texttt{W}]; \texttt{sw}; [\texttt{R}]$. Then, by definition, we have $\langle a, b\rangle \in \texttt{rs}; \texttt{rf}$. Using the previous item, we obtain that $\langle a, b\rangle \in \texttt{eco}^?; \texttt{eco} \subseteq \texttt{eco}$.
3. Let $\langle a, b\rangle \in [\texttt{W}]; \texttt{sw}; [\texttt{F}]$. Then, by definition, we have $\langle a, b\rangle \in \texttt{rs}; \texttt{rf}; \texttt{sb}$. Using the first item, we obtain that $\langle a, b\rangle \in \texttt{eco}^?; \texttt{eco}; \texttt{sb} \subseteq \texttt{eco}; \texttt{sb}$.
4. Let $\langle a, c\rangle \in \texttt{eco}; \texttt{hb}$, and let $b \in \texttt{E}$ be an $\texttt{eco}$-maximal event satisfying $\langle a, b\rangle \in \texttt{eco}$, and $\langle b, c\rangle \in \texttt{hb}^?$. If $b = c$ then $\langle a, c\rangle \in \texttt{eco}$, and we are done. Otherwise, the maximality of b ensures that $\langle b, b'\rangle \in \texttt{sb} \setminus \texttt{sw}$ and $\langle b', c\rangle \in \texttt{hb}^?$ for some $b' \in \texttt{E}$. Since $\texttt{rmw} \subseteq \texttt{rb} \subseteq \texttt{eco}$, it follows that $\langle a, c\rangle \in \texttt{eco}; (\texttt{sb} \setminus \texttt{rmw}); \texttt{hb}^?$. □

Lemma C.1 (Read at end). *Let* $a \in \mathtt{R} \setminus dom(\mathtt{sb})$. *Suppose that* $G' = G|_{G.\mathtt{E}\setminus\{a\}}$ *is* RC11*-consistent. Then, there exists an event* $b \in G'.\mathtt{W}$ *such that the execution* G'' *given by* $G''.\mathtt{c} = G.\mathtt{c}$ *for every* $\mathtt{c} \in \{\mathtt{E}, \mathtt{sb}, \mathtt{rmw}, \mathtt{data}, \mathtt{addr}, \mathtt{ctrl}, \mathtt{mo}\}$, $G''.\mathtt{lab} = G'.\mathtt{lab} \cup \{a \mapsto \mathtt{R}^{\mathtt{mod}(a)}(\mathtt{loc}(a), \mathtt{val}_{\mathtt{w}}(b))\}$, *and* $G''.\mathtt{rf} = G'.\mathtt{rf} \cup \{\langle b, a\rangle\}$ *is* RC11*-consistent.*

Proof. Take b to be the $\mathtt{mo}$-maximal event in $G.\mathtt{W}_{\mathtt{loc}(a)}$. It is straightforward to show that G'', as defined in the statement, is RC11-consistent. □

Proposition C.3. *Let* $a \in \mathtt{W}^{\sqsubseteq \mathtt{rlx}} \setminus dom(\mathtt{rf})$. *Let* $G' = G|_{G.\mathtt{E}\setminus\{a\}}$. *Then,* $[G'.\mathtt{E}]; G.\mathtt{hb}; [G'.\mathtt{E}] = G'.\mathtt{hb}$.

Proposition C.4. *Let* G' *be any execution obtained from* G *by possibly changing the value read at some* $a \in \mathtt{R}^{\mathtt{na}}$, *and the source of the* $\mathtt{rf}$ *edge entering the event* a. *Then,* $G'.\mathtt{hb} = G.\mathtt{hb}$.

Proposition C.5. *Let* G' *be an execution, such that* $G'.\mathtt{E} = G.\mathtt{E} \uplus \{a\}$ *for some event* a. *Suppose that* $a \in G'.\mathtt{R}^{\mathtt{na}}$, $G.\mathtt{sb} \subseteq G'.\mathtt{sb}$, $G.\mathtt{lab} \subseteq G'.\mathtt{lab}$, $G.\mathtt{rmw} = G'.\mathtt{rmw}$, *and* $G'.\mathtt{rf} = G.\mathtt{rf} \cup \{\langle b, a\rangle\}$ *for some* $b \in G.\mathtt{E}$. *Then,* $[G.\mathtt{E}]; G'.\mathtt{hb}; [G.\mathtt{E}] = G.\mathtt{hb}$.

D. The $\mathsf{RC}_{\mathtt{na}}$ Model

In this section we present a variant of RC11, which has a smaller $\mathtt{psc}_{\mathrm{base}}$ relation, and is useful in our correctness of compilation proofs. It is based on the following additional derived relations:

$$
\begin{aligned}
\mathtt{rb}^{\mathtt{na}} &\triangleq [\mathtt{R}^{\mathtt{na}}]; \mathtt{rb} \\
\mathtt{rb}^{\neq \mathtt{na}} &\triangleq \mathtt{rb} \setminus \mathtt{rb}^{\mathtt{na}} \\
\mathtt{eco}^{\neq \mathtt{na}} &\triangleq \mathtt{rf} \cup (\mathtt{mo} \cup \mathtt{rb}^{\neq \mathtt{na}}); \mathtt{rf}^{?} \\
\mathtt{scb}^{\neq \mathtt{na}} &\triangleq \mathtt{sb} \cup \mathtt{sb}|_{\neq \mathtt{loc}}; \mathtt{hb}; \mathtt{sb}|_{\neq \mathtt{loc}} \cup \mathtt{hb}|_{\mathtt{loc}} \cup \mathtt{mo} \cup \mathtt{rb}^{\neq \mathtt{na}} \\
\mathtt{psc}_{\mathrm{base}}^{\neq \mathtt{na}} &\triangleq \left([\mathtt{E}^{\mathtt{sc}}] \cup [\mathtt{F}^{\mathtt{sc}}]; \mathtt{hb}^{?}\right); \mathtt{scb}^{\neq \mathtt{na}}; \left([\mathtt{E}^{\mathtt{sc}}] \cup \mathtt{hb}^{?}; [\mathtt{F}^{\mathtt{sc}}]\right) \\
\mathtt{psc}_{\mathtt{F}}^{\neq \mathtt{na}} &\triangleq [\mathtt{F}^{\mathtt{sc}}]; (\mathtt{hb} \cup \mathtt{hb}; \mathtt{eco}^{\neq \mathtt{na}}; \mathtt{hb}); [\mathtt{F}^{\mathtt{sc}}]
\end{aligned}
$$

Proposition D.1. *If* $\mathtt{rb}^{\mathtt{na}} \subseteq \mathtt{hb}$ *then* $\mathtt{psc}_{\mathrm{base}} = \mathtt{psc}_{\mathrm{base}}^{\neq \mathtt{na}}$ *and* $\mathtt{psc}_{\mathtt{F}} = \mathtt{psc}_{\mathtt{F}}^{\neq \mathtt{na}}$.

Proof. Note that $\mathtt{rb}^{\mathtt{na}} \subseteq \mathtt{hb}$ implies that $\mathtt{rb}^{\mathtt{na}} \subseteq \mathtt{hb}|_{\mathtt{loc}}$, and $\mathtt{hb}; \mathtt{rb}^{\mathtt{na}}; \mathtt{rf}^{?}; \mathtt{hb} \subseteq \mathtt{hb} \cup \mathtt{hb}; \mathtt{eco}; \mathtt{hb}$. In addition, we have $\mathtt{psc}_{\mathrm{base}} \setminus \mathtt{psc}_{\mathrm{base}}^{\neq \mathtt{na}} \subseteq [\mathtt{F}^{\mathtt{sc}}]; \mathtt{hb}; \mathtt{rb}^{\mathtt{na}}; ([\mathtt{E}^{\mathtt{sc}}] \cup \mathtt{hb}; [\mathtt{F}^{\mathtt{sc}}])$, and $\mathtt{psc}_{\mathtt{F}} \setminus \mathtt{psc}_{\mathtt{F}}^{\neq \mathtt{na}} \subseteq [\mathtt{F}^{\mathtt{sc}}]; \mathtt{hb}; \mathtt{rb}^{\mathtt{na}}; \mathtt{rf}^{?}; \mathtt{hb}; [\mathtt{F}^{\mathtt{sc}}]$. Thus, this claim immediately follows from our definitions. □

We call an execution $\mathsf{RC}_{\mathtt{na}}$*-consistent* if it satisfies all conditions of Def. 1, except possibly for SC, and $\mathtt{psc}_{\mathrm{base}}^{\neq \mathtt{na}} \cup \mathtt{psc}_{\mathtt{F}}^{\neq \mathtt{na}}$ is acyclic.

Lemma D.1. *Let* G *be an* $\mathsf{RC}_{\mathtt{na}}$*-consistent execution of a program* P. *Then, either* G *is* RC11*-consistent, or* P *has undefined behavior under* RC11.

Proof. If $\mathtt{rb}^{\mathtt{na}} \subseteq \mathtt{hb}$, then, by Prop. D.1, G is RC11-consistent. Suppose otherwise. We show that P has undefined behavior under RC11. Let $a_1, \ldots, a_n$ be an enumeration of $\mathtt{E}$ that respects $\mathtt{sb} \cup \mathtt{rf}$ (that is, $i < j$ whenever $\langle a_i, a_j\rangle \in \mathtt{sb} \cup \mathtt{rf}$). For every $1 \leq i \leq n$, let $E_i = \mathtt{E}_0 \cup \{a_1, \ldots, a_i\}$ and $G_i = G|_{E_i}$. Let k be the minimal index such that $G_k.\mathtt{rb}^{\mathtt{na}} \not\subseteq G_k.\mathtt{hb}$. Then, by Prop. D.1, $G_{k-1}.\mathtt{psc}_{\mathrm{base}} \cup G_{k-1}.\mathtt{psc}_{\mathtt{F}} = G_{k-1}.\mathtt{psc}_{\mathrm{base}}^{\neq \mathtt{na}} \cup G_{k-1}.\mathtt{psc}_{\mathtt{F}}^{\neq \mathtt{na}}$ is acyclic, and so G_{k-1} is RC11-consistent. Let $\langle a_{\mathtt{R}}, a_{\mathtt{W}}\rangle \in G_k.\mathtt{rb}^{\mathtt{na}} \setminus G_k.\mathtt{hb}$. Then, we must have $a_k \in \{a_{\mathtt{R}}, a_{\mathtt{W}}\}$. Note also that $\langle a_{\mathtt{W}}, a_{\mathtt{R}}\rangle \notin G_k.\mathtt{hb}$ since G_k satisfies COHERENCE.

Now, if G_k is RC11-consistent, then we are done (it is a racy execution of P). Suppose otherwise. We show that $a_k \neq a_{\mathtt{W}}$. Indeed, otherwise, since G_k is $\mathsf{RC}_{\mathtt{na}}$-consistent but not RC11-consistent, and G_{k-1} is RC11-consistent, it must be the case that $\mathtt{mod}(a_k) = \mathtt{sc}$, and there exist $b, f \in E_{k-1}$ such that:

- $\langle a_k, b\rangle \in G_k.\mathtt{mo}; (G_{k-1}.\mathtt{hb}; [\mathtt{F}])^{?}; [G_{k-1}.\mathtt{E}^{\mathtt{sc}}]$
- $\langle b, f\rangle \in (G_{k-1}.\mathtt{psc}_{\mathrm{base}} \cup G_{k-1}.\mathtt{psc}_{\mathtt{F}})^{*}; [\mathtt{F}^{\mathtt{sc}}]$
- $\langle f, a_k\rangle \in G_{k-1}.\mathtt{hb}; G_k.\mathtt{rb}^{\mathtt{na}}$

Now, since we have $[E_{k-1}]; G_k.\mathtt{rb}; G_k.\mathtt{mo}; [E_{k-1}] \subseteq G_{k-1}.\mathtt{rb}$, it follows that $\langle f, b\rangle \in G_{k-1}.\mathtt{psc}_{\mathrm{base}}$. This, however, contradicts the fact that G_{k-1} is RC11-consistent.

Therefore, we have $a_k = a_{\mathtt{R}}$. Let $x = G.\mathtt{loc}(a_k)$. By Lemma C.1, there exists an event $b \in G_{k-1}.\mathtt{W}_x$ such that the execution G' given by $G'.\mathtt{c} = G_k.\mathtt{c}$ for every $\mathtt{c} \in \{\mathtt{E}, \mathtt{sb}, \mathtt{rmw}, \mathtt{data}, \mathtt{addr}, \mathtt{ctrl}, \mathtt{mo}\}$, $G'.\mathtt{lab} = G_{k-1}.\mathtt{lab} \cup \{a_k \mapsto \mathtt{R}^{\mathtt{na}}(x, \mathtt{val}_{\mathtt{w}}(b))\}$, and $G'.\mathtt{rf} = G_{k-1}.\mathtt{rf} \cup \{\langle b, a_k\rangle\}$ is RC11-consistent. By Assumption B.1, G' is an execution of P. In addition, we have $\langle a_{\mathtt{W}}, a_k\rangle \notin G'.\mathtt{hb}$ (since $\langle a_{\mathtt{W}}, a_k\rangle \notin G_k.\mathtt{hb}$ and $G'.\mathtt{hb} = G_k.\mathtt{hb}$ by Prop. C.4), and so G' is racy. Hence, P has undefined behavior under RC11. □

Next, we prove some lemmas that allow us (under some restrictions) to add a memory access inside a given execution. In what follows, we take G to be an arbitrary execution.

Proposition D.2. *If $a \notin dom(\mathtt{sb}^?; [\mathtt{E}^{\sqsupseteq \mathtt{rel}}])$, then for every $b \in \mathtt{E}$, we have $\langle a, b\rangle \in \mathtt{hb}$ iff $\langle a, b\rangle \in \mathtt{sb}$.*

Proof. The assumption that $a \notin dom(\mathtt{sb}^?; [\mathtt{E}^{\sqsupseteq \mathtt{rel}}])$ ensures that $a \notin dom(\mathtt{sb}^?; \mathtt{sw})$, and so we have $\langle a, b\rangle \in \mathtt{hb}$ iff $\langle a, b\rangle \in \mathtt{sb}$. □

Lemma D.2 (Add write)**.** *Let $a \in \mathtt{W} \setminus (dom(\mathtt{sb}^?; [\mathtt{E}^{\sqsupseteq \mathtt{rel}}]) \cup \mathtt{At})$. Suppose that $G' = G|_{G.\mathtt{E}\setminus\{a\}}$ is $\mathsf{RC}_{\mathtt{na}}$-consistent. Let $x = \mathtt{loc}(a)$, and suppose that $\langle a, b\rangle \in \mathtt{sb}; [\mathtt{R}_x]$ implies $\langle a, b\rangle \in \mathtt{sb}; [\mathtt{W}_x]; \mathtt{sb}$. Then, there exists a relation $T \subseteq G.\mathtt{W}_x \times G.\mathtt{W}_x$ such that the execution G'' given by $G''.\mathtt{c} = G.\mathtt{c}$ for every $\mathtt{c} \in \{\mathtt{E}, \mathtt{lab}, \mathtt{sb}, \mathtt{rmw}, \mathtt{data}, \mathtt{addr}, \mathtt{ctrl}\}$, $G''.\mathtt{rf} = G'.\mathtt{rf}$, and $G''.\mathtt{mo} = G'.\mathtt{mo} \cup T$ is $\mathsf{RC}_{\mathtt{na}}$-consistent.*

Proof. Let $C = \{c \in G'.\mathtt{W}_x \mid \langle a, c\rangle \in G.\mathtt{sb}; G'.\mathtt{mo}^?\}$, and take $T = (\{a\} \times C) \cup ((G'.\mathtt{W}_x \setminus C) \times \{a\})$. It is straightforward to show that G'', as defined in the statement, is $\mathsf{RC}_{\mathtt{na}}$-consistent. In particular, we have $G''.\mathtt{psc}_{\mathrm{base}}^{\neq \mathtt{na}} = G'.\mathtt{psc}_{\mathrm{base}}^{\neq \mathtt{na}}$ and $G''.\mathtt{psc}_{\mathtt{F}}^{\neq \mathtt{na}} = G'.\mathtt{psc}_{\mathtt{F}}^{\neq \mathtt{na}}$. □

Lemma D.3 (Add rmw write)**.** *Suppose that $\mathtt{rmw}^{-1}; \mathtt{rf}^{-1}; \mathtt{rf}; \mathtt{rmw} \subseteq [G.\mathtt{E}]$. Let $a \in (\mathtt{W} \cap \mathtt{At}) \setminus dom(\mathtt{sb}^?; [\mathtt{E}^{\sqsupseteq \mathtt{rel}}])$. Suppose that $G' = G|_{G.\mathtt{E}\setminus\{a\}}$ is $\mathsf{RC}_{\mathtt{na}}$-consistent. Let $x = \mathtt{loc}(a)$, and suppose that $\langle a, b\rangle \in \mathtt{sb}; [\mathtt{R}_x]$ implies $\langle a, b\rangle \in \mathtt{sb}; [\mathtt{W}_x]; \mathtt{sb}$. Then, there exists a relation $T \subseteq G.\mathtt{W}_x \times G.\mathtt{W}_x$ such that the execution G'' given by $G''.\mathtt{c} = G.\mathtt{c}$ for every $\mathtt{c} \in \{\mathtt{E}, \mathtt{lab}, \mathtt{sb}, \mathtt{rmw}, \mathtt{data}, \mathtt{addr}, \mathtt{ctrl}\}$, $G''.\mathtt{rf} = G'.\mathtt{rf}$, and $G''.\mathtt{mo} = G'.\mathtt{mo} \cup T$ is $\mathsf{RC}_{\mathtt{na}}$-consistent.*

Proof. Let $b, d \in G'.\mathtt{E}$ such that $\langle b, a\rangle \in G.\mathtt{rmw}$ and $\langle d, b\rangle \in G'.\mathtt{rf}$. Let $C = \{c \in G'.\mathtt{W}_x \mid \langle d, c\rangle \in G'.\mathtt{mo}\}$, and take $T = (\{a\} \times C) \cup ((G'.\mathtt{W}_x \setminus C) \times \{a\})$. It is straightforward to show that G'', as defined in the statement, is $\mathsf{RC}_{\mathtt{na}}$-consistent. □

Lemma D.4 (Add non-atomic read)**.** *Let $a \in \mathtt{R}^{\mathtt{na}} \setminus dom(\mathtt{sb}; [\mathtt{E}^{\sqsupseteq \mathtt{rel}}])$. Suppose that $G' = G|_{G.\mathtt{E}\setminus\{a\}}$ is $\mathsf{RC}_{\mathtt{na}}$-consistent. Then, there exists an event $b \in G'.\mathtt{W}$ such that the execution G'' given by $G''.\mathtt{E} = G.\mathtt{E}$, $G''.\mathtt{lab} = G'.\mathtt{lab} \cup \{a \mapsto \mathtt{R}^{\mathtt{na}}(\mathtt{loc}(a), \mathtt{val}_{\mathtt{w}}(b))\}$, $G''.\mathtt{c} = G.\mathtt{c}$ for every $\mathtt{c} \in \{\mathtt{sb}, \mathtt{rmw}, \mathtt{data}, \mathtt{addr}, \mathtt{ctrl}\}$, $G''.\mathtt{rf} = G'.\mathtt{rf} \cup \{\langle b, a\rangle\}$, and $G''.\mathtt{mo} = G.\mathtt{mo}$ is $\mathsf{RC}_{\mathtt{na}}$-consistent.*

Proof. Let $x = \mathtt{loc}(a)$. Let $B = \{b \in G.\mathtt{W}_x \mid \langle b, a\rangle \in G.\mathtt{rf}^?; G.\mathtt{hb}\}$, and take b be the $\mathtt{mo}$-maximal event in B. It is straightforward to show that G'', as defined in the statement, is $\mathsf{RC}_{\mathtt{na}}$-consistent. □

E. Proof of Global Transformation of SC accesses

In this section we prove the soundness of a global program transformation that either adds an SC fence before every SC access, or adds an SC fence after every SC access, and then replaces all SC accesses by release/acquire ones. This will allow us later to prove the correctness of compilation only for programs that do not contain any SC accesses.

