


Compiling Volatile Correctly in Java

Shuyang Liu ✉ 

University of California, Los Angeles, United States

John Bender ✉

Sandia National Laboratories, United States

Jens Palsberg ✉

University of California, Los Angeles, United States

Abstract

The compilation scheme for `Volatile` accesses in the OpenJDK 9 HotSpot Java Virtual Machine has a major problem that persists despite a recent bug report and a long discussion. One of the suggested fixes is to let Java compile `Volatile` accesses in the same way as C/C++11. However, we show that this approach is invalid for Java. Indeed, we show a set of optimizations that is valid for C/C++11 but invalid for Java, while the compilation scheme is similar. We prove the correctness of the compilation scheme to Power and x86 and a suite of valid optimizations in Java. Our proofs are based on a language model that we validate by proving key properties such as the DRF-SC theorem and by running litmus tests via our implementation of Java in Herd7.

2012 ACM Subject Classification Software and its engineering → Semantics

Keywords and phrases formal semantics, concurrency, compilation

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2022.23

Acknowledgements We thank Doug Lea for the helpful insights on the Java language semantics and compilers; we thank Ori Lahav, Anton Podkopaev and Viktor Vafeiadis for initially pointing out the issue of the Java Access Modes model; we thank all the reviewers of ECOOP'22 for their useful feedback. This material is based upon work supported by the National Science Foundation under Grant No. 1815496.

1 Introduction

In OpenJDK 9, the Java programming language introduced the `VarHandle` API with Access Modes to provide a standard set of operations that gives clear semantics to programs with shared object fields. Among the four available Access Modes (which we will explain in Section 3 in detail), programmers are allowed to use `Volatile` mode to ensure the consistency of updates on shared variables. Conceptually, the set of `Volatile` mode accesses in a program is totally ordered [9]. If all of the accesses in a program are in `Volatile` mode, then the program should only have sequentially consistent executions since all accesses in that program are totally ordered.

Sadly, this basic property of `Volatile` mode does not hold under the current implementation of the Java compiler in OpenJDK 9 HotSpot JVM. That is, marking all accesses as `Volatile` in a Java program can still result in behaviors that are not sequentially consistent when compiling to Power [14]. In particular, the C1 and the C2 compilers in HotSpot do not provide enough synchronization between a `Volatile` read and a `Volatile` write when compiling to the Power architecture. While we leave the details of their respective compilation schemes to Section 2, when a program includes a sequence of a `Volatile` read followed by a `Volatile` write, there is no `hwsync` instruction inserted in-between. Without the `hwsync`, it is possible for threads to disagree on the orders in which instructions are executed. As a consequence, the compilation schemes can still cause violations of sequential consistency in programs with



© Shuyang Liu, John Bender, Jens Palsberg;

licensed under Creative Commons License CC-BY 4.0

36th European Conference on Object-Oriented Programming (ECOOP 2022).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:71



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

all accesses marked `Volatile`. We have contacted the maintainers of the OpenJDK about this issue and a bug report has been filed [18].

One solution is to add the missing `hwsync` instruction to restore sequential consistency for `Volatile`. The resulting compilation scheme is similar to C/C++11 [4], which leads one to wonder whether Java compilers can simply handle Access Modes the same way as C/C++11 compilers handle atomic memory orders. However, there are significant differences in the semantics of `Volatile` access mode and the `seq_cst` memory order, which leads to differences in the valid compiler transformations applied to them respectively. In contrast to C/C++11 [6], Java does not allow certain compiler transformations to be applied to `Volatile` accesses. For example, register promotion cannot be applied to memory locations with `Volatile` accesses in Java while it can be applied in C/C++11. The differences provide Java programmers stronger synchronization guarantees and a more intuitive reasoning process: `Volatile` accesses (1) are equivalent to inserting `fullFence()`s, and (2) will not be optimized by the compiler in unexpected ways. We provide a detailed comparison along with soundness proofs and examples in Section 5.