We use the following additional notation:

$$\mathtt{sb}' \triangleq \mathtt{sb} \setminus \mathtt{rmw}$$

Lemma E.1. *Let G be an execution satisfying all conditions of Def. 1, except possibly for* SC*. Suppose that $[\mathtt{RW}^{\mathtt{sc}}]; (\mathtt{sb}' \cup \mathtt{sb}'; \mathtt{hb}; \mathtt{sb}'); [\mathtt{RW}^{\mathtt{sc}}] \subseteq \mathtt{hb}; [\mathtt{F}^{\mathtt{sc}}]; \mathtt{hb}$. Let $T = \mathtt{sb} \cup \mathtt{sb}'; \mathtt{hb}; \mathtt{sb}' \cup \mathtt{eco}$. Then:*

$$[\mathtt{F}^{\mathtt{sc}}]; \mathtt{hb}; \mathtt{eco}^?; ([\mathtt{RW}^{\mathtt{sc}}]; T; [\mathtt{RW}^{\mathtt{sc}}])^*; \mathtt{eco}^?; \mathtt{hb}; [\mathtt{F}^{\mathtt{sc}}] \subseteq \mathtt{psc}_{\mathtt{F}}^{+}.$$

Proof. We show by induction on n, that $[\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?;([\mathtt{RW^{sc}}];T;[\mathtt{RW^{sc}}])^n;\mathtt{eco}^?;\mathtt{hb};[\mathtt{F^{sc}}] \subseteq \mathtt{psc}_\mathtt{F}^+$ for every $n \geq 0$. For $n = 0$, the claim holds since $\mathtt{eco}^?;\mathtt{eco}^? \subseteq \mathtt{eco}^?$, and $[\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?;\mathtt{hb};[\mathtt{F^{sc}}] \subseteq \mathtt{psc}_\mathtt{F}$. Suppose now that $[\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?;([\mathtt{RW^{sc}}];T;[\mathtt{RW^{sc}}])^{n-1};\mathtt{eco}^?;\mathtt{hb};[\mathtt{F^{sc}}] \subseteq \mathtt{psc}_\mathtt{F}^+$, and let $R = [\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?;([\mathtt{RW^{sc}}];T;[\mathtt{RW^{sc}}])^n;\mathtt{eco}^?;\mathtt{hb};[\mathtt{F^{sc}}]$. Expanding the definition of T (keeping in mind that $\mathtt{rmw} \subseteq \mathtt{eco}$) we have $R \subseteq R_1 \cup R_2$, where:

$$R_1 = [\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?;([\mathtt{RW^{sc}}];T;[\mathtt{RW^{sc}}])^{n-1};[\mathtt{RW^{sc}}];(\mathtt{sb}' \cup \mathtt{sb}';\mathtt{hb};\mathtt{sb}');[\mathtt{RW^{sc}}];\mathtt{eco}^?;\mathtt{hb};[\mathtt{F^{sc}}],$$

$$R_2 = [\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?;([\mathtt{RW^{sc}}];T;[\mathtt{RW^{sc}}])^{n-1};\mathtt{eco};\mathtt{eco}^?;\mathtt{hb};[\mathtt{F^{sc}}].$$

Since $\mathtt{eco};\mathtt{eco}^? \subseteq \mathtt{eco}$, by the induction hypothesis, we have $R_2 \subseteq \mathtt{psc}_\mathtt{F}^+$. In addition, our assumption entails that R_1 is contained in

$$R_1' = [\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?;([\mathtt{RW^{sc}}];T;[\mathtt{RW^{sc}}])^{n-1};\mathtt{hb};[\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?;\mathtt{hb};[\mathtt{F^{sc}}],$$

which, in turn, using the induction hypothesis is also contained in $\mathtt{psc}_\mathtt{F}^+$. □

Lemma E.2. *Let G be an execution satisfying all conditions of Def. 1, except possibly for* SC*. Suppose that $[\mathtt{RW^{sc}}];(\mathtt{sb}' \cup \mathtt{sb}';\mathtt{hb};\mathtt{sb}');[\mathtt{RW^{sc}}] \subseteq \mathtt{hb};[\mathtt{F^{sc}}];\mathtt{hb}$. Then, if $\mathtt{psc}_\mathtt{F}$ is acyclic, then so is $\mathtt{psc}_{\mathrm{base}} \cup \mathtt{psc}_\mathtt{F}$.*

Proof. Contrapositively, suppose that $\mathtt{psc}_{\mathrm{base}} \cup \mathtt{psc}_\mathtt{F}$ is cyclic. Then, by definition, the union of the following relations is cyclic:

- $A_1 = [\mathtt{RW^{sc}}];\mathtt{scb};[\mathtt{RW^{sc}}]$
- $A_2 = [\mathtt{F^{sc}}];(\mathtt{hb} \cup \mathtt{hb};\mathtt{eco};\mathtt{hb});[\mathtt{F^{sc}}]$
- $A_3 = [\mathtt{RW^{sc}}];\mathtt{scb};\mathtt{hb}^?;[\mathtt{F^{sc}}]$
- $A_4 = [\mathtt{F^{sc}}];\mathtt{hb}^?;\mathtt{scb};[\mathtt{RW^{sc}}]$

Consider first the case that A_1 is cyclic. Then, since $\mathtt{rmw} \subseteq \mathtt{eco}$ and $\mathtt{hb}|_{\mathtt{loc}} \subseteq \mathtt{eco} \cup \mathtt{sb}';\mathtt{hb};\mathtt{sb}'$, the relation $[\mathtt{RW^{sc}}];(\mathtt{sb}' \cup \mathtt{sb}';\mathtt{hb};\mathtt{sb}');[\mathtt{RW^{sc}}] \cup \mathtt{eco}$ is cyclic. Our assumption on G entails that $\mathtt{hb};[\mathtt{F^{sc}}];\mathtt{hb} \cup \mathtt{eco}$ is cyclic. Since both $\mathtt{hb};[\mathtt{F^{sc}}];\mathtt{hb}$ and $\mathtt{eco}$ are transitive and irreflexive, we obtain that $\mathtt{hb};[\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}$ is cyclic, which in turn implies that $[\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco};\mathtt{hb};[\mathtt{F^{sc}}] \subseteq \mathtt{psc}_\mathtt{F}$ is cyclic.

Now, consider the case that A_1 is acyclic. Let $T = \mathtt{sb} \cup \mathtt{sb}';\mathtt{hb};\mathtt{sb}' \cup \mathtt{eco}$. It is easy to see that $\mathtt{scb} \subseteq T$ (since we have $\mathtt{rmw} \subseteq \mathtt{eco}$ and $\mathtt{hb}|_{\mathtt{loc}} \subseteq \mathtt{eco} \cup \mathtt{sb}';\mathtt{hb};\mathtt{sb}'$). Then, the union of $\mathtt{psc}_\mathtt{F}$ and the following relation must be cyclic:

$$B = [\mathtt{F^{sc}}];\mathtt{hb}^?;\mathtt{scb};([\mathtt{RW^{sc}}];T;[\mathtt{RW^{sc}}])^*;\mathtt{scb};\mathtt{hb}^?;[\mathtt{F^{sc}}]$$

Now, we have $[\mathtt{F^{sc}}];\mathtt{hb}^?;\mathtt{scb} \subseteq [\mathtt{F^{sc}}];\mathtt{hb};\mathtt{eco}^?$ and $\mathtt{scb};\mathtt{hb}^?;[\mathtt{F^{sc}}] \subseteq \mathtt{eco}^?;\mathtt{hb};[\mathtt{F^{sc}}]$. By Lemma E.1, it follows that $B \subseteq \mathtt{psc}_\mathtt{F}^+$, and so $\mathtt{psc}_\mathtt{F}$ is cyclic. □

Lemma E.3. *Let G be an* RC11*-consistent execution without any SC accesses. Let $A \subseteq \mathtt{R}^{\sqsupseteq \mathtt{acq}} \cup \mathtt{W}^{\sqsupseteq \mathtt{rel}}$, such that $[A];(\mathtt{sb}' \cup \mathtt{sb}';\mathtt{hb};\mathtt{sb}');[A] \subseteq \mathtt{hb};[\mathtt{F^{sc}}];\mathtt{hb}$, and $[A];\mathtt{rmw} = \mathtt{rmw};[A]$. Then, the execution G' obtained from G by changing all modes of events in A to* sc *is* RC11*-consistent.*

Proof. The only constraint that is affected by such modification is SC. Now, in G' we have $[G'.\mathtt{RW^{sc}}];(G'.\mathtt{sb}' \cup G'.\mathtt{sb}';G'.\mathtt{hb};G'.\mathtt{sb}');[G'.\mathtt{RW^{sc}}] \subseteq G'.\mathtt{hb};[G'.\mathtt{F^{sc}}];G'.\mathtt{hb}$, and by Lemma E.2 it suffices to show that $G'.\mathtt{psc}_\mathtt{F}$ is acyclic. This follows from the fact that G satisfies SC, since $G'.\mathtt{psc}_\mathtt{F} = G.\mathtt{psc}_\mathtt{F}$. □

F. Properties of the Power and ARMv7 Models

In this appendix we provide the full definition of preserved program order (ppo) used by Power and ARMv7, and prove various properties of these models that are needed in our compilation correctness proof.

Notation F.1. For every relation $\mathtt{c}$ (e.g., rf, mo, etc.), we denote by $\mathtt{ci}$ and $\mathtt{ce}$ (internal $\mathtt{c}$ and external $\mathtt{c}$) its thread-internal and thread-external restrictions. Formally, $\mathtt{ci} = \mathtt{c} \cap \mathtt{sb}$ and $\mathtt{ce} = \mathtt{c} \setminus \mathtt{sb}$.

F.1 Preserved Program Order

ppo is defined based on the four dependencies — data, addr, ctrl, $\mathtt{ctrl}_{\mathtt{isync}}$ — that satisfy the following properties:

1. $\texttt{data} \subseteq \texttt{R} \times \texttt{W}$.
2. $\texttt{addr} \subseteq \texttt{R} \times (\texttt{R} \cup \texttt{W})$.
3. $\texttt{ctrl}_{\texttt{isync}} \subseteq \texttt{ctrl} \subseteq \texttt{R} \times \texttt{E}$.
4. $\texttt{ctrl}; \texttt{sb} \subseteq \texttt{ctrl}$.
5. $\texttt{ctrl}_{\texttt{isync}}; \texttt{sb} \subseteq \texttt{ctrl}_{\texttt{isync}}$.
6. $\texttt{rmw} \subseteq \texttt{data} \cup \texttt{addr} \cup \texttt{ctrl}$
7. $\texttt{rmw}; \texttt{sb} \subseteq \texttt{ctrl}$

$1-5$ hold by definition (see [4]). $6-7$ hold due to the compilation scheme: it always places a dependency from the load to the store that form an RMW pair, and a branch after each (conditional) store in such pairs.

The relation $\texttt{deps}$ includes all types of dependencies:

$$\texttt{deps} \triangleq \texttt{data} \cup \texttt{addr} \cup \texttt{ctrl}$$

Herd's definition of $\texttt{ppo}$ is as follows:

$$\texttt{rdw} \triangleq (\texttt{rbe}; \texttt{rfe}) \cap \texttt{sb} \qquad \texttt{detour} \triangleq (\texttt{moe}; \texttt{rfe}) \cap \texttt{sb}$$
$$\texttt{ii}_0 \triangleq \texttt{addr} \cup \texttt{data} \cup \texttt{rdw} \cup \texttt{rfi} \qquad \texttt{ic}_0 \triangleq \emptyset$$
$$\texttt{ci}_0 \triangleq \texttt{ctrl}_{\texttt{isync}} \cup \texttt{detour} \qquad \texttt{cc}_0^{\mathsf{Power}} \triangleq \texttt{data} \cup \texttt{ctrl} \cup \texttt{addr}; \texttt{sb}^? \cup \texttt{sb}|_{\texttt{loc}}$$
$$\texttt{cc}_0^{\mathsf{ARMv7}} \triangleq \texttt{data} \cup \texttt{ctrl} \cup \texttt{addr}; \texttt{sb}^?$$
$$\texttt{ppo} \triangleq [\texttt{R}]; \texttt{ii}; [\texttt{R}] \cup [\texttt{R}]; \texttt{ic}; [\texttt{W}]$$

where, $\texttt{ii}, \texttt{ic}, \texttt{ci}, \texttt{cc}$ are inductively defined as follows:

$$\frac{\texttt{ii}_0}{\texttt{ii}} \qquad \frac{\texttt{ci}}{\texttt{ii}} \qquad \frac{\texttt{ic}; \texttt{ci}}{\texttt{ii}} \qquad \frac{\texttt{ii}; \texttt{ii}}{\texttt{ii}}$$
$$\frac{\texttt{ic}_0}{\texttt{ic}} \qquad \frac{\texttt{ii}}{\texttt{ic}} \qquad \frac{\texttt{cc}}{\texttt{ic}} \qquad \frac{\texttt{ic}; \texttt{cc}}{\texttt{ic}} \qquad \frac{\texttt{ii}; \texttt{ic}}{\texttt{ic}}$$
$$\frac{\texttt{ci}_0}{\texttt{ci}} \qquad \frac{\texttt{ci}; \texttt{ii}}{\texttt{ci}} \qquad \frac{\texttt{cc}; \texttt{ci}}{\texttt{ci}}$$
$$\frac{\texttt{cc}_0}{\texttt{cc}} \qquad \frac{\texttt{ci}}{\texttt{cc}} \qquad \frac{\texttt{ci}; \texttt{ic}}{\texttt{cc}} \qquad \frac{\texttt{cc}; \texttt{cc}}{\texttt{cc}}$$

Note that $\texttt{ci} \subseteq \texttt{ii} \subseteq \texttt{ic}$, as well as $\texttt{ci} \subseteq \texttt{cc} \subseteq \texttt{ic}$.

Alternatively the relations $\texttt{ii}, \texttt{ic}, \texttt{ci}, \texttt{cc}$ can be defined as follows:

$$\texttt{xy} \triangleq \bigcup_{n \geq 1} \texttt{x}^1\texttt{y}^1{}_0; \texttt{x}^2\texttt{y}^2{}_0; \ldots; \texttt{x}^n\texttt{y}^n{}_0$$

where:

- $\texttt{x}, \texttt{y}, \texttt{x}^1, \ldots, \texttt{x}^n, \texttt{y}^1, \ldots, \texttt{y}^n \in \{\texttt{i}, \texttt{c}\}$.
- If $\texttt{x} = \texttt{c}$ then $\texttt{x}^1 = \texttt{c}$.
- For every $1 \leq i \leq n-1$, if $\texttt{y}^i = \texttt{c}$ then $\texttt{x}^{i+1} = \texttt{c}$.
- If $\texttt{y} = \texttt{i}$ then $\texttt{y}^n = \texttt{i}$.

Note that the only difference between Power and ARMv7 is in the definition of $\texttt{cc}_0$. Henceforth, we only assume ARMv7's definition, which is weaker, so our proofs apply for both Power and ARMv7.

Next, we prove some useful properties of $\texttt{ppo}$. In all propositions below we assume some Power-consistent execution.

Proposition F.1. $\texttt{ppo}$ *is transitive.*

Proof. Immediately follows from the definition. □

Proposition F.2. $[\mathtt{W}];\mathtt{psbloc} \subseteq \mathtt{ii}$.

Proof. Let $\langle a,b\rangle \in [\mathtt{W}];\mathtt{psbloc}$ and let $x = \mathtt{loc}(a)$. Then, by definition, $a \in \mathtt{W}_x$, $b \in \mathtt{R}_x$, $\langle a,b\rangle \in \mathtt{sb}$, and there is no $c \in \mathtt{W}_x$ such that $\langle a,c\rangle, \langle c,b\rangle \in \mathtt{sb}$. Since G is complete, there exists some $d \in \mathtt{W}_x$ such that $\langle d,b\rangle \in \mathtt{rf}$. If $d = a$, then we are done since $\mathtt{rfi} \subseteq \mathtt{ii}$. Otherwise, since G satisfies SC-PER-LOC, we have $\langle a,d\rangle \in \mathtt{mo}$, $\langle d,a\rangle \notin \mathtt{sb}$, and $\langle b,d\rangle \notin \mathtt{sb}$. It follows that $\langle a,d\rangle \in \mathtt{moe}$ and $\langle d,b\rangle \in \mathtt{rfe}$. Thus, we have $\langle a,b\rangle \in \mathtt{detour} \subseteq \mathtt{ii}$. □

Proposition F.3. $(\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});[\mathtt{W}];\mathtt{psbloc};\mathtt{ppo};[\mathtt{W}] \subseteq \mathtt{ppo}$.

Proof. Let $a,b,c,d \in \mathtt{E}$ such that $\langle a,b\rangle \in (\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});[\mathtt{W}]$, $\langle b,c\rangle \in \mathtt{psbloc}$, and $\langle c,d\rangle \in \mathtt{ppo};[\mathtt{W}]$. If $\langle a,b\rangle \in \mathtt{ctrl}$, then by definition, we have $\langle a,d\rangle \in \mathtt{ctrl}$, and so $\langle a,d\rangle \in \mathtt{ppo}$. If $\langle a,b\rangle \in \mathtt{addr};\mathtt{sb}$, then by definition, we have $\langle a,d\rangle \in \mathtt{cc}$, and so $\langle a,d\rangle \in \mathtt{ppo}$. Otherwise, $\langle a,b\rangle \in \mathtt{addr} \cup \mathtt{data} \subseteq \mathtt{ii}$. By Prop. F.2, we also have $\langle b,c\rangle \in \mathtt{ii}$. Hence, $\langle a,c\rangle \in \mathtt{ii}$, and so $\langle a,c\rangle \in \mathtt{ppo}$. It follows that $\langle a,d\rangle \in \mathtt{ppo}$. □

Proposition F.4. $(\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});[\mathtt{R}];\mathtt{sb};[\mathtt{W}] \subseteq \mathtt{ppo}$.

Proof. Let $a,b,c \in \mathtt{E}$ such that $\langle a,b\rangle \in (\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});[\mathtt{R}]$ and $\langle b,c\rangle \in \mathtt{sb};[\mathtt{W}]$. If $\langle a,b\rangle \in \mathtt{ctrl}$, then by definition, we have $\langle a,c\rangle \in \mathtt{ctrl}$, and so $\langle a,c\rangle \in \mathtt{ppo}$. Otherwise, $\langle a,b\rangle \in \mathtt{addr};\mathtt{sb}^?$. In this case, we have $\langle a,c\rangle \in \mathtt{addr};\mathtt{sb}$, and so $\langle a,c\rangle \in \mathtt{ppo}$. □

Proposition F.5. *Let* $R = \mathtt{deps} \cup \mathtt{addr};\mathtt{sb} \cup \mathtt{psbloc}$. *Then,* $(\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});R^*;[\mathtt{W}] \subseteq \mathtt{ppo}$.

Proof. We prove by induction that for every $n \geq 0$, $(\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});R^n;[\mathtt{W}] \subseteq \mathtt{ppo}$. For $n = 0$, we have $(\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});[\mathtt{W}] \subseteq \mathtt{ppo}$ by definition. Let $n \geq 1$ and suppose that $(\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});R^k;[\mathtt{W}] \subseteq \mathtt{ppo}$ for every $k < n$. Let $\langle a,b\rangle \in (\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});R^n;[\mathtt{W}]$. Let $c \in \mathtt{E}$ such that $\langle a,c\rangle \in (\mathtt{deps} \cup \mathtt{addr};\mathtt{sb})$, and $\langle c,b\rangle \in R^n$. If $c \in \mathtt{R}$, then we are done using Prop. F.4. Otherwise, $c \in \mathtt{W}$, and $\langle c,b\rangle \in \mathtt{psbloc};R^{n-1}$. Let d be the $\mathtt{sb}$-maximal event satisfying $\langle c,d\rangle \in \mathtt{psbloc}$ and $\langle d,b\rangle \in R^k$ for some $k \leq n-1$. The maximality of d ensures that $\langle d,b\rangle \in (\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});R^{k-1}$. By the induction hypothesis, we have $\langle d,b\rangle \in \mathtt{ppo}$. Hence, we have $\langle a,b\rangle \in (\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});[\mathtt{W}];\mathtt{psbloc};\mathtt{ppo};[\mathtt{W}]$, and the claim follows by Prop. F.3. □

Proposition F.6. *Let* $R = \mathtt{deps} \cup \mathtt{addr};\mathtt{sb} \cup \mathtt{psbloc}$. *Then,* $\mathtt{rfe};R^+;[\mathtt{W}] \subseteq \mathtt{rfe};\mathtt{ppo}$.

Proof. Let $\langle a,c\rangle \in \mathtt{rfe};R^+;[\mathtt{W}]$. Let b be the $\mathtt{sb}$-maximal event satisfying $\langle a,b\rangle \in \mathtt{rfe}$ and $\langle b,c\rangle \in R^+$. If $\langle b,c\rangle \in (\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});R^*$, then we are done by Prop. F.5. Otherwise, let d be the $\mathtt{sb}$-maximal element such that $\langle b,d\rangle \in \mathtt{psbloc}$ and $\langle d,c\rangle \in R^*$. Then, $d \in \mathtt{R}$, and since $c \in \mathtt{W}$, we have $\langle d,c\rangle \in R^+$. The maximality of b and SC-PER-LOC ensure that $\langle b,d\rangle \in \mathtt{rdw}$, and so $\langle b,d\rangle \in \mathtt{ppo}$. The maximality of d ensures that $\langle d,c\rangle \in (\mathtt{deps} \cup \mathtt{addr};\mathtt{sb});R^*$. By Prop. F.5, we have $\langle d,c\rangle \in \mathtt{ppo}$, and so $\langle a,c\rangle \in \mathtt{rfe};\mathtt{ppo}$. □

Proposition F.7. $\mathtt{ppo};\mathtt{rbi} \subseteq \mathtt{ppo} \cup \mathtt{mo} \cup \mathtt{ppo};\mathtt{mo} \cup \mathtt{rbi}$.

Proof. For any $n \geq 0$, let $\mathtt{ppo}_n$ denote $\mathtt{ppo}$ edges that are formed by at most n basic $\mathtt{ppo}$ edges ($\mathtt{ii}_0$, $\mathtt{ic}_0$, $\mathtt{ci}_0$, and $\mathtt{cc}_0$). Then, $\mathtt{ppo}^? = \bigcup_{n\geq 0} \mathtt{ppo}_n$. The proof proceeds by induction on n. For $n = 0$, the claim obviously holds. Suppose now that it holds for $n-1$, and let $\langle a,b\rangle \in \mathtt{ppo}_n$ and $\langle b,c\rangle \in \mathtt{rbi}$. Then, b must be a read event, and so there exists a' such that $\langle a,a'\rangle \in \mathtt{ppo}_{n-1}$ and $\langle a',b\rangle \in \mathtt{ii}_0 \cup \mathtt{ci}_0$. This leads to five cases:

- $\langle a',b\rangle \in \mathtt{addr}$. In this case we have $\langle a',c\rangle \in \mathtt{cc}_0$, and so $\langle a,c\rangle \in \mathtt{ppo}$.
- $\langle a',b\rangle \in \mathtt{rdw}$. In this case we have $\langle a',c\rangle \in \mathtt{rbi}$, and the claim follows by the induction hypothesis.
- $\langle a',b\rangle \in \mathtt{rfi}$. In this case we have $\langle a',c\rangle \in \mathtt{mo}$, and so $\langle a,c\rangle \in \mathtt{ppo}^?;\mathtt{mo}$.
- $\langle a',b\rangle \in \mathtt{ctrl}_{\mathtt{isync}}$. In this case we have $\langle a',c\rangle \in \mathtt{ci}_0$, and so $\langle a,c\rangle \in \mathtt{ppo}$.
- $\langle a',b\rangle \in \mathtt{detour}$. In this case we have $\langle a',c\rangle \in \mathtt{mo}$, and so $\langle a,c\rangle \in \mathtt{ppo}^?;\mathtt{mo}$. □

F.2 Additional Properties

Proposition F.8. $\texttt{rmw} \cap (\texttt{rb}; \texttt{mo}) = \emptyset$.

Proof. POWER-ATOMICITY condition ensures that $\texttt{rmw} \cap (\texttt{rbe}; \texttt{moe}) = \emptyset$. In addition, in every execution we have $\texttt{rmw} \subseteq \texttt{sb}$, $\texttt{rbe}; \texttt{sb} \not\subseteq \texttt{sb}$, $\texttt{sb}; \texttt{moe} \not\subseteq \texttt{sb}$, and $\texttt{sb}; \texttt{sb} \not\subseteq \texttt{rmw}$. It follows that $\texttt{rmw} \cap (\texttt{rb}; \texttt{mo}) = \emptyset$. □

Proposition F.9. *Let* $R \in \{\texttt{sync}, \texttt{fence}\}$. *Then,* $R; \texttt{hb}^*; \texttt{rbi} \subseteq R; \texttt{hb}^*; \texttt{mo}^?$.

Proof. We prove by induction on n that for every $n \geq 0$, we have $R; \texttt{hb}^n; \texttt{rbi} \subseteq R; \texttt{hb}^*; \texttt{mo}^?$. For $n = 0$, the claim follows since $R; \texttt{rbi} \subseteq R$. Now, suppose it holds for $n-1$, and let a, b, c, d such that $\langle a, b\rangle \in R; \texttt{hb}^{n-1}$, $\langle b, c\rangle \in \texttt{hb}$, and $\langle c, d\rangle \in \texttt{rbi}$. If $\langle b, c\rangle \in \texttt{rfe}$, then we have $\langle b, d\rangle \in \texttt{mo}$, and so $\langle a, d\rangle \in R; \texttt{hb}^*; \texttt{mo}$. If $\langle b, c\rangle \in \texttt{fence}$, then we have $\langle b, d\rangle \in \texttt{fence}$, and so $\langle a, d\rangle \in R; \texttt{hb}^*$. Otherwise, we have $\langle b, c\rangle \in \texttt{ppo}$, and the claim follows using Prop. F.7 and the induction hypothesis. □

Proposition F.10. $\texttt{fence}$ *is transitive.*

Proof. Immediately follows from the definition of $\texttt{fence}$. □

Proposition F.11. $\texttt{fence}; \texttt{hb}^* \subseteq \texttt{sb} \cup \texttt{fence}; [\texttt{W}]; \texttt{hb}^*$.

Proof. Let $a, b, c \in \texttt{E}$ such that $\langle a, b\rangle \in \texttt{fence}$ and $\langle b, c\rangle \in \texttt{hb}^*$. If $\langle b, c\rangle \in \texttt{sb}$, then the claim follows since $\texttt{fence} \subseteq \texttt{sb}$. Suppose otherwise. Then, there exists $\langle d, e\rangle \in \texttt{rfe}$ such that $\langle b, d\rangle \in \texttt{hb}^* \cap \texttt{sb}^?$ and $\langle e, c\rangle \in \texttt{hb}^*$. It follows that $\langle a, d\rangle \in \texttt{fence}$, and so $\langle a, c\rangle \in \texttt{fence}; [\texttt{W}]; \texttt{hb}^*$. □

Proposition F.12. $[\texttt{RW}]; \texttt{sb}; (\texttt{fence}; \texttt{hb}^*)^?; \texttt{sync} \subseteq (\texttt{fence}; \texttt{hb}^*)^?; \texttt{sync}$.

Proof. Immediately follows from the definition of $\texttt{sync}$ and Prop. F.11. □

Proposition F.13. $\texttt{eco}^?; (\texttt{fence}; \texttt{hb}^*)^?; \texttt{sync}; \texttt{hb}^*$ *is acyclic.*

Proof. By definition, we have $\texttt{eco}^? = (\texttt{mo} \cup \texttt{rbe})^?; \texttt{rf}^? \cup \texttt{rbi}; \texttt{rfi}^? \cup \texttt{rbi}; \texttt{rfe}$. Thus, it suffices to show that the union of the following relations is acyclic:

- $A = ((\texttt{mo} \cup \texttt{rbe})^?; \texttt{rf}^? \cup \texttt{rbi}; \texttt{rfi}^?); (\texttt{fence}; \texttt{hb}^*)^?; \texttt{sync}; \texttt{hb}^*$
- $B = \texttt{rbi}; \texttt{rfe}; (\texttt{fence}; \texttt{hb}^*)^?; \texttt{sync}; \texttt{hb}^*$

By Prop. F.9, $A; B \subseteq A$; A and $B; B \subseteq B; A$. Hence, it suffices to show that A is acyclic and B is irreflexive. Acyclicity of A follows from Power's PROPAGATION condition, since we have $A \subseteq \texttt{mo}^?; \texttt{prop}_2$ (using Prop. F.12). Irreflexivity of B also follows from PROPAGATION, using Prop. F.9. □

Proposition F.14. *Let* $S = \texttt{sb}^?; [\texttt{F}]; \texttt{sb}; \texttt{rfe}; \texttt{hb}^*; (\texttt{sb}; [\texttt{F}] \cup \texttt{ecoi})^?$. *Then,* S *is a strict partial order.*

Proposition F.15. $\texttt{eco}; (\texttt{sb} \cup \texttt{fence}; \texttt{hb}^*)$ *is irreflexive.*

Proof. $\texttt{eco}; \texttt{sb}$ is irreflexive using SC-PER-LOC. By Prop. F.11, it suffices to show that $\texttt{eco}; \texttt{fence}; [\texttt{W}]; \texttt{hb}^*$ is irreflexive. Suppose otherwise, and let $a, b \in \texttt{E}$ such that $\langle a, b\rangle \in \texttt{eco}$ and $\langle b, a\rangle \in \texttt{fence}; [\texttt{W}]; \texttt{hb}^*$. First, if $\langle a, b\rangle \in \texttt{sb}$, then we have $\langle a, a\rangle \in \texttt{fence}; \texttt{hb}^* \subseteq \texttt{hb}^+$, which contradicts POWER-NO-THIN-AIR. Suppose otherwise, and consider the possible cases:

- $\langle a, b\rangle \in \texttt{rfe}$. In this case we obtain $\langle a, a\rangle \in \texttt{hb}^+$, which contradicts POWER-NO-THIN-AIR.
- $\langle a, b\rangle \in \texttt{mo}; \texttt{rf}^?$. Let $c \in \texttt{E}$ such that $\langle a, c\rangle \in \texttt{mo}$ and $\langle c, b\rangle \in \texttt{rf}^?$. Then, we have $\langle c, a\rangle \in \texttt{prop}_1$, and we obtain that $\texttt{mo}; \texttt{prop}_1$ is not irreflexive, which contradicts PROPAGATION.
- $\langle a, b\rangle \in \texttt{rbe}; \texttt{rf}^?$. Let $c \in \texttt{W}$ such that $\langle a, c\rangle \in \texttt{rbe}$ and $\langle c, b\rangle \in \texttt{rf}^?$. Let $d \in \texttt{W}$ such that $\langle b, d\rangle \in \texttt{fence}$ and $\langle d, a\rangle \in \texttt{hb}^*$. Then, we have $\langle c, d\rangle \in \texttt{prop}_1$, and obtain a violation of OBSERVATION.
- $\langle a, b\rangle \in \texttt{rbi}; \texttt{rfe}$. Let $c \in \texttt{W}$ such that $\langle a, c\rangle \in \texttt{rbi}$ and $\langle c, b\rangle \in \texttt{rfe}$. By Prop. F.9, we have $\langle b, c\rangle \in \texttt{fence}; \texttt{hb}^*; \texttt{mo}^?$. Let $d \in \texttt{E}$ such that $\langle b, d\rangle \in \texttt{fence}; \texttt{hb}^*$ and $\langle d, c\rangle \in \texttt{mo}^?$. Then, we have $\langle c, d\rangle \in \texttt{prop}_1$, and we obtain that $\texttt{mo}^?; \texttt{prop}_1$ is not irreflexive, which contradicts PROPAGATION. □

F.3 Removing Redundant Fences

Lemma F.1. *Let G be a* Power *execution, and let $\langle a, b\rangle \in [\mathtt{F}^{\mathtt{sync}}]; \mathtt{sb}|_{imm}; [\mathtt{F}^{\mathtt{lwsync}}]$. Let G' be the execution obtained from G by removing b ($G' = G|_{G.\mathtt{E}\setminus\{b\}}$). If G' is* Power*-consistent, then so is G.*

Proof. Since b's immediate $\mathtt{sb}$-predecessor is a full fence, we have $G'.\mathtt{fence} = G.\mathtt{fence}$. Then, it is easy to see that for every relation $\mathtt{c}$ mentioned in Def. 6, we have $G'.\mathtt{c} = G.\mathtt{c}$, and so if G' is Power-consistent, then so is G. □

Lemma F.2. *Let G be a* Power *execution, and let $\langle a, b\rangle \in [\mathtt{R}]; (\mathtt{sb}|_{imm} \cap \mathtt{ctrl}_{\mathtt{isync}}); [\mathtt{F}]$. Let G' be the execution obtained from G by removing the $\mathtt{ctrl}_{\mathtt{isync}}$ dependency edges from a onwards ($G'.\mathtt{ctrl}_{\mathtt{isync}} = G.\mathtt{ctrl}_{\mathtt{isync}} \setminus (\{a\} \times \mathtt{E})$). If G' is* Power*-consistent, then so is G.*

Proof. Since a's immediate $\mathtt{sb}$-successor is a fence, we have $\langle a, c\rangle \in G.\mathtt{fence}$ for every $c \in \mathtt{RW}$ such that $\langle a, c\rangle \in \mathtt{sb}$. Now, by omitting $\mathtt{ctrl}_{\mathtt{isync}}$ dependency edges from a onwards, we may remove ppo edges from a, but whenever ppo is used to form an hb edge, it can be replaced by a $\mathtt{fence}$ edge. Consequently, for every relation $\mathtt{c}$ mentioned in Def. 6, we have $G'.\mathtt{c} = G.\mathtt{c}$, and so if G' is Power-consistent, then so is G. □

G. Power-before Relation

In this section, we define a relation that we call *Power-before* (pb), and show that if pb is acyclic in some execution G of a program P, then either G is RC11-consistent, or P has undefined behavior under RC11. This relation is the key for showing that NO-THIN-AIR holds when proving compilation correctness. (Thus, if one is only interested in weakRC11-consistency, this section can be completely ignored.)