While the change to the compilation scheme appears to be simple, the work of verifying its soundness is challenging. First, the formal language model *JAM* (hereafter *JAM*₁₉) [3] exhibits the same issue as the HotSpot compilers. That is, it cannot guarantee sequential consistency for programs with all accesses marked `Volatile`. Therefore, we revise the language model to fix this issue. To ensure the change to the model is valid we formally verify its key properties, such as the standard DRF-SC theorem, and leverage a set of empirical litmus tests via our implementation of Java in Herd7 [1] that keeps the model valid. We call the revised model *JAM*₂₁ to distinguish from the original version. Second, the language model defines the semantics of `fullFence()` with a total order. However, many target-level architectures such as the Power memory model [14] only specify a partial observable order among their synchronization mechanisms (fence cumulativity). Therefore, we develop an intermediate language model, *JAM'*₂₁, to bridge *JAM*₂₁ with the target level models. We show that *JAM'*₂₁ yields the same observable program executions as *JAM*₂₁ but does not specify a total order among `fullFence()`s, which simplifies the proof for compilation correctness.

1.1 Outline

The rest of the paper is structured as follows. Section 2 explains the bug in the current Java compiler to Power with an example. In Section 3, we explain the formal model that we use in this paper. Section 4 provides a correctness proof for our proposed compilation scheme to Power. Section 5 presents a set of program transformations that are valid/invalid for Java and a comparison with C/C++11. We include a discussion on expected performance impact in Section 6. Section 7 details some recent related work and finally, Section 8 concludes the paper.

1.2 Supplementary Material

The proofs of the theorems appear in this paper are available in the appendices (which are available in the full version of the paper). The following are also available as artifact of this paper at <https://github.com/ShuyangLiu/EC00P22-Supplementary-Material>.

- The extended Herd7 tool suite with the Java architecture.
- The litmus tests that appear in this paper.
- The Coq proofs for some of the theorems in this paper.

2 The Problem of Compiling Volatile and How to Fix it

In this section we use an example to demonstrate that the approach implemented by the HotSpot JVM compilers does not provide sequentially consistent semantics even when all accesses use `Volatile` mode.

Consider the `volatile-non-sc.4` example shown as an execution in Fig.1. In this example, there are four concurrent threads (P1, P2, P3, and P4) accessing two shared integer variables (`x` and `y`). The notation $Wx = 1$ means “writing to variable `x` with value 1”. The notation $Rx = 0$ means “reading from variable `x` and the value returned is 0”. In addition, each variable is initialized to 0 at the beginning before the threads start execution. The small superscript on each memory access denotes the access mode that the access uses. For example, Rx^v means “reading with `Volatile` mode”.

If all of the read and write accesses in this program use `Volatile` mode, would the reads ever return the values that are specified in the figure?

According to the specification [9], the program must exhibit sequentially consistent behavior because all accesses are marked `Volatile`:

“When all accesses use `Volatile` mode, program execution is sequentially consistent, in which case, for two `Volatile` mode accesses `A` and `B`, it must be that `A` precedes execution of `B`, or vice versa.”

Therefore, we are interested in whether the example in Fig. 1 is sequentially consistent. Sequential consistency, as first defined by [7], requires a total sequential order that preserves program order and the values returned by the reads are compatible with this total order. Following the definition, the execution in Fig. 1 does not satisfy sequential consistency. To see this, we demonstrate a contradiction under the guarantees of sequential consistency. Consider the following order constraints:

1. By program order, we know that (a) occurs before (b).
2. Since the value (b) gets is the initial value 0, it must occur before (c) writes to the location `y`.
3. Then, (d) reads the value written by (c), so (c) occurs before (d).
4. By program order, (d) occurs before (e).
5. Now, looking at P4, we know that the value of `x` changed from 1 to 2. Therefore, we can infer that (e) occurs before (a) since (e) is the only write to `x` with a value of 1 and (a) is the only write to `x` with a value of 2.