In what follows we assume an execution G.

pb is given by:

$$\mathtt{psbloc} \triangleq [\mathtt{E} \setminus \mathtt{E}_0]; \mathtt{sb}|_{\mathtt{loc}}; [\mathtt{R}] \setminus \mathtt{sb}|_{\mathtt{loc}}; [\mathtt{W}]; \mathtt{sb} \qquad \textit{(preserved sb-loc)}$$

$$\mathtt{pbi} \triangleq \mathtt{deps} \cup \mathtt{addr}; \mathtt{sb} \cup [\mathtt{R}^{\sqsupseteq \mathtt{rlx}} \cup \mathtt{W}^{\sqsupseteq \mathtt{rel}} \cup \mathtt{F}]; \mathtt{sb} \cup \mathtt{psbloc} \cup \mathtt{sb}; [\mathtt{E}^{\sqsupseteq \mathtt{rel}}] \qquad \textit{(internal Power-before)}$$

$$\mathtt{pb} \triangleq \mathtt{pbi} \cup \mathtt{rfe} \qquad \textit{(Power-before)}$$

Clearly, $\mathtt{pb} \subseteq \mathtt{sb} \cup \mathtt{rf}$, and so pb is acyclic in every RC11-consistent execution.

Proposition G.1. *If G is* weakRC11*-consistent, then $\mathtt{rf} \subseteq \mathtt{pb}$.*

Proof. COHERENCE guarantees that $\mathtt{rfi} \subseteq \mathtt{psbloc} \subseteq \mathtt{pbi}$, and by definition we have $\mathtt{rfe} \subseteq \mathtt{pb}$. □

Proposition G.2. *For every* weakRC11*-consistent execution G, $\mathtt{hb} \subseteq \mathtt{sb} \cup \mathtt{pb}^+$.*

Proof. It suffices to show that $\mathtt{sb}^?; \mathtt{swe}; \mathtt{sb}^? \subseteq \mathtt{pb}^+$. By definition, we have

$$\mathtt{sb}^?; \mathtt{swe}; \mathtt{sb}^? \subseteq \mathtt{sb}^?; [\mathtt{E}^{\sqsupseteq \mathtt{rel}}]; \mathtt{sb}^?; (\mathtt{rf} \cup \mathtt{rmw})^+; [\mathtt{R}^{\sqsupseteq \mathtt{rlx}}]; \mathtt{sb}^?.$$

The claim follows because we have:

- $\mathtt{sb}^?; [\mathtt{E}^{\sqsupseteq \mathtt{rel}}]; \mathtt{sb}^? \subseteq \mathtt{pbi}^*$
- $\mathtt{rf} \subseteq \mathtt{pb}$ and $\mathtt{rmw} \subseteq \mathtt{deps} \subseteq \mathtt{pbi}$.
- $[\mathtt{R}^{\sqsupseteq \mathtt{rlx}}]; \mathtt{sb}^? \subseteq \mathtt{pbi}^?$. □

Proposition G.3. *If pb is acyclic, but $\mathtt{sb} \cup \mathtt{rf}$ is cyclic, then $(\mathtt{rfe}; [\mathtt{R}^{\mathtt{na}}] \setminus \mathtt{hb}); \mathtt{sb} \neq \emptyset$.*

Proof. A cycle in $\mathtt{sb} \cup \mathtt{rf}$ implies a cycle in $\mathtt{rfe}; \mathtt{sb}$. Since $\mathtt{rfe}; [\mathtt{R}^{\sqsupseteq \mathtt{rlx}}]; \mathtt{sb}$ and $(\mathtt{rfe} \cap \mathtt{hb}); \mathtt{sb}$ are contained in $\mathtt{pb}^+$ (using Prop. G.2 for the latter), there must exist an edge $\langle a, b\rangle \in \mathtt{rfe}; \mathtt{sb}$ that is neither in $\mathtt{rfe}; [\mathtt{R}^{\sqsupseteq \mathtt{rlx}}]; \mathtt{sb}$ nor in $(\mathtt{rfe} \cap \mathtt{hb}); \mathtt{sb}$. Then, we have $\langle a, b\rangle \in (\mathtt{rfe}; [\mathtt{R}^{\mathtt{na}}] \setminus \mathtt{hb}); \mathtt{sb}$. □

Lemma G.1. *Suppose that G is a* weakRC11*-consistent execution of a program P, and that pb is acyclic, but G is not* RC11*-consistent. Then, P has undefined behavior under* RC11.

Proof. Since G is weakRC11-consistent but not RC11-consistent, we have that $\mathtt{sb} \cup \mathtt{rf}$ is cyclic. By Prop. G.3, $\mathtt{rf}; [\mathtt{R}^{\mathtt{na}}] \not\subseteq \mathtt{hb}$. We show that this implies that P has undefined behavior under RC11.

Let $a_1, \ldots, a_n$ be an enumeration of $\mathtt{E}$ that respects $\mathtt{pb}$ (that is, $i < j$ whenever $\langle a_i, a_j \rangle \in \mathtt{pb}^+$). For every $1 \leq i \leq n$, let $E_i = \{a_1, \ldots, a_i\}$. Let k be the minimal index such that $[E_k]; \mathtt{rf}; [\mathtt{R}^{\mathtt{na}}]; [E_k] \not\subseteq \mathtt{hb}$. Then, we have $\langle a_j, a_k \rangle \in \mathtt{rf}; [\mathtt{R}^{\mathtt{na}}] \setminus \mathtt{hb}$ for some $j < k$. Let $B = \mathit{dom}(\mathtt{sb}^?; [E_k])$ and $H = B \setminus E_k$.

Claim 1: $h \in \mathtt{R}^{\mathtt{na}} \cup \mathtt{W}^{\sqsubseteq \mathtt{rlx}}$ for every $h \in H$.

Proof: Otherwise, since $[\mathtt{R}^{\sqsupseteq \mathtt{rlx}} \cup \mathtt{W}^{\sqsupseteq \mathtt{rel}} \cup \mathtt{F}]; \mathtt{sb} \subseteq \mathtt{pb}$, we would obtain $\langle h, a \rangle \in \mathtt{pb}$ for some $a \in E_k$. This contradicts the fact that $h \notin E_k$. □

Claim 2: $\langle h, b \rangle \notin \mathtt{sb}^?$ for every $h \in H$ and $b \in B \cap (\mathtt{E}^{\sqsupseteq \mathtt{rel}})$.

Proof: Suppose otherwise. Let $a \in E_k$ such that $\langle b, a \rangle \in \mathtt{sb}^?$. It follows that $\langle h, a \rangle \in \mathtt{sb}^?; \mathtt{E}^{\sqsupseteq \mathtt{rel}}; \mathtt{sb}^?$, and so $\langle h, a \rangle \in \mathtt{pb}^*$. Hence, $h \in E_k$ as well, which contradicts our assumption. □

Claim 3: $\langle h, b \rangle \notin \mathtt{deps}^*; \mathtt{ctrl}$ for every $h \in H$ and $b \in B$.

Proof: Suppose otherwise. Let $a \in E_k$ such that $\langle b, a \rangle \in \mathtt{sb}^?$. Since $\mathtt{ctrl}; \mathtt{sb}^? \subseteq \mathtt{ctrl}$, it follows that $\langle h, a \rangle \in \mathtt{deps}^+$, and so $\langle h, a \rangle \in \mathtt{pb}^+$. This contradicts the fact that $h \notin E_k$. □

Claim 4: $\langle h, b \rangle \notin \mathtt{deps}^*; \mathtt{addr}$ for every $h \in H$ and $b \in B$.

Proof: Suppose otherwise. Let $a \in E_k$ such that $\langle b, a \rangle \in \mathtt{sb}^?$. Then, $\langle h, a \rangle \in \mathtt{deps}^*; \mathtt{addr}; \mathtt{sb}^? \subseteq \mathtt{pb}^+$. This contradicts the fact that $h \notin E_k$. □

Let $h_1, \ldots, h_m$ be an enumeration of H that respects $\mathtt{sb}$, and let $H_i = \{h_1, \ldots, h_i\}$ for every $0 \leq i \leq m$.

Claim 5: For every $1 \leq i \leq m$, $h_i \notin \mathit{dom}(\mathtt{deps}^+; [E_k \cup H_{i-1}])$.

Proof: Suppose otherwise, and let $a \in E_k \cup H_{i-1}$ such that $\langle h_i, a \rangle \in \mathtt{deps}^+$. Then, $\langle h_i, a \rangle \in \mathtt{pb}^+$. If $a \in E_k$, then $h_i \in E_k$ as well, which contradicts our assumption. Hence, we have $a \in H_{i-1}$. This contradicts the fact that the h_i's enumeration respects $\mathtt{sb}$. □

Claim 6: Let $1 \leq i \leq m$, and let $x = \mathtt{loc}(h_i)$. Let $a \in (E_k \cup H_{i-1}) \cap \mathtt{R}_x$ and suppose that $\langle h_i, a \rangle \in \mathtt{sb}$. Then, $\langle h_i, a \rangle \in \mathtt{sb}; [(E_k \cup H_{i-1}) \cap \mathtt{W}_x]; \mathtt{sb}$.

Proof: Suppose otherwise. Let $i \leq j \leq m$ be the maximal index satisfying $h_j \in \mathtt{E}_x$, $\langle h_i, h_j \rangle \in \mathtt{sb}^?$ and $\langle h_j, a \rangle \in \mathtt{sb}$. Then, $\langle h_j, a \rangle \in \mathtt{psbloc}$, and so $\langle h_j, a \rangle \in \mathtt{pb}$. If $a \in E_k$, then $h_j \in E_k$ as well, which contradicts our assumption. Hence, we have $a \in H_{i-1}$. This contradicts the fact that the h_i's enumeration respects $\mathtt{sb}$. □

For every $1 \leq i \leq n$, let and $G_i = G|_{E_i}$. Since $G.\mathtt{rf} \subseteq G.\mathtt{pb}$ (Prop. G.1), all the G_i's are weakRC11-consistent. Additionally, $G_i.\mathtt{pb}$ is acyclic for every $1 \leq i \leq n$.

We inductively construct a sequences of labeling functions $lab_0, \ldots, lab_m : B \to \mathsf{Label}$ and executions $G'_0, \ldots, G'_m$ such that the following hold:

1. For every $0 \leq i \leq m$, $G'_i.\mathtt{E} = E_k \cup H_i$.
2. For every $0 \leq i \leq m$, $G'_i.\mathtt{lab} = lab_i|_{G'_i.\mathtt{E}}$.
3. For every $0 \leq i \leq m$, G'_i is $\mathsf{RC}_{\mathtt{na}}$-consistent.
4. For every $0 \leq i \leq m$, $\langle a_j, a_k \rangle, \langle a_k, a_j \rangle \notin G'_i.\mathtt{hb}$.
5. For every $0 \leq i \leq m$, $lab_i(G|_B)$ is an execution of P.
6. For every $0 \leq i \leq m$, $G.\mathtt{rmw}^{-1}; G'_i.\mathtt{rf}^{-1}; G'_i.\mathtt{rf}; G.\mathtt{rmw} \subseteq [G.\mathtt{E}]$.

Finally, we would obtain that G'_m is a racy $\mathsf{RC}_{\mathtt{na}}$-consistent execution with $G'_m.\mathtt{E} = B$, and $lab_m(G|_B) = G'_m.\mathtt{lab}(G|_B)$ is an execution of P. Hence, G'_m is an execution of P, and by Lemma D.1, G'_m is RC11-consistent or P has undefined behavior under RC11. Since G'_m is racy, in any case we would obtain that P has undefined behavior under RC11.

First, we define lab_0 and G'_0. The minimality of k and Prop. G.3 ensure that G_{k-1} is RC11-consistent. Hence, Lemma D.4 ensures that there exists some event $b \in E_k$ such that the execution G' given by $G'.\mathtt{c} = G_k.\mathtt{c}$ for every $\mathtt{c} \in \{\mathtt{E}, \mathtt{sb}, \mathtt{rmw}, \mathtt{data}, \mathtt{addr}, \mathtt{ctrl}, \mathtt{mo}\}$, $G'.\mathtt{lab} = G_k.\mathtt{lab}[a_k \mapsto \mathtt{R}^{\mathtt{na}}(G.\mathtt{loc}(a_k), G.\mathtt{val}_{\mathtt{w}}(b))]$,

and $G'.\mathtt{rf} = G_k.\mathtt{rf} \cup \{\langle b, a_k\rangle\}$ is $\mathsf{RC}_{\mathtt{na}}$-consistent. In addition, $a_k \notin dom(G|_B.\mathtt{deps})$ (since it is $G.\mathtt{pb}$ maximal in $G|_B$). By Assumption B.1, there exists a reevaluation lab of $G.\mathtt{lab}$ such that $lab(G|_B)$ is an execution of P, $lab(G|_B).\mathtt{val}_{\mathtt{r}}(a_k) = G.\mathtt{val}_{\mathtt{w}}(b)$, and $lab(c) = G|_B.\mathtt{lab}(c)$ for every $c \in B \setminus \{a_k\}$. We take $lab_0 = lab$ and $G'_0 = G'$. It is straightforward to see that lab and G' satisfy the six conditions above. In particular, $G|_B.\mathtt{rmw}^{-1}; G'.\mathtt{rf}^{-1}; G'.\mathtt{rf}; G|_B.\mathtt{rmw} \subseteq [G.\mathtt{E}]$ follows from the fact that G satisfies ATOMICITY. Additionally, by Prop. C.4, $G'.\mathtt{hb} = G_k.\mathtt{hb}$, and so, we have $\langle a_j, a_k\rangle, \langle a_k, a_j\rangle \notin G'.\mathtt{hb}$.

Next, let $1 \leq i \leq m$, and suppose that lab_{i-1} and G'_{i-1} are defined. We construct lab_i and G'_i. By Claim 1 above, we have $h_i \in G.\mathtt{R}^{\mathtt{na}} \cup G.\mathtt{W}^{\sqsubseteq \mathtt{rlx}}$. Let G^*_i be the execution obtained from G'_{i-1} by adding the event h_i, labeled with $lab_{i-1}(h_i)$, and the $\mathtt{sb}$, $\mathtt{rmw}$, and dependency edges from/to h_i as in $G|_B$. By Claim 2 above, we also have $h_i \notin dom(G^*_i.\mathtt{sb}; [G^*_i.\mathtt{E}^{\mathtt{rel}}])$. Let $x = G.\mathtt{loc}(h_i)$, and consider the two cases:

$h_i \in G.\mathtt{R}^{\mathtt{na}}$**:** Since G'_{i-1} is $\mathsf{RC}_{\mathtt{na}}$-consistent, Lemma D.4 ensures that there exists some event $b \in E_k \cup H_{i-1}$ such that the execution G' given by $G'.\mathtt{E} = E_k \cup H_i$, $G'.\mathtt{lab} = G'_{i-1}.\mathtt{lab} \cup \{h_i \mapsto \mathtt{R}^{\mathtt{na}}(x, G^*_i.\mathtt{val}_{\mathtt{w}}(b))\}$, $G'.\mathtt{c} = G^*_i.\mathtt{c}$ for every $\mathtt{c} \in \{\mathtt{sb}, \mathtt{rmw}, \mathtt{data}, \mathtt{addr}, \mathtt{ctrl}, \mathtt{mo}\}$, and $G'.\mathtt{rf} = G^*_i.\mathtt{rf} \cup \{\langle b, a_k\rangle\}$ is $\mathsf{RC}_{\mathtt{na}}$-consistent. In addition, by Claims 3 and 4 above, we have that $h_i \notin dom(G|_B.\mathtt{deps}^*; (G|_B.\mathtt{ctrl} \cup G|_B.\mathtt{addr}))$. By Assumption B.1, there exists a reevaluation lab of lab_{i-1} such that $lab(G|_B)$ is an execution of P, $lab(G|_B).\mathtt{val}_{\mathtt{r}}(h_i) = G.\mathtt{val}_{\mathtt{w}}(b)$, and $lab(c) = lab_{i-1}(c)$ for every c such that $\langle h_i, c\rangle \notin G|_B.\mathtt{deps}^+$. We take $lab_i = lab$ and $G'_i = G'$. Again, it is straightforward to see that lab and G' satisfy the required conditions. In particular, $G'_i.\mathtt{lab} = lab_i|_{G'_i.\mathtt{E}}$ follows from the fact that $G'_{i-1}.\mathtt{lab} = lab_{i-1}|_{E_k \cup H_{i-1}}$, and Claim 5 above. In addition, by Prop. C.5, we have $[G'_{i-1}.\mathtt{E}]; G'.\mathtt{hb}; [G'_{i-1}.\mathtt{E}] = G'_{i-1}.\mathtt{hb}$, and so, we have $\langle a_j, a_k\rangle, \langle a_k, a_j\rangle \notin G'.\mathtt{hb}$.