In this execution, we find a cycle: (a) \rightarrow (b) \rightarrow (c) \rightarrow (d) \rightarrow (e) \rightarrow (a) which appears in Fig. 1 with the “occurs before” relation represented as edges in the execution graph. Sequential consistency requires an irreflexive total order among all instructions. Therefore, the chain formed by the total order should be acyclic, i.e., a valid execution should not exhibit any cycle in its graph. Thus, this execution is inconsistent under sequential consistency and should be forbidden.

However, despite the promise of sequential consistency given by the source-level `Volatile` semantics, the compilation scheme found in the Java compilers for Power allows the example execution in Fig. 1. To see this, we present the compilation scheme from the C1 compiler which is the more conservative of HotSpot’s two compilers. We then give a Power-consistent execution graph corresponding to the example in Fig. 1.

The Power architecture adopts a relaxed memory model and provides fence instructions to recover sequential consistency. Two main types of fence instructions, the stronger fence `hwsync`




Figure 1 volatile-non-sc.4 under the sequential consistency model, Forbidden




Figure 2 volatile-non-sc.4.ppc translated to Power by HotSpot C1, Allowed

and the weaker fence `lwsync`, are usually used by the compilers to enforce synchronization guarantees. Using `lwsync` usually gives better performance but the synchronization guarantee of `lwsync` is weaker than `hwsync`. In particular, while both fence instructions carries a set of writes (Group A writes) when propagating to another thread, `lwsync` does not require an acknowledgement to continue executing the instructions after it. On the other hand, a `hwsync` requires an acknowledgment marking that it (along with its Group A writes) has propagated to all threads before proceeding to the next instruction.

The compilation to Power for Volatile accesses on C1 is the following ¹:

```

RV ~> hwsync ; lwz ; lwsync
WV ~> lwsync ; stw ; hwsync
    
```

A Volatile read is compiled to a `hwsync` instruction followed by a load instruction and a `lwsync` instruction; a Volatile write is compiled to a `lwsync` instruction followed by a store instruction and a `hwsync` instruction.

Fig. 2 shows the example from Fig. 1 according the compilation scheme in the C1 compiler².

The Power memory model [14] allows the behavior annotated in Fig. 2. The full trace of the execution can be found in Appendix M. Here we give a brief explanation. First note that a write operation is split into multiple steps and can be propagated to foreign threads in different orders if not properly synchronized. Furthermore, the `lwsync` in P3 is not sufficient in this case. In particular, the `lwsync` does not require an acknowledgement

¹ This compilation scheme was found in the OpenJDK 13 HotSpot compiler and it follows from a previously inaccurate description in the documentation [9] regarding the semantics of Volatile accesses. We have contacted the author and the documentation has been corrected in the latest version while the compiler bug (although reported) is still not fixed at the time of writing.

² The C2 compiler yields a slightly different compilation scheme for Volatile reads: Instead of inserting a `lwsync` fence after the load instruction, it emits a control dependency followed by an `isync` instruction, which we denote as `ctrlisync`. But in this example, the resulting execution graph is effectively the same as C1's because the effect of `ctrlisync` is subsumed into the `lwsync` or the `hwsync` instruction that it follows. In addition, we have simplified the compiled code (such as eliminating the fence instructions at the beginning or end of the threads and merging consecutive fence instructions) without changing its semantics for clarity here.

		Wx = 0	Wy = 0		
P1	P2			P3	P4
(a) Wx = 2	(c) Wy = 1			(d) Ry = 1	(f) Rx = 1
(B1) <code>hwsync</code>				(B2) <code>hwsync</code>	(B3) <code>hwsync</code>
(b) Ry = 0				(e) Wx = 1	(g) Rx = 2

■ **Figure 3** `volatile-non-sc.4.ppc` translated to Power using the revised compilation scheme, Forbidden

before proceeding to the next instructions and it only requires (c) `Wy = 1` to be propagated when itself needs to be propagated to the thread (the cumulativeness of `lwsync`). Since P4 needs to read from (e) `Wx = 1`, which is subsequent to (B2), (B2) needs to be propagated to P4 before (e) `Wx = 1` is propagated to P4. The propagation of (B2) `lwsync` makes sure that (c) `Wy = 1` is propagated to P4 before it can read x (even though it doesn't really need to read the value of y). On the other hand, P1 does not have any instructions reads from an instruction of P3 that comes after (in program order) (B2). Therefore, it does not require (c) and (B2) to be propagated to it when it executes (b). As a result, (c) can be propagated to P1 long after reaching P3 and hence letting P3 and P1 have different views of the memory during the execution. When P1 tries to read the value of y, it can only get an initial value of 0 since the newer value has not been propagated to P1 yet. Consequently, this non-SC execution is allowed (consistent) under the Power memory model, despite that the semantics of the "all-Volatile" source program requires it to be forbidden.

The solution to fix this issue is quite straightforward. Instead of letting `Volatile` read be translated using "leading fence" while `Volatile` write be translated using "trailing fence", they should both use the same fence inserting strategy (both leading fence or both trailing fence).³ Therefore, the correct compiler scheme for `Volatile` should be:

```

RV ↪ hwsync ; lwz ; lwsync
WV ↪ hwsync ; stw
```

With the revised compilation scheme we can demonstrate that the example of Fig. 1 is forbidden in accordance with the required SC semantics. The resulting execution graph is shown in Fig. 3. While most of this example matches Fig. 2, (B2) now is a `hwsync` instruction. As an effect of this change, (B2) is now required to be propagated to every thread and get acknowledged before start executing (e). As a result, at the time when (c) is propagated to P4 (as a result of the cumulative effect of (B2) just like in Section. 2), it must also have propagated to P1 due to the acknowledgement required by the `hwsync` at (B2). Therefore, it becomes impossible for (b) to read the value 0 because Power requires reads to always read from the latest value that has been propagated to the thread. That is, this execution is now forbidden by Power, aligning with the sequentially consistent semantics promised by the Java `Volatile` mode. Note that the reasoning is the same if we use a "trailing fence" scheme. The key is to deploy a fence insertion strategy such that there is a `hwsync` fence inserted between every pair of `Volatile` accesses.

³ Here we choose to show the leading fence strategy for simplicity. However, the trailing fence strategy is symmetric to leading fence and the same correctness proof works for both conventions given it's used consistently (more details can be found in Section 4.1). In practice, it is usually preferable to use trailing fence strategy for better performance.

Interestingly, we found similar compilation schemes applied to other architectures in HotSpot as well. This is not an accident. The source of this compiling behavior stems from the IR phase of the compiler. At the IR (called the Ideal Graph IR in HotSpot) level, a `Volatile` read is translated to a `fullFence()` followed by an `Acquire` read; a `Volatile` write is translated to a `Release` write followed by a `fullFence()`. Then each compiler back end translates the code further using the corresponding template file that maps the IR to specific architecture instructions. In the case of Power, a `fullFence()` is mapped to the `hwsync` instruction and `Release-Acquire` accesses are implemented using the `lwsync` instruction. While the example we provide here focuses on the compilation to Power, the more fundamental issue here is a lack of `fullFence()` between a `Volatile` read and a `Volatile` write at the IR encoding level. JAM_{19} aligns with this encoding when specifying the semantics of `Volatile` memory operations. As a result, JAM_{19} also exhibits the same problem. That is, when all memory accesses are `Volatile`, JAM_{19} does not guarantee sequential consistency.

3 Formal Model

In this section we present the revised model JAM_{21} , which we use as our theoretical foundation for proving compiler correctness in the rest of the paper. We begin by introducing the basic syntax (Section 3.1) used in the rest of the paper. Then we give the formal definition of JAM_{21} in Section 3.2.

3.1 Basic Syntax

We adopt the syntax of [3] and the `cat` language [1] in addition to some utility functions.