$h_i \in G.\mathtt{W}^{\sqsubseteq \mathtt{rlx}}$**:** By Claim 6 above, we have that for every $b \in G^*_i.\mathtt{E}$, if $\langle h_i, b\rangle \in G^*_i.\mathtt{sb}; [G^*_i.\mathtt{R}_x]$ then $\langle a, b\rangle \in G^*_i.\mathtt{sb}; [G^*_i.\mathtt{W}_x]; G^*_i.\mathtt{sb}$. Thus, since G'_{i-1} is $\mathsf{RC}_{\mathtt{na}}$-consistent, and $G.\mathtt{rmw}^{-1}; G'_{i-1}.\mathtt{rf}^{-1}; G'_{i-1}.\mathtt{rf}; G.\mathtt{rmw} \subseteq [G.\mathtt{E}]$, Lemmas D.2 and D.3 ensure that there exists $T \subseteq G^*_i.\mathtt{W}_x \times G^*_i.\mathtt{W}_x$ such that the execution G' given by $G'.\mathtt{E} = E_k \cup H_i$, $G'.\mathtt{lab} = lab_{i-1}|_{G'.\mathtt{E}}$, $G'.\mathtt{c} = G^*_i.\mathtt{c}$ for every $\mathtt{c} \in \{\mathtt{sb}, \mathtt{rmw}, \mathtt{data}, \mathtt{addr}, \mathtt{ctrl}\}$, $G'.\mathtt{rf} = G'_{i-1}.\mathtt{rf}$, and $G^*_i.\mathtt{mo} = G'_{i-1}.\mathtt{mo} \cup T$ is $\mathsf{RC}_{\mathtt{na}}$-consistent. We take $lab_i = lab_{i-1}$ and $G'_i = G'$. It is straightforward to see that lab_{i-1} and G' satisfy the required conditions. In particular, Prop. C.3 guarantees that $\langle a_j, a_k\rangle, \langle a_k, a_j\rangle \notin G'.\mathtt{hb}$. □

H. Proof of Compilation Correctness

In this section we prove the correctness of compilation from RC11 to Power. Lemma E.3 allows us to assume that there are no SC-accesses in the source execution. In addition, the next lemma allows us to assume that there are no release writes:

Lemma H.1. *Let G_{tgt} be an* RC11*-consistent execution. Let $\langle f, w\rangle \in [\mathtt{F}^{\mathtt{rel}}]; G_{\mathsf{tgt}}.\mathtt{sb}|_{imm}; \mathtt{rmw}^?$. Suppose that $G_{\mathsf{tgt}}.\mathtt{lab}(w) = \mathtt{W}^{\mathtt{rlx}}(x, v)$. Let G_{src} be the execution satisfying:*

- $G_{\mathsf{src}}.\mathtt{E} = G_{\mathsf{tgt}}.\mathtt{E} \setminus \{f\}$.
- $G_{\mathsf{src}}.\mathtt{lab} = G_{\mathsf{tgt}}.\mathtt{lab}|_{G_{\mathsf{src}}.\mathtt{E}}[w \mapsto \mathtt{W}^{\mathtt{rel}}(x, v)]$.
- $G_{\mathsf{src}}.\mathtt{sb} = [G_{\mathsf{src}}.\mathtt{E}]; G_{\mathsf{tgt}}.\mathtt{sb}; [G_{\mathsf{src}}.\mathtt{E}]$.
- $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf}$.
- $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo}$.

Then, G_{src} is RC11*-consistent, and it is racy if G_{tgt} is racy.*

Lemma H.2. *Let G be an execution without SC accesses and release writes. Let G_p be a* Power *execution. Suppose that the following hold:*

- $G.\mathtt{E} = G_p.\mathtt{E}$, $G.\mathtt{sb} = G_p.\mathtt{sb}$, $G.\mathtt{rmw} = G_p.\mathtt{rmw}$, $G.\mathtt{rf} = G_p.\mathtt{rf}$, *and* $G.\mathtt{mo} = G_p.\mathtt{mo}$.
- $G_p.\mathtt{lab}$ *is obtained from* $G.\mathtt{lab}$ *by removing the access mode (for accesses), replacing* $\mathtt{F}^{\mathtt{sc}}$ *by full fences, and replacing all other fence labels by lightweight fences.*
- $G.\mathtt{data} = G_p.\mathtt{data}$, $G.\mathtt{addr} = G_p.\mathtt{addr}$, *and* $G.\mathtt{ctrl} \subseteq G_p.\mathtt{ctrl}$.
- $G.\mathtt{rmw}; G.\mathtt{sb} \subseteq G_p.\mathtt{ctrl}$.
- $[G.\mathtt{R}^{\mathtt{rlx}} \setminus G.\mathtt{At}]; G.\mathtt{sb} \subseteq G_p.\mathtt{ctrl}$.

- $[G.\mathtt{R}^{\mathtt{acq}}]; G.\mathtt{sb} \subseteq G_p.\mathtt{rmw} \cup G_p.\mathtt{ctrl}_{\mathtt{isync}}$.

Then:

- *G and G_p have the same outcome.*
- *If G_p is* Power*-consistent, then G is* weakRC11*-consistent and $G.\mathtt{pb}$ is acyclic.*

Proof. The first claim easily follows from our definitions. Suppose that G_p is Power-consistent. Before proving the second claim, we present some properties relating G and G_p. Let $S = G_p.\mathtt{sb}^?; [G_p.\mathtt{F}]; G_p.\mathtt{sb}; G_p.\mathtt{rfe}; G_p.\mathtt{hb}^*; (G_p.\mathtt{sb}; [G_p.\mathtt{F}] \cup G_p.\mathtt{ecoi})^?$. Then, the following hold:

1. $G.\mathtt{swe}; G.\mathtt{sb}^? \subseteq S$
   (follows from the definition of $\mathtt{sw}$)
2. $G.\mathtt{hb} \subseteq G_p.\mathtt{sb} \cup S$
   (follows from Item 1 using Prop. F.14)
3. $[G.\mathtt{RW}]; (G.\mathtt{sb} \setminus G.\mathtt{rmw}); G.\mathtt{hb}^? \subseteq G_p.\mathtt{sb} \cup G_p.\mathtt{fence}; G_p.\mathtt{hb}^*; (G_p.\mathtt{sb}; [G_p.\mathtt{F}] \cup G_p.\mathtt{ecoi})^?$
   (follows from Item 1 using Prop. F.14)
4. $[G.\mathtt{F}^{\mathtt{sc}}]; G.\mathtt{hb}; [G.\mathtt{RW}] \subseteq [G_p.\mathtt{F}^{\mathtt{sync}}]; G_p.\mathtt{sb}; G_p.\mathtt{hb}^*; G_p.\mathtt{ecoi}^?; [G_p.\mathtt{RW}]$
   (easily follows from Item 2)

In addition, in order to apply Prop. C.2 in the proof below, we note that:

- $[G.\mathtt{W}]; G.\mathtt{sb}|_{\mathtt{loc}}; [G.\mathtt{W}] \subseteq G.\mathtt{mo}$: Indeed, we have $[G.\mathtt{W}]; G.\mathtt{sb}|_{\mathtt{loc}}; [G.\mathtt{W}] = [G_p.\mathtt{W}]; G_p.\mathtt{sb}|_{\mathtt{loc}}; [G_p.\mathtt{W}]$ and $G.\mathtt{mo} = G_p.\mathtt{mo}$, and the claim follows by Power's SC-PER-LOC condition.
- $G.\mathtt{rmw} \subseteq G.\mathtt{rb}$: Indeed, we have $G.\mathtt{rmw} = G_p.\mathtt{rmw}$ and $G.\mathtt{rb} = G_p.\mathtt{rb}$, and the claim follows by Power's SC-PER-LOC and the fact that G is complete.

Next, we show that G is weakRC11-consistent. Clearly, it is complete (since $G.\mathtt{R} = G_p.\mathtt{R}$ and $G.\mathtt{rf} = G_p.\mathtt{rf}$).

COHERENCE. We show that $G.\mathtt{eco}^?; G.\mathtt{hb}$ is irreflexive. The irreflexivity of $G.\mathtt{hb}$ follows from Prop. F.14. Now, applying Prop. C.2, it suffices to show that $G.\mathtt{eco} \cup G.\mathtt{eco}; (G.\mathtt{sb} \setminus G.\mathtt{rmw}); G.\mathtt{hb}^?$ is irreflexive. First, $G.\mathtt{eco} = G_p.\mathtt{eco}$ is irreflexive because of SC-PER-LOC. Second, by Item 3 above, we have $G.\mathtt{eco}; (G.\mathtt{sb} \setminus G.\mathtt{rmw}); G.\mathtt{hb}^?; [G.\mathtt{RW}] \subseteq G_p.\mathtt{eco}; (G_p.\mathtt{sb} \cup G_p.\mathtt{fence}; G_p.\mathtt{hb}^*; G_p.\mathtt{ecoi}^?)$. By Prop. F.15, the latter relation is irreflexive.

ATOMICITY. By Prop. F.8, we have $G_p.\mathtt{rmw} \cap (G_p.\mathtt{rb}; G_p.\mathtt{mo}) = \emptyset$. Then, $G.\mathtt{rmw} \cap (G.\mathtt{rb}; G.\mathtt{mo}) = \emptyset$ immediately follows since $G.\mathtt{rmw} = G_p.\mathtt{rmw}$, $G.\mathtt{rb} = G_p.\mathtt{rb}$, and $G.\mathtt{mo} = G_p.\mathtt{mo}$.

SC. We show that $G.\mathtt{psc}$ is acyclic. Assuming no SC accesses, we have $G.\mathtt{psc} = R_1 \cup R_2$ where $R_1 = [G.\mathtt{F}^{\mathtt{sc}}]; G.\mathtt{hb}; G.\mathtt{eco}; G.\mathtt{hb}; [G.\mathtt{F}^{\mathtt{sc}}]$ and $R_2 = [G.\mathtt{F}^{\mathtt{sc}}]; G.\mathtt{hb}; [G.\mathtt{F}^{\mathtt{sc}}]$. Since R_2 is irreflexive and $R_2^+; R_1 \subseteq R_1$, it suffices to prove the acyclicity of R_1. To this end, we show that $G.\mathtt{eco}; G.\mathtt{hb}; [G.\mathtt{F}^{\mathtt{sc}}]; G.\mathtt{hb}; [G.\mathtt{RW}]$ is acyclic. Applying Prop. C.2, it suffices to show that $G.\mathtt{eco}; (G.\mathtt{sb} \setminus G.\mathtt{rmw}); G.\mathtt{hb}^?; [G.\mathtt{F}^{\mathtt{sc}}]; G.\mathtt{hb}; [G.\mathtt{RW}]$ is acyclic. Using Items 3 and 4 above (and applying several simple simplifications), it suffices to show that the following relation is acyclic:

$$G_p.\mathtt{eco}; (G_p.\mathtt{fence}; G_p.\mathtt{hb}^*)^?; G_p.\mathtt{sb}; [G_p.\mathtt{F}^{\mathtt{sync}}]; G_p.\mathtt{sb}; G_p.\mathtt{hb}^*; [G_p.\mathtt{RW}].$$

Using the definition of sync, this relation is equal to:

$$G_p.\mathtt{eco}; (G_p.\mathtt{fence}; G_p.\mathtt{hb}^*)^?; G_p.\mathtt{sync}; G_p.\mathtt{hb}^*; [G_p.\mathtt{RW}].$$

Its acyclicity then follows by Prop. F.13.

Next, we show that $G.\mathtt{pb}$ is acyclic. Suppose otherwise. Then, there are $a_1, \ldots, a_n$ such that $\langle a_i, a_{i+1}\rangle \in G.\mathtt{rfe}; G.\mathtt{pbi}^+$ for every $1 \leq i \leq n$ (where $a_{n+1} = a_1$). We show that $\langle a_i, a_{i+1}\rangle \in G_p.\mathtt{hb}^+$ for every $1 \leq i \leq n$ (which contradicts POWER-NO-THIN-AIR). Let $1 \leq i \leq n$, and let $b \in \mathtt{E}$ such that $\langle a_i, b\rangle \in G.\mathtt{rfe} = G_p.\mathtt{rfe}$ and $\langle b, a_{i+1}\rangle \in G.\mathtt{pbi}^+$. If $\langle b, a_{i+1}\rangle \in G_p.\mathtt{fence}$, then we are done since $G_p.\mathtt{rfe}, G_p.\mathtt{fence} \subseteq G_p.\mathtt{hb}$. Otherwise, it follows that $\langle b, a_{i+1}\rangle \in (G_p.\mathtt{deps} \cup G_p.\mathtt{addr}; G_p.\mathtt{sb} \cup G_p.\mathtt{psbloc})^+$. By Prop. F.6, we have $\langle a_i, a_{i+1}\rangle \in G_p.\mathtt{rfe}; G_p.\mathtt{ppo} \subseteq G.\mathtt{hb}^+$. □

Lemma H.3. *Given a program P without SC accesses, every outcome of $(\!|P|\!)$ under* Power *is an outcome of P under* RC11.

Proof. Given a full Power-consistent Power execution G_p of $(\!|P|\!)$, the compilation scheme (see Fig. 9) ensures that there exists some full execution G of P for which the properties of Lemma H.2 hold. Here we assumed that all RMW write attempts (`stwcx.`) succeed in the first attempt. Indeed, otherwise, one could always remove the RMW reads (`lwarx`) that precede the failed `stwcx.` attempts while preserving Power-consistency as well as the outcome of the execution. Now, Lemma H.2 ensures that G has the same outcome as G_p, G is weakRC11-consistent, and $G.\mathtt{pb}$ is acyclic. By Lemma G.1, either $G.\mathtt{sb} \cup G.\mathtt{rf}$ is acyclic (and NO-THIN-AIR holds) or P has undefined behavior under RC11. In any case, we obtain that the outcome of G_p is an outcome of P under RC11. □

I. Proofs for §7 (Correctness of Program Transformations)

In this appendix, we state (and outline the proofs of) the properties that ensure the soundness of the transformations discussed in §7. For this purpose, it is technically convenient to employ a different presentation of RMWs, that treat them as *single events* (like in C11). To this end, we consider RMW-*executions*, defined as the executions in §3, with the following exceptions:

- Labels in RMW-executions may also be $\mathtt{RMW}^o(x, v_r, v_w)$ where $o \in \{\mathtt{rlx}, \mathtt{acq}, \mathtt{rel}, \mathtt{acqrel}, \mathtt{sc}\}$. Both sets $G.\mathtt{R}$ and $G.\mathtt{W}$ include all events a with $\mathtt{typ}(a) = \mathtt{RMW}$, while $G.\mathtt{RMW}$ denotes the set of all events a with $\mathtt{typ}(a) = \mathtt{RMW}$.
- RMW-executions do not include an `rmw` component.

RC11-consistency for RMW-executions is also defined as for executions, with the following exceptions:

- $G.\mathtt{rb} \triangleq \mathtt{rf}^{-1}; \mathtt{mo} \setminus [\mathtt{E}]$.
- Instead of ATOMICITY we now require:

  $\mathtt{rf} \cap (\mathtt{mo}; \mathtt{mo}) = \emptyset$. (ATOMICITY-RMW)

The rest of the notions are defined for RMW-executions exactly as for executions above.

There exists a trivial one-to-one correspondence, denoted by $\sim$, between executions according to §3 and RMW-executions (the latter are obtained by collapsing `rmw` edges to single RMW events).

Proposition I.1. *Suppose that $G \sim G^{\mathtt{RMW}}$ for some execution G and* RMW*-execution $G^{\mathtt{RMW}}$. Then:*

- *G is* RC11*-consistent iff $G^{\mathtt{RMW}}$ is* RC11*-consistent.*
- *G is racy iff $G^{\mathtt{RMW}}$ is racy.*

Using this correspondence, we may define and prove the correctness of transformations on RMW-executions.

Lemma I.1 (Strengthening)**.** *Let G_{tgt} be an* RMW*-execution, obtained from an* RMW*-execution G_{src} by strengthening some access/fence modes ($G_{\mathsf{src}}.\mathtt{mod}(a) \sqsubseteq G_{\mathsf{tgt}}.\mathtt{mod}(a)$ for every $a \in G_{\mathsf{src}}.\mathtt{E}$). Then:*

- *If G_{tgt} is* RC11*-consistent, then so is G_{src}.*
- *If G_{tgt} is racy, then so is G_{src}.*

Proof. Easily follows from our definitions, because both properties are monotone with respect to the mode ordering. □

Lemma I.2 (Sequentialization)**.** *Let G_{tgt} be an* RMW*-execution, and let $\langle a, b\rangle \in \mathtt{sb} \setminus \mathtt{sb}; \mathtt{sb}$. Let G_{src} be the* RMW*-execution obtained from G by removing the* sb *edge $\langle a, b\rangle$. Then:*

- *If G_{tgt} is* RC11*-consistent, then so is G_{src}.*
- *If G_{tgt} is racy, then so is G_{src}.*

Proof. Easily follows from our definitions, because both properties are monotone with respect to sb. □

Next, to state the soundness of *deordering* transformations, we use the following definition of adjacency.

Definition I.1. Let R be a strict partial order on a set A. A pair $\langle a, b\rangle \in A \times A$ is called *R-adjacent* if the following hold for every $c \in A$:

- If $\langle c, a\rangle \in R$ then $\langle c, b\rangle \in R$.
- If $\langle b, c\rangle \in R$ then $\langle a, c\rangle \in R$.