Given a program $P \in \mathbb{P}$, there is a set of executions (run-time traces) associated with P . We call the executions *histories* of P and use H to denote a single history. Each execution history consists of sets of memory access events specified by P . In particular:

- $H.E$ denotes the whole set of memory events of H .
- $H.F$ denotes the whole set of fence events of H .
- $H.IW$ denotes the set of initialization writes of H .
- $H.FW$ denotes the set of final writes of H .
- $H.W$ denotes the set of write events in H .
- $H.R$ denotes the set of read events in H .
- $H.RMW$ denotes the set of read-modify-write events in H .

Note that we treat each `RMW` events as a single event and $H.RMW \subseteq H.W$ and $H.RMW \subseteq H.R$. In addition, for `RMW` operations such as *compare-and-swap* (`CAS`), we assume the operation is on its success comparison path. They are sometimes implemented using `LL/SC` instructions on hardware, which cannot guarantee atomicity if the comparison fails. We assume each write event to the same memory location has an unique value for simplicity.

For each memory event i , we define the following utility functions to extract memory event attributes:

- $H.AccessMode(i)$ returns the Access Mode of event i in H .
- $H.val(i)$ returns the value of event i in H .
- $H.loc(i)$ returns the shared memory location of event i in H .
- $H.Tid(i)$ returns the thread identifier of which i is executed from

Finally, we use the symbol \mathbb{H} to denote the set of all execution histories.

The memory events of each H are related by order relations.

- The program order (**po**) is a partial order relation ($\text{po} \subseteq H.E \times H.E$) specified by P . We use the notation $i_1 \xrightarrow{\text{po}} i_2$ to denote the pair of events $\langle i_1, i_2 \rangle$ related by **po** and $H.\text{po}$ to denote the set of all pairs relates by **po** in H .
- The reads-from (**rf**) order is a partial order relation ($\text{rf} \subseteq H.W \times H.R$). For each read event i_2 , there exists a unique write event i_1 such that $H.\text{val}(i_1) = H.\text{val}(i_2)$ and $H.\text{loc}(i_1) = H.\text{loc}(i_2)$. We use the notation $i_1 \xrightarrow{\text{rf}} i_2$ to denote the pair of events $\langle i_1, i_2 \rangle$ related by **rf** and $H.\text{rf}$ to denote the set of all pairs relates by **rf** in H .
- Model-Specific relations. There are sets of relations that are specifically defined by the memory model. They are derived from the event attributes, **po**, and **rf** using the semantic rules of the memory model. We will detail them in the next few sections. We use the notation $i_1 \xrightarrow{R} i_2$ to denote the pair of events $\langle i_1, i_2 \rangle \in H.R$.

We also use operations on relations: given relations R_1 and R_2 , we use composition $R_1 ; R_2$, union $R_1 | R_2$, intersection $R_1 \& R_2$, complement $\sim R_1$, transitive closure R_1^+ , and inversion R_1^{-1} .

We may present an execution history H as a graph. An execution graph consists of a set of nodes labeled with unique identifiers, and a set of labeled edges. Each labeled node refers to an executed memory access.

Lastly, we use the notation $\text{acyclic}(\xrightarrow{R})$ to denote that R is acyclic in the execution history.

3.2 The JAM_{21} Model

In this section, we present the JAM_{21} model. The full definition of the relations in JAM_{21} can be found in Appendix A. We explain several excerpts of the formal model.

There are five available access modes in JAM_{21} : Plain mode, Opaque mode, Release mode, Acquire mode, and Volatile mode. The synchronization effect of the access modes are partially ordered using \sqsubseteq :

$$\text{Plain} \sqsubseteq \text{Opaque} \sqsubseteq \{\text{Release, Acquire}\} \sqsubseteq \text{Volatile}.$$

3.2.1 Visibility

At the center of JAM_{21} is the notion of *visibility orders* (**vo**). The most basic form of visibility, **vo** includes the reads-from (**rf**) relation. Intuitively, a read has certainly seen the effects of the write it takes its value from. Otherwise, visibility comes from synchronization⁴. Both Volatile (V) and Release(REL)-Acquire(ACQ), (RA as the union) accesses provide synchronization and thus visibility. Note that Volatile accesses are also included in the set of accesses that are considered Release-Acquire by the model. Further, **vo** can be derived from **ra** or **svo** orders, which captures the synchronization effects of Release-Acquire memory events or fences, **spush** or **volint** orders, which capture the synchronization effects of Volatile memory events or **fullFence()**s. In addition, the **pushto** order is trace order (**to**) restricted to the domain of **spush** and **volint**. Composing **pushto** with **spush** or **volint** emulates the cross-thread total order among **fullFence()**s, which is also part of the **vo** order. Finally, **po** to the same location is also included as part of the **vo** definition.

$$\text{ra} \triangleq \text{po} ; [\text{REL} | \text{V}] | [\text{ACQ} | \text{V}] ; \text{po}$$

⁴ Here, we use the high-level term "synchronization" for any memory consistency guarantee among instructions. We noticed that the usage of this term might differ outside of this paper. Therefore, we try to avoid using this term ambiguously to avoid confusion.

$$\begin{aligned}
\text{svo} &\triangleq \text{po} ; [\text{F} \ \& \ \text{REL}] ; \text{po} ; [\text{W}] \mid [\text{R}] ; \text{po} ; [\text{F} \ \& \ \text{ACQ}] ; \text{po} \\
\text{spush} &\triangleq \text{po} ; [\text{F} \ \& \ \text{V}] ; \text{po} \\
\text{volint} &\triangleq [\text{V}] ; \text{po} ; [\text{V}] \\
\text{vvo} &\triangleq \text{rf} \mid \text{svo} \mid \text{ra} \mid \text{spush} \mid \text{volint} \mid \text{pushto} ; (\text{spush} \mid \text{volint}) \\
\text{vo} &\triangleq \text{vvo}^+ \mid \text{po-loc}
\end{aligned}$$

Note that the definition of `volint` has been corrected from JAM_{19} to ensure sequential consistency for Volatile.

3.2.2 Coherence

The coherence order, `co-jom`, is an order among writes to the same location. Coherence order edges can be derived using the `vo` order and the `po` order among memory accesses.

$$\begin{aligned}
\text{WWco}(\text{rel}) &\triangleq \{ \langle i_1, i_2 \rangle \mid \langle i_1, i_2 \rangle \in H.\text{rel} \wedge i_1, i_2 \in H.W \wedge H.\text{loc}(i_1) = H.\text{loc}(i_2) \wedge i_1 \neq i_2 \} \\
\text{coww} &\triangleq \text{WWco}(\text{vo}) \\
\text{cowr} &\triangleq \text{WWco}(\text{vo} ; \text{rf}^{-1}) \\
\text{corw} &\triangleq \text{WWco}(\text{vo} ; \text{po}) \\
\text{corr} &\triangleq [0 \mid \text{RA} \mid \text{V}] ; \text{WWco}(\text{rf} ; \text{po} ; \text{rf}^{-1}) ; [0 \mid \text{RA} \mid \text{V}] \\
\text{co-jom} &\triangleq \text{coww} \mid \text{cowr} \mid \text{corw} \mid \text{corr}
\end{aligned}$$

Note that `co-jom` is different from the definition of `co` in other memory models such as Power and x86-TSO. Instead of enumerating all possible total coherence order to check the consistency of a given execution history, JAM_{21} derives coherence order `co-jom` among memory events from their known relations. Therefore, `co-jom` is a partial order among writes to the same location in JAM_{21} . We use the notation $i_1 \xrightarrow{\text{co-jom}} i_2$ to denote the pair of events $\langle i_1, i_2 \rangle$ related by `co-jom` and $H.\text{co-jom}$ to denote the set of all pairs relates by `co-jom` in H . We use the simpler name `co` to denote `co-jom` when the context is clear.

In addition, different from JAM_{19} , Plain mode reads to the same location ordered by `po` can be reordered by compiler and therefore cannot be used to derive `co-jom` order.

3.2.3 Execution Consistency

Axiomatic models define program semantics as the set of allowed executions. We adopt the same definition of *candidate execution* from [1].