Lemma I.3 (Non-load-store deordering)**.** *Let G_{tgt} be an* RMW*-execution, and let $a, b \in G_{\mathsf{tgt}}.\mathtt{E}$ such that $\langle a, b\rangle$ is $G_{\mathsf{tgt}}.\mathtt{sb}$-adjacent. Let G_{src} be the* RMW*-execution obtained from G_{src} by adding an* sb *edge $\langle a, b\rangle$. Suppose that the labels of a and b form a deorderable pair according to Table 1, except for the load-store deorderable pairs (*R; W*,* R; RMW*, and* RMW; W*). Then:*

- *If G_{tgt} is* RC11*-consistent, then so is G_{src}.*
- *If G_{tgt} is racy, then so is G_{src}.*

Proof. It is straightforward to verify that all components and derived relations in G_{src} are identical to those of G_{tgt} except for: $G_{\mathsf{src}}.\mathtt{sb} = G_{\mathsf{tgt}}.\mathtt{sb} \cup \{\langle a, b\rangle\}$ and $G_{\mathsf{src}}.\mathtt{hb} = G_{\mathsf{tgt}}.\mathtt{hb} \cup \{\langle a, b\rangle\}$. Then, the fact that G_{src} is RC11-consistent, easily follows from the fact that G_{tgt} is RC11-consistent. In particular, since a, b is not a load-store deorderable pair, assuming that G_{tgt} satisfies NO-THIN-AIR, we cannot have $\langle b, a\rangle \in (G_{\mathsf{src}}.\mathtt{sb} \cup G_{\mathsf{src}}.\mathtt{rf})^+$, so the additional sb edge $\langle a, b\rangle$ cannot close an $\mathtt{sb} \cup \mathtt{rf}$ cycle. Finally, since $G_{\mathsf{src}}.\mathtt{race} = G_{\mathsf{tgt}}.\mathtt{race}$, we have that G_{src} is racy if G_{tgt} is racy. □

Lemma I.4 (Load-store deordering)**.** *Let G_{tgt} be an* RMW*-execution, and let $a, b \in G_{\mathsf{tgt}}.\mathtt{E}$ such that $\langle a, b\rangle$ is $G_{\mathsf{tgt}}.\mathtt{sb}$-adjacent. Let G_{src} be the* RMW*-execution obtained from G_{src} by adding an* sb *edge $\langle a, b\rangle$. Suppose that the labels of a and b form a load-store deorderable pair (*R; W*,* R; RMW*, or* RMW; W*) according to Table 1. Then:*

- *If G_{tgt} is* RC11*-consistent, then G_{src} is* weakRC11*-consistent and $G_{\mathsf{src}}.\mathtt{pb}$ is acyclic.*
- *If G_{tgt} is racy, then so is G_{src}.*

Proof. The proof is similar to the proof of Lemma I.3. The fact that G_{src} is weakRC11-consistent follows from the fact that G_{tgt} is RC11-consistent. In addition, since $G_{\mathsf{src}}.\mathtt{pb} = G_{\mathsf{tgt}}.\mathtt{pb} \subseteq G_{\mathsf{tgt}}.\mathtt{sb} \cup G_{\mathsf{tgt}}.\mathtt{rf}$, assuming that G_{tgt} satisfies NO-THIN-AIR, we have that $G_{\mathsf{src}}.\mathtt{pb}$ is acyclic. □

Using Lemma G.1, one obtains the soundness of load-store deordering according to Table 1.

Notation I.1. For a binary relation R on a set A and an element $a \in A$, we denote by $R^{\uparrow}_a$ the set $\{b \in A \mid \langle b, a\rangle \in R\}$, and by $R^{\downarrow}_a$ the set $\{b \in A \mid \langle a, b\rangle \in R\}$.

Lemma I.5 (Read-read merging)**.** *Let G_{tgt} be an* RC11*-consistent* RMW*-execution. Let $a \in \mathtt{R} \setminus \mathtt{RMW}$, and let $a' \in \mathtt{E}$ such that $\langle a', a\rangle \in \mathtt{rf}$. Let $b \notin \mathtt{E}$, and let G_{src} be the* RMW*-execution satisfying:*

- $G_{\mathsf{src}}.\mathtt{E} = G_{\mathsf{tgt}}.\mathtt{E} \uplus \{b\}$.
- $G_{\mathsf{src}}.\mathtt{lab} = G_{\mathsf{tgt}}.\mathtt{lab} \cup \{b \mapsto G_{\mathsf{tgt}}.\mathtt{lab}(a)\}$.
- $G_{\mathsf{src}}.\mathtt{sb} = G_{\mathsf{tgt}}.\mathtt{sb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{sb}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{sb}^{\downarrow}_a)$.
- $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf} \cup \{\langle a', b\rangle\}$.
- $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo}$.

Then, G_{src} is RC11*-consistent, and it is racy if G_{tgt} is racy.*

Proof. By definition, G_{src} is complete, and ATOMICITY-RMW holds (since $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo}$ and $b \notin G_{\mathsf{src}}.\mathtt{R} \setminus G_{\mathsf{src}}.\mathtt{RMW}$). It is also easy to see that we have:

- $G_{\mathsf{src}}.\mathtt{eco} = G_{\mathsf{tgt}}.\mathtt{eco} \cup (G_{\mathsf{tgt}}.\mathtt{eco}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{eco}^{\downarrow}_a)$.
- $G_{\mathsf{src}}.\mathtt{hb} = G_{\mathsf{tgt}}.\mathtt{hb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{hb}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{hb}^{\downarrow}_a)$.

Hence, G_{src} satisfies COHERENCE. To see that NO-THIN-AIR holds, note that if we had $\langle b, a\rangle \in (G_{\mathsf{src}}.\mathtt{sb} \cup G_{\mathsf{src}}.\mathtt{rf})^+$, then we would have $\langle a, a\rangle \in (G_{\mathsf{tgt}}.\mathtt{sb} \cup G_{\mathsf{tgt}}.\mathtt{rf})^+$; and, similarly, if we had $\langle b, a'\rangle \in (G_{\mathsf{src}}.\mathtt{sb} \cup G_{\mathsf{src}}.\mathtt{rf})^+$, then we would have $\langle a, a'\rangle \in (G_{\mathsf{tgt}}.\mathtt{sb} \cup G_{\mathsf{tgt}}.\mathtt{rf})^+$. It remains to show that $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} \cup G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{F}}$ is acyclic. First, note that we have $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{F}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{F}}$. Now, if $G_{\mathsf{tgt}}.\mathtt{mod}(a) \neq \mathtt{sc}$, then we also have $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}$, and the claim follows since G_{tgt} satisfies SC. Otherwise, we have:

- $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}^{\uparrow}_{a} \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}^{\downarrow}_{a})$.

This implies that a $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} \cup G_{\mathsf{src}}.\mathtt{psc}_{\mathtt{F}}$ cycle would imply a $G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}} \cup G_{\mathsf{tgt}}.\mathtt{psc}_{\mathtt{F}}$ cycle. Finally, if $\langle c, b\rangle \in G_{\mathsf{src}}.\mathtt{race}$, then we have $\langle c, a\rangle \in G_{\mathsf{tgt}}.\mathtt{race}$. □

Lemma I.6 (Write-write merging)**.** *Let G_{tgt} be an* RC11*-consistent* RMW*-execution. Let $b \in \mathtt{W} \setminus \mathtt{RMW}$, $a \notin \mathtt{E}$, and $v \in \mathsf{Val}$. Let G_{src} be the* RMW*-execution satisfying:*

- $G_{\mathsf{src}}.\mathtt{E} = G_{\mathsf{tgt}}.\mathtt{E} \uplus \{a\}$.
- $G_{\mathsf{src}}.\mathtt{lab} = G_{\mathsf{tgt}}.\mathtt{lab} \cup \{a \mapsto \mathtt{W}^{G_{\mathsf{tgt}}.\mathtt{mod}(b)}(G_{\mathsf{tgt}}.\mathtt{loc}(b), v)\}$.
- $G_{\mathsf{src}}.\mathtt{sb} = G_{\mathsf{tgt}}.\mathtt{sb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{sb}^{\uparrow}_{b} \times \{a\}) \cup (\{a\} \times G_{\mathsf{tgt}}.\mathtt{sb}^{\downarrow}_{b})$.
- $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf}$.
- $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{mo}^{\uparrow}_{b} \times \{a\}) \cup (\{a\} \times G_{\mathsf{tgt}}.\mathtt{mo}^{\downarrow}_{b})$.

Then, G_{src} is RC11*-consistent, and it is racy if G_{tgt} is racy.*

Proof. By definition, G_{src} is complete. To see that ATOMICITY-RMW holds, note that we have $G_{\mathsf{src}}.\mathtt{mo}; G_{\mathsf{src}}.\mathtt{mo}; [\mathtt{RMW}] \subseteq G_{\mathsf{tgt}}.\mathtt{mo}; G_{\mathsf{tgt}}.\mathtt{mo} \cup (\{a\} \times G_{\mathsf{src}}.\mathtt{E})$, and that a has no outgoing $\mathtt{rf}$ edges. It is also easy to see that we have:

- $G_{\mathsf{src}}.\mathtt{eco} = G_{\mathsf{tgt}}.\mathtt{eco} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{eco}^{\uparrow}_{b} \times \{a\}) \cup (\{a\} \times G_{\mathsf{tgt}}.\mathtt{eco}^{\downarrow}_{b})$.
- $G_{\mathsf{src}}.\mathtt{hb} = G_{\mathsf{tgt}}.\mathtt{hb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{hb}^{\uparrow}_{b} \times \{a\}) \cup (\{a\} \times G_{\mathsf{tgt}}.\mathtt{hb}^{\downarrow}_{b})$.

Hence, G_{src} satisfies COHERENCE. To see that NO-THIN-AIR holds, note that if we had $\langle b, a\rangle \in (G_{\mathsf{src}}.\mathtt{sb} \cup G_{\mathsf{src}}.\mathtt{rf})^{+}$, then we would have $\langle b, b\rangle \in (G_{\mathsf{tgt}}.\mathtt{sb} \cup G_{\mathsf{tgt}}.\mathtt{rf})^{+}$. It remains to show that $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} \cup G_{\mathsf{src}}.\mathtt{psc}_{\mathtt{F}}$ is acyclic. First, note that we have $G_{\mathsf{src}}.\mathtt{psc}_{\mathtt{F}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathtt{F}}$. Now, if $G_{\mathsf{tgt}}.\mathtt{mod}(a) \neq \mathtt{sc}$, then we also have $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}$, and the claim follows since G_{tgt} satisfies SC. Otherwise, we have:

- $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}^{\uparrow}_{b} \times \{a\}) \cup (\{a\} \times G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}^{\downarrow}_{b})$.

This implies that a $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} \cup G_{\mathsf{src}}.\mathtt{psc}_{\mathtt{F}}$ cycle would imply a $G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}} \cup G_{\mathsf{tgt}}.\mathtt{psc}_{\mathtt{F}}$ cycle. Finally, if $\langle c, b\rangle \in G_{\mathsf{src}}.\mathtt{race}$, then we have $\langle c, a\rangle \in G_{\mathsf{tgt}}.\mathtt{race}$. □

Lemma I.7 (Write/RMW-read merging)**.** *Let G_{tgt} be an* RC11*-consistent* RMW*-execution. Let $a \in \mathtt{W}$ and $b \notin \mathtt{E}$. Let $o \in \mathsf{Ord}$, such that:*

- *If $\mathtt{typ}(a) = \mathtt{W}$ and $o = \mathtt{sc}$, then $\mathtt{mod}(a) = \mathtt{sc}$.*
- *If $\mathtt{typ}(a) = \mathtt{RMW}$, then $o \sqsubseteq \mathtt{mod}(a)$.*

Let G_{src} be the RMW*-execution satisfying:*

- $G_{\mathsf{src}}.\mathtt{E} = G_{\mathsf{tgt}}.\mathtt{E} \uplus \{b\}$.
- $G_{\mathsf{src}}.\mathtt{lab} = G_{\mathsf{tgt}}.\mathtt{lab} \cup \{b \mapsto \mathtt{R}^{o}(G_{\mathsf{tgt}}.\mathtt{loc}(a), G_{\mathsf{tgt}}.\mathtt{val}_{\mathtt{w}}(a))\}$.
- $G_{\mathsf{src}}.\mathtt{sb} = G_{\mathsf{tgt}}.\mathtt{sb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{sb}^{\uparrow}_{a} \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{sb}^{\downarrow}_{a})$.
- $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf} \cup \{\langle a, b\rangle\}$.
- $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo}$.

Then, G_{src} is RC11*-consistent, and it is racy if G_{tgt} is racy.*

Proof. Similar to the proof of Lemma I.5. □

Lemma I.8 (Write-RMW merging)**.** *Let G_{tgt} be an* RC11*-consistent* RMW*-execution. Let $b \in \mathtt{W} \setminus \mathtt{RMW}$, $a \notin \mathtt{E}$, $v \in \mathsf{Val}$, and $o \in \mathsf{Ord}$ such that $o_w = \mathtt{mod}(b)$. Let G_{src} be the* RMW*-execution satisfying:*

- $G_{\mathsf{src}}.\mathtt{E} = G_{\mathsf{tgt}}.\mathtt{E} \uplus \{a\}$.
- $G_{\mathsf{src}}.\mathtt{lab} = G_{\mathsf{tgt}}.\mathtt{lab}[b \mapsto \mathtt{RMW}^{o}(G_{\mathsf{tgt}}.\mathtt{loc}(b), v, G_{\mathsf{tgt}}.\mathtt{val}_{\mathtt{w}}(b))] \cup \{a \mapsto \mathtt{W}^{o_w}(G_{\mathsf{tgt}}.\mathtt{loc}(b), v)\}$.
- $G_{\mathsf{src}}.\mathtt{sb} = G_{\mathsf{tgt}}.\mathtt{sb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{sb}^{\uparrow}_{b} \times \{a\}) \cup (\{a\} \times G_{\mathsf{tgt}}.\mathtt{sb}^{\downarrow}_{b})$.
- $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf} \cup \{\langle a, b\rangle\}$.
- $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{mo}^{\uparrow}_{b} \times \{a\}) \cup (\{a\} \times G_{\mathsf{tgt}}.\mathtt{mo}^{\downarrow}_{b})$.

Then, G_{src} is RC11*-consistent, and it is racy if G_{tgt} is racy.*

Proof. By definition, G_{src} is complete. To see that ATOMICITY-RMW holds, note that we have $G_{\mathsf{src}}.\mathtt{mo}; G_{\mathsf{src}}.\mathtt{mo}; [\mathtt{RMW}] \subseteq G_{\mathsf{tgt}}.\mathtt{mo}; G_{\mathsf{tgt}}.\mathtt{mo} \cup (\{a\} \times G_{\mathsf{src}}.\mathtt{E}) \cup (G_{\mathsf{src}}.\mathtt{E} \times \{b\})$, and that a has only an $\mathtt{rf}$ edge to its immediate $G_{\mathsf{src}}.\mathtt{mo}$-successor b. The rest of the properties are proved as in the proof of Lemma I.6. □

Lemma I.9 (RMW-RMW merging)**.** *Let G_{tgt} be an* RC11*-consistent* RMW*-execution. Let $a \in \mathtt{E}$ with $\mathtt{lab}(a) = \mathtt{RMW}^o(x, v_r, v_w)$. Let $b \notin \mathtt{E}$ and $v \in \mathsf{Val}$, and let G_{src} be the* RMW*-execution satisfying:*

- $G_{\mathsf{src}}.\mathtt{E} = G_{\mathsf{tgt}}.\mathtt{E} \uplus \{b\}$.
- $G_{\mathsf{src}}.\mathtt{lab} = G_{\mathsf{tgt}}.\mathtt{lab}[a \mapsto \mathtt{RMW}^o(x, v_r, v)] \cup \{b \mapsto \mathtt{RMW}^o(x, v, v_w)\}$.
- $G_{\mathsf{src}}.\mathtt{sb} = G_{\mathsf{tgt}}.\mathtt{sb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{sb}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{sb}^{\downarrow}_a)$.
- $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf} \cup \{\langle a, b\rangle\}$.
- $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{mo}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{mo}^{\downarrow}_a)$.

Then, G_{src} is RC11*-consistent, and it is racy if G_{tgt} is racy.*

Lemma I.10 (Fence-fence merging)**.** *Let G_{tgt} be an* RC11*-consistent* RMW*-execution. Let $a \in \mathtt{F}$, $b \notin \mathtt{E}$, and let G_{src} be the* RMW*-execution satisfying:*

- $G_{\mathsf{src}}.\mathtt{E} = G_{\mathsf{tgt}}.\mathtt{E} \uplus \{b\}$.
- $G_{\mathsf{src}}.\mathtt{lab} = G_{\mathsf{tgt}}.\mathtt{lab} \cup \{b \mapsto G_{\mathsf{tgt}}.\mathtt{lab}(a)\}$.
- $G_{\mathsf{src}}.\mathtt{sb} = G_{\mathsf{tgt}}.\mathtt{sb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{sb}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{sb}^{\downarrow}_a)$.
- $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf}$.
- $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo}$.

Then, G_{src} is RC11*-consistent, and it is racy if G_{tgt} is racy.*

Proof. By definition, G_{src} is complete, and ATOMICITY-RMW holds since $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf}$ and $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo}$. It is also easy to see that we have $G_{\mathsf{src}}.\mathtt{eco} = G_{\mathsf{tgt}}.\mathtt{eco}$ and:

- $G_{\mathsf{src}}.\mathtt{hb} = G_{\mathsf{tgt}}.\mathtt{hb} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{hb}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{hb}^{\downarrow}_a)$.

Hence, G_{src} satisfies COHERENCE. To see that NO-THIN-AIR holds, note that if we had $\langle b, a\rangle \in G_{\mathsf{src}}.\mathtt{sb} \cup G_{\mathsf{src}}.\mathtt{rf}$, then we would have $\langle a, a\rangle \in G_{\mathsf{tgt}}.\mathtt{sb} \cup G_{\mathsf{tgt}}.\mathtt{rf}$. It remains to show that $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}}$ is acyclic. If $G_{\mathsf{tgt}}.\mathtt{mod}(a) \neq \mathtt{sc}$, then we have $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}$ and $G_{\mathsf{src}}.\mathtt{psc}_{\mathtt{F}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathtt{F}}$, and the claim follows since G_{tgt} satisfies SC. Otherwise, we have:

- $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}} \cup (G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}{}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}}{}^{\downarrow}_a)$.
- $G_{\mathsf{src}}.\mathtt{psc}_{\mathtt{F}} = G_{\mathsf{tgt}}.\mathtt{psc}_{\mathtt{F}} \cup \{\langle a, b\rangle\} \cup (G_{\mathsf{tgt}}.\mathtt{psc}_{\mathtt{F}}{}^{\uparrow}_a \times \{b\}) \cup (\{b\} \times G_{\mathsf{tgt}}.\mathtt{psc}_{\mathtt{F}}{}^{\downarrow}_a)$.