► **DEFINITION 1 (CONSISTENT CANDIDATE EXECUTION).** Given a program P and a memory model M , an execution history H is a **M-consistent candidate execution of P** if and only if:

- H is a candidate execution of P (specified by the architecture of the programming language of which P is written in).
- H is M -consistent.

We denote the set of all M -consistent candidate executions of P by $\text{Histories}_M(P)$.

We now have all the definitions needed to define execution consistency under JAM_{21} .

► **DEFINITION 2 (JAM_{21} -CONSISTENCY).** An execution history H is **JAM_{21} -consistent** if it is trace coherent and satisfies the following two requirements:

1. **NO-THIN-AIR:** `po` \mid `rf` is acyclic. $\text{acyclic}(\text{po} \mid \text{rf})$

2. COHERENCE.: `co-jom` is acyclic, `acyclic(-co-jom→)`

We say such an execution history H is **allowed** by JAM_{21} . Otherwise, it is **forbidden**.

For the JAM_{21} model, we use $Histories_{JAM_{21}}(P)$ to denote the set of all JAM_{21} -consistent execution histories of P .

JAM_{21} satisfies a set of properties such as the DRF-SC Theorem. We show the theorems and the proofs in Appendix H and Appendix I.

3.2.4 Validation with Litmus Tests

The experimental validation of the JAM_{21} model includes two parts.

First, we implement the Java *architecture* in Herd7. Herd7 [1] was developed to simulate program executions with user-defined memory models. An *architecture* in Herd7 provides the parser for litmus tests written in the language corresponding to the architecture and an operational semantics of the instructions that appear in litmus tests. Herd7 uses the parser and the instruction semantics from the architecture to form an internal representation of the input litmus test and generate the set of all possible executions. Then, Herd7 checks the consistency of the executions using memory models written in the `cat` language. As of today, several mainstream architectures, such as C/C++11 [6], x86 [15], ARM [2], and Power [14], have been implemented and included in Herd7’s official repository. Unfortunately, Java is not. JAM_{19} [3] validated its formalization by mapping memory events to other architectures’ events that exists in the Herd7 repository and run the litmus tests in the architecture’s language. The mapping roughly captures part of the compilation scheme but it is neither complete nor proven sound. For example, in its mapping to ARMv8, `Volatile` accesses are ignored and not mapped to any memory event. Hence this approach is invalid and the results cannot be trusted though they show intentions on how JAM_{19} was expected to behave. Therefore, we extend the Herd7 tool suite with the Java architecture and translate the set of litmus tests used for testing JAM_{19} to Java⁵. A detailed description of each supported instruction is shown in Appendix K.1.

Second, we validate the JAM_{21} model using the Java translation of the set of litmus tests that was originally used to validate JAM_{19} and compare their outcomes. The results are mostly the same as the results from JAM_{19} except for three cases that are relevant to the inconsistency issue discussed earlier in this paper because we wish to fix the issue while keeping other parts of the model unchanged. The three exceptions reveal another aspect of the change, accommodating both the leading fence convention and the trailing fence convention, whereas JAM_{19} forced the compiler to choose a particular (problematic) convention. Since the compiler is free to choose either convention, a full synchronisation is only guaranteed to appear between a pair of `Volatile` accesses. In effect, certain executions that was forbidden by JAM_{19} are allowed by JAM_{21} since it is no longer guaranteed that `Volatile` writes are *followed* by a full synchronisation and `Volatile` reads are *prepending* with a full synchronisation. In addition, we have added new litmus tests for showing the change in the semantics of `Volatile`, `volatile-non-sc.4` and `volatile-non-sc.5`. While JAM_{19} allows the non-sequentially consistent behavior, JAM_{21} correctly forbids them. We further translated the examples to Power using the problematic compilation scheme, `volatile-non-sc.4.ppc` and

⁵ Note that not all tests are translatable. For example, for the cases that test address dependencies, there is no corresponding Java version since the notion of address dependency does not exist in Java. We drop a small set of litmus tests due to this reason.

volatile-non-sc.5.ppc, and the tests are indeed allowed by the Power memory model. Please see Appendix K.2 for a detailed report.