This implies that a $G_{\mathsf{src}}.\mathtt{psc}_{\mathrm{base}} \cup G_{\mathsf{src}}.\mathtt{psc}_{\mathtt{F}}$ cycle would imply a $G_{\mathsf{tgt}}.\mathtt{psc}_{\mathrm{base}} \cup G_{\mathsf{tgt}}.\mathtt{psc}_{\mathtt{F}}$ cycle. Finally, $G_{\mathsf{src}}.\mathtt{race} = G_{\mathsf{tgt}}.\mathtt{race}$, so G_{src} is racy if G_{tgt} is racy. □

Soundness of register promotion is proved in two steps. First, we show that if all accesses to some location are in one thread, then they can be safely weakened to non-atomic accesses. Second, we show that these non-atomic accesses can be safely removed (replaced by register assignments at the program level).

Lemma I.11 (Register promotion-a)**.** *Let G_{tgt} be an* RC11*-consistent* RMW*-execution. Suppose that all accesses to some location x are related by $G_{\mathsf{tgt}}.\mathtt{sb}$. Let G_{src} be the* RMW*-execution obtained by strengthening the accesses mode of all accesses to x to* sc*. Then, G_{src} is* RC11*-consistent, and it is racy if G_{tgt} is racy.*

Proof. By definition, we have $G_{\mathsf{src}}.\mathtt{c} = G_{\mathsf{tgt}}.\mathtt{c}$ for $\mathtt{c} \in \{\mathtt{sb}, \mathtt{rf}, \mathtt{mo}, \mathtt{eco}\}$. It is also easy to see that $G_{\mathsf{src}}.\mathtt{hb} = G_{\mathsf{tgt}}.\mathtt{hb}$. Hence, G_{src} is complete, and ATOMICITY,COHERENCE,NO-THIN-AIR hold for G_{src} since they hold for G_{tgt}. To see that $G_{\mathsf{src}}.\mathtt{psc}$ is acyclic, it suffices to note that $G_{\mathsf{src}}.\mathtt{psc} \subseteq G_{\mathsf{tgt}}.\mathtt{psc} \cup G_{\mathsf{tgt}}.\mathtt{sb}$ (acyclicity of $G_{\mathsf{tgt}}.\mathtt{psc} \cup G_{\mathsf{tgt}}.\mathtt{sb}$ follows from the acyclicity of $G_{\mathsf{tgt}}.\mathtt{psc}$ since $\mathtt{psc}; \mathtt{sb}; \mathtt{psc} \subseteq \mathtt{psc}^+$ in every execution). Finally, if $\langle a, b\rangle \in G_{\mathsf{tgt}}.\mathtt{race}$ and $\mathtt{na} \in \{G_{\mathsf{tgt}}.\mathtt{mod}(a), G_{\mathsf{tgt}}.\mathtt{mod}(b)\}$, then the same holds in G_{src}: we must have $\mathtt{loc}(a) \neq x$ if $\langle a, b\rangle \notin G_{\mathsf{tgt}}.\mathtt{hb} \cup (G_{\mathsf{tgt}}.\mathtt{hb})^{-1}$. □

Lemma I.12 (Register promotion-b)**.** *Let G_{tgt} be an* RC11*-consistent* RMW*-execution. Let $x \in \mathsf{Loc}$ and let $X = \{b \in \mathtt{E} \mid \mathtt{loc}(b) = x\}$. Suppose that all accesses in X are related by $G_{\mathsf{tgt}}.\mathtt{sb}$. Let $a \notin \mathtt{E}$, let G_{src} be an* RMW*-execution satisfying:*

- $G_{\mathsf{src}}.\mathtt{E} = G_{\mathsf{tgt}}.\mathtt{E} \uplus \{a\}$.

- $G_{\mathsf{src}}.\mathtt{lab} = G_{\mathsf{tgt}}.\mathtt{lab} \cup \{a \mapsto L\}$ *where* L *is some access label with mode* `na` *and location* x.
- $G_{\mathsf{src}}.\mathtt{sb} \supset G_{\mathsf{tgt}}.\mathtt{sb}$ *and every ecent in* X *is* $G_{\mathsf{src}}.\mathtt{sb}$*-related to* a.
- $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf}$ *if* $G_{\mathsf{src}}.\mathtt{typ}(a) = \mathtt{W} \setminus \mathtt{RMW}$,
  and otherwise $G_{\mathsf{src}}.\mathtt{rf} = G_{\mathsf{tgt}}.\mathtt{rf} \cup \{\langle \max_{G_{\mathsf{src}}.\mathtt{sb}} G_{\mathsf{src}}.\mathtt{sb}^{\uparrow}_{a}, a\rangle\}$.
- $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo}$ *if* $G_{\mathsf{src}}.\mathtt{typ}(a) = \mathtt{R} \setminus \mathtt{RMW}$,
  and otherwise $G_{\mathsf{src}}.\mathtt{mo} = G_{\mathsf{tgt}}.\mathtt{mo} \cup (G_{\mathsf{src}}.\mathtt{sb}^{\uparrow}_{a} \times \{a\}) \cup (\{a\} \times G_{\mathsf{src}}.\mathtt{sb}^{\downarrow}_{a})$.

Then, G_{src} *is* RC11*-consistent, and it is racy if* G_{tgt} *is racy.*

Proof. Easily follows from our definitions. □

J. Proofs for §8 (Programming Guarantees)

Theorem 4. *If in all SC-consistent executions of a program* P*, every race* $\langle a, b\rangle$ *has* $\mathtt{mod}(a) = \mathtt{mod}(b) = \mathtt{sc}$*, then the outcomes of* P *under* RC11 *coincide with those under SC.*

Proof. Let P be a program, and suppose that every race $\langle a, b\rangle$ in some SC-consistent execution of P has $\mathtt{mod}(a) = \mathtt{mod}(b) = \mathtt{sc}$. We prove that P has no weak behaviors. Suppose toward a contradiction that there exists an execution G of P that is RC11-consistent but not SC-consistent. (Note that if P has undefined behavior under RC11, then there exists a racy RC11-consistent execution of P, and our assumption ensures that this execution is not SC-consistent.)

We call an execution G' is a *prefix* of an execution G if it is obtained by restricting G to a set E of events that contains the set $\mathtt{E}_0$ of initialization events, and is closed with respect to $G.\mathtt{sb} \cup G.\mathtt{rf}$ ($a \in E$ whenever $b \in E$ and $\langle a, b\rangle \in G.\mathtt{sb} \cup G.\mathtt{rf}$). It is easy to show that G' is RC11-consistent, provided that G is RC11-consistent.

Notation J.1. For an execution G, $G.\mathtt{rf}|_{\mathtt{sc}}$ denotes the restriction of $G.\mathtt{rf}$ to SC accesses ($G.\mathtt{rf}|_{\mathtt{sc}} = [G.\mathtt{E}^{\mathtt{sc}}]; G.\mathtt{rf}; [G.\mathtt{E}^{\mathtt{sc}}]$). A similar notation is used for $G.\mathtt{mo}$ and $G.\mathtt{rb}$.

For a set of events E, let $\Pi(E)$ denote the set of all pairs $\langle a, b\rangle \in E \times E$ of conflicting events, such that $\{G.\mathtt{mod}(a), G.\mathtt{mod}(b)\} \neq \{\mathtt{sc}\}$ and $\langle a, b\rangle, \langle b, a\rangle \notin (G.\mathtt{sb} \cup G.\mathtt{rf}|_{\mathtt{sc}})^+$. Let $a_1, \ldots, a_n$ be an enumeration of $\mathtt{E} \setminus \mathtt{E}_0$ that respects $G.\mathtt{sb} \cup G.\mathtt{rf}$ (that is, $i < j$ whenever $\langle a_i, a_j\rangle \in G.\mathtt{sb} \cup G.\mathtt{rf}$). For every $1 \leq i \leq n$, let $E_i = \mathtt{E}_0 \cup \{a_1, \ldots, a_i\}$ and $G_i = G|_{E_i}$. Since the G_i's are all prefixes of G, all of them are RC11-consistent.

Claim: For every $1 \leq i \leq n$, if $\Pi(E_i) = \emptyset$ then G_i is SC-consistent.

Proof: Suppose that $\Pi(E_i) = \emptyset$. Since G satisfies COHERENCE, it follows that:

- $G_i.\mathtt{rf} \subseteq (G.\mathtt{sb} \cup G.\mathtt{rf}|_{\mathtt{sc}})^+$.
- $G_i.\mathtt{mo} \subseteq (G.\mathtt{sb} \cup G.\mathtt{rf}|_{\mathtt{sc}})^+ \cup G.\mathtt{mo}|_{\mathtt{sc}}$.
- $G_i.\mathtt{rb} \subseteq (G.\mathtt{sb} \cup G.\mathtt{rf}|_{\mathtt{sc}})^+ \cup G.\mathtt{rb}|_{\mathtt{sc}}$.

Hence, we have $G_i.\mathtt{sb} \cup G_i.\mathtt{rf} \cup G_i.\mathtt{mo} \cup G_i.\mathtt{rb} \subseteq R^+$, where $R = G.\mathtt{sb} \cup G.\mathtt{rf}|_{\mathtt{sc}} \cup G.\mathtt{mo}|_{\mathtt{sc}} \cup G.\mathtt{rb}|_{\mathtt{sc}}$. Since G satisfies the SC condition, we have that R is acyclic, and so G_i is SC-consistent (ATOMICITY holds since it holds for G). □

Now, since G is not SC-consistent, we have $\Pi(G.\mathtt{E}) \neq \emptyset$. Let $k = \min\{i \mid \Pi(E_i) \neq \emptyset\}$. Then, $\Pi(E_{k-1}) = \emptyset$ (and so, G_{k-1} is SC-consistent), and there exists some $j < k$, such that a_j and a_k are conflicting, $\{G.\mathtt{mod}(a_j), G.\mathtt{mod}(a_k)\} \neq \{\mathtt{sc}\}$, and $\langle a_j, a_k\rangle, \langle a_k, a_j\rangle \notin (G.\mathtt{sb} \cup G.\mathtt{rf}|_{\mathtt{sc}})^+$. Let $B = \{b \in E_k \mid \langle b, a_k\rangle \in G.\mathtt{sb}\}$. Since $\langle a_j, a_k\rangle \notin (G.\mathtt{sb} \cup G.\mathtt{rf}|_{\mathtt{sc}})^+$, and $G_{k-1}.\mathtt{rf} \subseteq (G.\mathtt{sb} \cup G.\mathtt{rf}|_{\mathtt{sc}})^+$, we have $\langle a_j, b\rangle \notin (G.\mathtt{sb} \cup G.\mathtt{rf})^+$ for every event $b \in B$. Let $x = \mathtt{loc}(a_k)$, and consider two cases:

- $\mathtt{typ}(a_k) = \mathtt{W}$:

  Claim: $\langle a_j, a_k\rangle \in G_k.\mathtt{race}$.

  Proof: Clearly, we have $\langle a_k, a_j\rangle \notin (G_k.\mathtt{sb} \cup G_k.\mathtt{rf})^+$ (a_k has no outgoing `sb` and `rf` edges in G_k). In addition, we have $\langle a_j, a_k\rangle \notin (G_k.\mathtt{sb} \cup G_k.\mathtt{rf})^+$ (otherwise, $\langle a_j, b\rangle \in (G.\mathtt{sb} \cup G.\mathtt{rf})^+$ for some $b \in B$). □

  Claim: G_k is not SC-consistent.

Proof: Since $\langle a_j, a_k\rangle \in G_k.\texttt{race}$ and $\{G.\texttt{mod}(a_j), G.\texttt{mod}(a_k)\} \neq \{\texttt{sc}\}$, the claim follows from our assumption. □

Claim: $a_k \notin G.\texttt{At}$.

Proof: Suppose otherwise, and let $b \in G.\texttt{E}$ such that $\langle b, a_k\rangle \in \texttt{rmw}$. Since G_k is not SC-consistent, but G_{k-1} is SC-consistent, it must be the case that $\langle a_k, c\rangle \in G.\texttt{mo}$ and $\langle c, a_k\rangle \in (G.\texttt{sb} \cup G.\texttt{rf} \cup G.\texttt{mo} \cup G.\texttt{rb})^+$ for some $c \in E_{k-1}$. Let $d \in E_{k-1}$ such that $\langle c, d\rangle \in (G.\texttt{sb} \cup G.\texttt{rf} \cup G.\texttt{mo} \cup G.\texttt{rb})^*$ and $\langle d, a_k\rangle \in G.\texttt{sb} \cup G.\texttt{mo} \cup G.\texttt{rb}$. Then, we also have $\langle c, d\rangle \in (G_{k-1}.\texttt{sb} \cup G_{k-1}.\texttt{rf} \cup G_{k-1}.\texttt{mo} \cup G_{k-1}.\texttt{rb})^*$. If $\langle d, a_k\rangle \in G.\texttt{mo} \cup G.\texttt{rb}$, then we obtain $\langle d, c\rangle \in G.\texttt{mo} \cup G.\texttt{rb}$, and so $\langle d, c\rangle \in G_{k-1}.\texttt{mo} \cup G_{k-1}.\texttt{rb}$, which contradicts the fact that G_{k-1} is SC-consistent. Otherwise, $\langle d, a_k\rangle \in G.\texttt{sb}$. It follows that $\langle d, b\rangle \in G.\texttt{sb}^?$. Now, COHERENCE ensures that $G.\texttt{rmw} \subseteq G.\texttt{rb}$, and it follows that $\langle b, c\rangle \in G.\texttt{rb}$. Hence, $\langle b, c\rangle \in G_{k-1}.\texttt{rb}$, which again contradicts the fact that G_{k-1} is SC-consistent. □

Let G'_k be the extension of G_{k-1} with the event a_k (with the same label as in G_k), the $\texttt{sb}$ edges of G_k, and the $\texttt{mo}$ edges $\{\langle a, a_k\rangle \mid a \in G_{k-1}.\texttt{W}_x\}$. It is easy to see that G'_k is SC-consistent as well (in particular, it is important here that $a_k \notin G.\texttt{At}$). Except for $\texttt{mo}$, it is identical to G_k, and so it is an execution of P and $\langle a_j, a_k\rangle \in G'_k.\texttt{race}$. Since $\{G.\texttt{mod}(a_j), G.\texttt{mod}(a_k)\} \neq \{\texttt{sc}\}$, this contradicts our assumption.

- $\texttt{typ}(a_k) = \texttt{R}$:

  In this case, we must have $\texttt{typ}(a_j) = \texttt{W}$. Let

  $$E = \{a \in G.\texttt{E} \mid \langle a, a_k\rangle \in (G.\texttt{sb} \cup G_{k-1}.\texttt{rf})^* \vee \langle a, a_j\rangle \in (G.\texttt{sb} \cup G_{k-1}.\texttt{rf})^*\}.$$

  Let G' be the restriction of G_k to the events in E. Since $G'|_{E\setminus\{a_k\}}$ is a prefix of G_{k-1}, it is SC-consistent. Let $c = \max_{G.\texttt{mo}} G'.\texttt{W}_x$, and consider two cases.

  - $c \neq a_j$:

    Let G'' be the execution obtained from G' by (i) modifying the value read at a_k to $\texttt{val}_\texttt{w}(c)$, and (ii) adding the reads-from edge $\langle c, a\rangle$. It is easy to see that G'' is SC-consistent, and Assumption B.1 ensures that it is an execution of P. Additionally, $\langle a_j, a_k\rangle \notin (G''.\texttt{sb} \cup G''.\texttt{rf})^+$ (there are no outgoing $\texttt{sb}$ and $\texttt{rf}$ edges from a_j in G''), and so, $\langle a_j, a_k\rangle \in G''.\texttt{race}$. Since $\{G.\texttt{mod}(a_j), G.\texttt{mod}(a_k)\} \neq \{\texttt{sc}\}$, this contradicts our assumption.

  - $c = a_j$:

    Let d be the immediate $G.\texttt{mo}$-predecessor of c, and let G'' be the execution obtained from G' by (i) modifying the value read at a_k to $\texttt{val}_\texttt{w}(d)$, and (ii) adding the reads-from edge $\langle d, a\rangle$. Again, it is easy to see that G'' is SC-consistent, and Assumption B.1 ensures that it is an execution of P. As in the previous case we obtain a contradiction to our assumption. □

Theorem 5. *Let G be an* RC11*-consistent execution. Suppose that for every two distinct shared locations x and y, $[\texttt{E}_x]; \texttt{sb}; [\texttt{E}_y] \subseteq \texttt{sb}; [\texttt{F}^{\texttt{sc}}]; \texttt{sb}$. Then, G is SC-consistent.*

Proof. It suffices to show that $\texttt{sb} \cup \texttt{ecoe}$ is acyclic (where $\texttt{ecoe} \triangleq \texttt{eco} \setminus \texttt{sb}$). Consider a cycle in $\texttt{sb} \cup \texttt{ecoe}$ of a minimal length. Cycles with at most one $\texttt{ecoe}$ edge are ruled out by COHERENCE. Hence, our cycle must have at least two $\texttt{ecoe}$ edges. Let $a_1, b_1, a_2, b_2, \ldots, a_n, b_n \in \texttt{E}$ (where $n \geq 2$) such that $\langle a_i, b_i\rangle \in \texttt{ecoe}$ and $\langle b_i, a_{i+1}\rangle \in \texttt{sb}$ for every $1 \leq i \leq n$ (where we take a_{n+1} to be a_1). The events $a_1, b_1, \ldots, a_n, b_n$ are all accesses to shared locations (since $\langle a_i, b_i\rangle \in \texttt{ecoe} \subseteq =_{\texttt{loc}}$ for every $1 \leq i \leq n$). In addition, we have $\texttt{loc}(b_i) \neq \texttt{loc}(a_{i+1})$ for every $1 \leq i \leq n$ (otherwise we would have $\langle a_i, a_{i+1}\rangle \in \texttt{ecoe}; \texttt{sb}|_{\texttt{loc}} \subseteq \texttt{ecoe}$, which contradicts the minimality of the cycle). Therefore, our assumption entails that there exist $f_1, \ldots, f_n \in \texttt{F}^{\texttt{sc}}$ such that $\langle b_i, f_i\rangle \in \texttt{sb}$ and $\langle f_i, a_{i+1}\rangle \in \texttt{sb}$ for every $1 \leq i \leq n$. It follows that $\langle f_i, f_{i+1}\rangle \in [\texttt{F}^{\texttt{sc}}]; \texttt{sb}; \texttt{ecoe}; \texttt{sb}; [\texttt{F}^{\texttt{sc}}] \subseteq \texttt{psc}_\texttt{F}$ for every $1 \leq i \leq n$ (where we take f_{n+1} to be f_1). This contradicts the fact that G satisfies the SC constraint. □

K. Proofs for §4 (Compilation to x86-TSO)

The following proposition is useful in our proof below:

Proposition K.1. *The following hold in every* TSO*-consistent* TSO *execution:*

- $\mathtt{rb};\mathtt{mo};\mathtt{hb}^?;\mathtt{rfe};\mathtt{hb}^?$ *is irreflexive.*
- $\mathtt{rb};\mathtt{mo};\mathtt{hb}^?;[\mathtt{RMW}\cup\mathtt{F}];\mathtt{hb}^?$ *is irreflexive.*
- *The relation* $T = [\mathtt{R}];\mathtt{hb}\cup\mathtt{hb}^?;\mathtt{rfe};\mathtt{hb}^?\cup\mathtt{hb};[\mathtt{F}]\cup[\mathtt{F}];\mathtt{hb}\cup\mathtt{mo}\cup\mathtt{rb}$ *is acyclic.*

Proof. The first two are straightforward. We prove the third claim. Consider a cycle $\langle a_1,\ldots,a_n\rangle$ in T with a minimal number of events $a_i \in \mathtt{W}\cup\mathtt{F}$: The minimality of the cycle entails that at most two events in $\mathtt{W}\cup\mathtt{F}$ participate in this cycle (otherwise, it can be shortened since $\mathtt{mo}$ is total on $\mathtt{W}\cup\mathtt{F}$). Since $\mathtt{hb}$ and $\mathtt{rb};\mathtt{hb}$ are irreflexive, there must be at least two such events in the cycle. Hence, we have exactly two indices $1 \le i < j \le n$ such that $a_i, a_j \in \mathtt{W}\cup\mathtt{F}$. W.l.o.g., we may assume that $\langle a_i, a_j\rangle \in \mathtt{mo}$. Since the rest of the events are not in $\mathtt{W}\cup\mathtt{F}$, and $\mathtt{mo};\mathtt{hb}$ is irreflexive, we obtain that $\langle a_j, a_i\rangle \in ([\mathtt{R}];\mathtt{hb}\cup\mathtt{hb}^?;\mathtt{rfe};\mathtt{hb}^?\cup\mathtt{hb};[\mathtt{F}];\mathtt{hb}^?)^+;\mathtt{rb}$. Since $[\mathtt{W}];([\mathtt{R}];\mathtt{hb}\cup\mathtt{hb}^?;\mathtt{rfe};\mathtt{hb}^?\cup\mathtt{hb};[\mathtt{F}]\cup[\mathtt{F}];\mathtt{hb})^+ \subseteq \mathtt{hb}^?;\mathtt{rfe};\mathtt{hb}^?\cup\mathtt{hb};[\mathtt{F}];\mathtt{hb}^?$, we obtain that $\langle a_j, a_i\rangle \in (\mathtt{hb}^?;\mathtt{rfe};\mathtt{hb}^?\cup\mathtt{hb};[\mathtt{F}];\mathtt{hb}^?);\mathtt{rb}$. This contradicts the previous claims. □

Lemma K.1. *Let* G *be an* RMW*-execution satisfying* $G.\mathtt{F}^{\neq\mathtt{sc}} = \emptyset$, $G.\mathtt{mod}(a) \sqsupset \mathtt{rlx}$ *for every* $a \in \mathtt{E}$*, and* $G.\mathtt{mod}(a) \sqsupset \mathtt{acqrel}$ *for every* $a \in \mathtt{RMW}$*. Let* G_t *be a* TSO *execution. Suppose that there exists an injective function* $f : (G.\mathtt{W}^{\mathtt{sc}} \setminus G.\mathtt{RMW}) \to \mathbb{N}$ *assigning a fresh event* $f(a) \notin G.\mathtt{E}$ *to every* $a \in G.\mathtt{W}^{\mathtt{sc}} \setminus G.\mathtt{RMW}$*, such that the following hold:*

- $G_t.\mathtt{E} = G.\mathtt{E} \cup f[G.\mathtt{W}^{\mathtt{sc}} \setminus G.\mathtt{RMW}]$.
- $G_t.\mathtt{lab}(a) = \mathtt{R}(G.\mathtt{loc}(a), G.\mathtt{val}_\mathtt{r}(a))$ *for every* $a \in G.\mathtt{R} \setminus G.\mathtt{RMW}$.
- $G_t.\mathtt{lab}(a) = \mathtt{W}(G.\mathtt{loc}(a), G.\mathtt{val}_\mathtt{w}(a))$ *for every* $a \in G.\mathtt{W} \setminus G.\mathtt{RMW}$.
- $G_t.\mathtt{lab}(a) = \mathtt{RMW}(G.\mathtt{loc}(a), G.\mathtt{val}_\mathtt{r}(a), G.\mathtt{val}_\mathtt{w}(a))$ *for every* $a \in G.\mathtt{RMW}$.
- $G_t.\mathtt{lab}(a) = \mathtt{F}$ *for every* $a \in G.\mathtt{F}^{\mathtt{sc}} \cup f[G.\mathtt{W}^{\mathtt{sc}} \setminus G.\mathtt{RMW}]$.
- $G_t.\mathtt{sb} = G.\mathtt{sb} \cup \{\langle b, f(a)\rangle \mid \langle b, a\rangle \in G.\mathtt{sb}^?;[G.\mathtt{W}^{\mathtt{sc}}\setminus G.\mathtt{RMW}]\} \cup \{\langle f(a), b\rangle \mid \langle a, b\rangle \in [G.\mathtt{W}^{\mathtt{sc}}\setminus G.\mathtt{RMW}];G.\mathtt{sb}\}$.
- $G.\mathtt{rf} = G_t.\mathtt{rf}$.
- $G.\mathtt{mo} = \bigcup_{x\in\mathsf{Loc}}[G_t.\mathtt{W}_x];G_t.\mathtt{mo};[G_t.\mathtt{W}_x]$.

Then:

- G *and* G_t *have the same outcome.*
- *If* G_t *is* TSO*-consistent, then* G *is* RC11*-consistent.*

Proof. The first claim easily follows from our definitions. Suppose that G_t is TSO-consistent. We show that G is RC11-consistent. Clearly, it is complete (since $G.\mathtt{R} = G_t.\mathtt{R}$ and $G.\mathtt{rf} = G_t.\mathtt{rf}$).

COHERENCE. Easily follows using Prop. 1 from the fact that $G_t.\mathtt{hb}$, $G_t.\mathtt{mo};G_t.\mathtt{hb}$ and $G_t.\mathtt{rb};G_t.\mathtt{hb}$ are all irreflexive (note that $G.\mathtt{rf} \subseteq G.\mathtt{hb} \subseteq G_t.\mathtt{hb}$, $G.\mathtt{mo} \subseteq G_t.\mathtt{mo}$, and $G.\mathtt{rb} \subseteq G_t.\mathtt{rb}$).

ATOMICITY-RMW. Trivially follows from the fact that $G_t.\mathtt{rb};G_t.\mathtt{mo}$ is irreflexive.

SC. Let $\mathtt{psc}'_{\mathrm{base}} = \left([G.\mathtt{E}^{\mathtt{sc}}]\cup[G.\mathtt{F}^{\mathtt{sc}}];G.\mathtt{hb}^?\right);(G.\mathtt{hb}\cup G.\mathtt{mo}\cup G.\mathtt{rb});\left([G.\mathtt{E}^{\mathtt{sc}}]\cup G.\mathtt{hb}^?;[G.\mathtt{F}^{\mathtt{sc}}]\right)$. We show that $\mathtt{psc}'_{\mathrm{base}}\cup G.\mathtt{psc}_\mathtt{F} \subseteq T^+$, where T is the relation defined in Prop. K.1. This implies that SC holds (as well as that Batty *et al.*'s [5] condition holds). To prove $\mathtt{psc}'_{\mathrm{base}}\cup G.\mathtt{psc}_\mathtt{F} \subseteq T^+$, note that the following hold (in G):

- $[\mathtt{F}];(\mathtt{psc}'_{\mathrm{base}}\cup\mathtt{psc}_\mathtt{F});[\mathtt{F}] \subseteq \mathtt{psc}_\mathtt{F} \subseteq [\mathtt{F}];\mathtt{hb}\cup[\mathtt{F}];\mathtt{hb};\mathtt{eco};\mathtt{hb};[\mathtt{F}] \subseteq ([\mathtt{F}];(\mathtt{sb}\cup\mathtt{rf})^+\cup(\mathtt{sb}\cup\mathtt{rf})^+;[\mathtt{F}]\cup\mathtt{mo}\cup\mathtt{rb})^+$.
- $[\mathtt{F}];(\mathtt{psc}'_{\mathrm{base}}\cup\mathtt{psc}_\mathtt{F});[\mathtt{R}\cup\mathtt{W}] = [\mathtt{F}];\mathtt{psc}'_{\mathrm{base}};[\mathtt{R}\cup\mathtt{W}] \subseteq [\mathtt{F}];\mathtt{hb};(\mathtt{mo}\cup\mathtt{rb})^?$
- $[\mathtt{R}\cup\mathtt{W}];(\mathtt{psc}'_{\mathrm{base}}\cup\mathtt{psc}_\mathtt{F});[\mathtt{F}] = [\mathtt{R}\cup\mathtt{W}];\mathtt{psc}'_{\mathrm{base}};[\mathtt{F}] \subseteq [\mathtt{R}\cup\mathtt{W}];(\mathtt{hb}\cup\mathtt{mo}\cup\mathtt{rb});\mathtt{hb}^?;[\mathtt{F}] \subseteq (\mathtt{mo}\cup\mathtt{rb})^?;\mathtt{hb};[\mathtt{F}]$
- $[\mathtt{R}];(\mathtt{psc}'_{\mathrm{base}}\cup\mathtt{psc}_\mathtt{F});[\mathtt{R}\cup\mathtt{W}] = [\mathtt{R}];\mathtt{psc}'_{\mathrm{base}};[\mathtt{R}\cup\mathtt{W}] \subseteq [\mathtt{R}];\mathtt{hb}\cup\mathtt{mo}\cup\mathtt{rb}$
- $[\mathtt{W}];(\mathtt{psc}'_{\mathrm{base}}\cup\mathtt{psc}_\mathtt{F});[\mathtt{R}\cup\mathtt{W}] = [\mathtt{W}^{\mathtt{sc}}];(\mathtt{hb}\cup\mathtt{mo}\cup\mathtt{rb});[\mathtt{R}^{\mathtt{sc}}\cup\mathtt{W}^{\mathtt{sc}}] \subseteq [\mathtt{W}^{\mathtt{sc}}];\mathtt{sb};[\mathtt{R}^{\mathtt{sc}}]\cup(\mathtt{sb}\cup\mathtt{rf})^*;\mathtt{rfe};(\mathtt{sb}\cup\mathtt{rf})^*\cup\mathtt{mo}\cup\mathtt{rb}$

Using these facts, since we have $G.\mathtt{F} \subseteq G_t.\mathtt{F}$, $G.\mathtt{R} \subseteq G_t.\mathtt{R}$, $G.\mathtt{hb} \subseteq G_t.\mathtt{hb}$, $(G.\mathtt{sb} \cup G.\mathtt{rf})^+ \subseteq G_t.\mathtt{hb}$, $G.\mathtt{mo} \subseteq G_t.\mathtt{mo}$, $G.\mathtt{rb} \subseteq G_t.\mathtt{rb}$, $G.\mathtt{rfe} \subseteq G_t.\mathtt{rfe}$, and $[G.\mathtt{W}^{\mathtt{sc}}]; G.\mathtt{sb}; [G.\mathtt{R}^{\mathtt{sc}}] \subseteq G_t.\mathtt{sb}; [G_t.\mathtt{F}]; G_t.\mathtt{sb} \subseteq G_t.\mathtt{hb}; [G_t.\mathtt{F}]; G_t.\mathtt{hb}$, it immediately follows that $\mathtt{psc}'_{\mathrm{base}} \cup G.\mathtt{psc}_{\mathtt{F}} \subseteq T^+$.

NO-THIN-AIR. Trivially follows from the facts that $G.\mathtt{sb} \cup G.\mathtt{rf} \subseteq G_t.\mathtt{hb}$, and $G_t.\mathtt{hb}$ is irreflexive. □

Lemma K.2. *Let G be an* RMW*-execution, such that* $\mathtt{mod}(a) \sqsupset \mathtt{rlx}$ *for every $a \in G.\mathtt{E}$, and* $\mathtt{mod}(a) \sqsupseteq \mathtt{acqrel}$ *for every $a \in G.\mathtt{RMW}$. Let G' be the* RMW*-execution obtained from G by removing all the non-SC fences (that is: $G'.\mathtt{E} = G.\mathtt{E} \setminus G.\mathtt{F}^{\neq \mathtt{sc}}$, $G'.\mathtt{sb} = [G'.\mathtt{E}]; G.\mathtt{sb}; [G'.\mathtt{E}]$, $G'.\mathtt{rf} = G.\mathtt{rf}$, and $G'.\mathtt{mo} = G.\mathtt{mo}$). Then, G and G' have the same outcome, and if G' is* RC11*-consistent then so is G.*

Proof. The conditions on the modes of accesses imply that $[G'.\mathtt{E}]; G.\mathtt{hb}; [G'.\mathtt{E}] = G'.\mathtt{hb}$. Then, the RC11-consistency of G' trivially implies the RC11-consistency of G. □

Theorem 1. *For a program P, denote by $(\!|P|\!)$ the* TSO *program obtained by compiling P using the scheme in Fig. 8. Then, given a program P, every outcome of $(\!|P|\!)$ under* TSO *is an outcome of P under* RC11.

Proof. First, let P' be the program obtained from P by (i) strengthening all read/write accesses in P to be at least release/acquire ones, (ii) all RMWs to be acquire-release RMWs, and (iii) omitting all non-SC fences. Note that $(\!|P|\!) = (\!|P'|\!)$, and by Lemmas I.1 and K.2, every outcome of P' under RC11 is an outcome of P under RC11. Hence, it suffices to show that every outcome of $(\!|P'|\!)$ under TSO is an outcome of P' under RC11. Given a full TSO-consistent TSO execution G_t of $(\!|P'|\!)$, the compilation scheme ensures that there exists some full execution G of P' for which the properties of Lemma K.1 hold. The claim then follows by Lemma K.1. □

K.1 Alternative correctness of "fences before SC reads" correctness

Here, we follow a different simpler approach utilizing the recent result of Lahav and Vafeiadis [19]. That result provides an alternative characterization of the TSO memory model, in terms of program transformations (or "compiler optimizations"). They show that every weak behavior of TSO can be explained by a sequence of:

- load-after-store reorderings
  (*e.g.*, $\mathtt{MOV}\ [x]\ \mathtt{1}\mathtt{;}\ \ \mathtt{MOV}\ r\ [y] \rightsquigarrow \mathtt{MOV}\ r\ [y]\mathtt{;}\ \mathtt{MOV}\ [x]\ \mathtt{1}$); and
- load-after-store eliminations
  (*e.g.*, $\mathtt{MOV}\ [x]\ \mathtt{1}\mathtt{;}\ \ \mathtt{MOV}\ r\ [x] \rightsquigarrow \mathtt{MOV}\ [x]\ \mathtt{1}\mathtt{;}\ \ \mathtt{MOV}\ r\ \mathtt{1}$).

They further outline an application of this characterization to prove compilation correctness, which we follow here. Accordingly, we have to meet three conditions:

1. Every outcome of the compiled program under SC is an outcome of the source program under RC11. This trivially holds, since obviously RC11 is weaker than SC (even if arbitrary fences are added to the source).
2. Every store-load reordering that can be applied on the compiled program corresponds to a transformation on the source program that is sound under RC11. Indeed, the compilation scheme ensures that adjacent load after store in the compiled program $(\!|P|\!)$ correspond to adjacent read after non-SC write in the source P. These can be soundly reordered under RC11 (see §7), resulting in a program P' whose compilation $(\!|P'|\!)$ is identical the reordered $(\!|P|\!)$.
3. Every load-after-store elimination that can be applied on the compiled program corresponds to a transformation on the source program that is sound under RC11. Again, the compilation scheme ensures that a load adjacently after a store in the compiled program $(\!|P|\!)$ corresponds to an adjacent non-SC read after a write in the source P. The read can be soundly eliminated under RC11 (see §7).

Note that this simple argument cannot be applied for the compilation scheme that places fences after SC writes, since a load adjacently after a store in the compiled program $(\!|P|\!)$ corresponds in this case to an adjacent read after a non-SC write in the source P. However, if the read is SC but the write is not SC, it is *unsound* to eliminate the read under RC11 (see Remark 6).