4 Compilation Correctness to Power

In this section, we show that the revised compilation scheme for Power is correct with respect to the Power memory model [14]. We use an intermediate model for the Java Access Modes that is observationally equivalent to JAM_{21} , which we call JAM'_{21} . We include the detailed definition of JAM'_{21} and the proofs for their observational equivalence to Appendix D.1 and its full definition in Appendix B. We use JAM'_{21} to prove that the revised compilation scheme to Power is correct.

4.1 The Power Memory Model

We use the Power memory model defined in Herd7 [1], which consists of the following basic order definitions (Please see Appendix C for the full semantics):

- `po` and `rf` follows the same definitions as in JAM_{21} (as described in Section. 3).
- `co` is the union of total orders among writes to the same location. Additionally, if i_1 and i_2 are events on different threads and $i_1 \xrightarrow{\text{co}} i_2$, then $i_1 \xrightarrow{\text{coe}} i_2$.
- `ctrl` is the control dependency between memory accesses.
- `ppo` is the set of preserved program orders. The detailed definition can be found in Appendix C.
- $\text{chapo} \triangleq \text{rfe} | \text{fre} | \text{coe} | (\text{fre} ; \text{rfe}) | (\text{coe} ; \text{rfe})$
- $\text{com} \triangleq \text{rf} | \text{fr} | \text{co}$
- `po-loc` is a subset of `po` that relates accesses to the same locations.
- `rmw` relates the read and the write access from the same RMW memory event.
- $\text{hb} \triangleq \text{ppo} | (\text{sync} | \text{lwsync}) | \text{rfe}$
- $\text{propbase} \triangleq ((\text{sync} | \text{lwsync}) | (\text{rfe} ; (\text{sync} | \text{lwsync}))) ; \text{hb}^*$
- $\text{prop} \triangleq \text{propbase} \& (\text{W} * \text{W}) | (\text{chapo?} ; \text{propbase}^* ; \text{sync} ; \text{hb}^*)$
- Additional order definitions can be found in Appendix C.

► **DEFINITION 3 (POWER CONSISTENCY).** An execution history H is **Power-consistent** if it is trace coherent and satisfies the following six requirements:

1. **SC-PER-LOCATION:** `po-loc` | `com` is acyclic.
2. **ATOMICITY:** `rmw` & `(fre ; coe)` is empty.
3. **NO-THIN-AIR:** `hb` is acyclic.
4. **PROPAGATION:** `(co` | `prop)` is acyclic.
5. **OBSERVATION:** `fre; prop; hb*` is irreflexive.
6. **SCXX:** `co` | `(po & (X * X))` is acyclic (where X denotes atomic accesses)

We say such an execution history H is **allowed** by Power. Otherwise, it is **forbidden**.

4.2 Compilation Scheme

We use the compilation scheme in Fig. 4. Note that this is slightly different from the compilation scheme found in OpenJDK HotSpot compiler in that each `Opaque` mode read is translated to a load instruction followed by a conditional branch. This enables us to ensure the **NO-THIN-AIR** property as it is not guaranteed in the Power memory model. The problem of Out-of-Thin-Air in axiomatic models has been an active research area for a long time and

```

    getOpaque() ~ lwz ; cmp ; bc
    setOpaque() ~ stw
    getAcquire() ~ lwz ; lwsync
    setRelease() ~ lwsync ; stw
    getVolatile() ~ hwsync ; lwz ; lwsync
    (Or getVolatile() ~ lwz ; hwsync for trailing fence convention)
    setVolatile() ~ hwsync ; stw
    (Or setVolatile() ~ lwsync ; stw ; hwsync for trailing fence convention)
    AcquireFence() ~ lwsync
    ReleaseFence() ~ lwsync
    fullFence() ~ hwsync
    getAndAdd() ~ hwsync ; _1: ldarx ; add ; stdcx. ; bne _1 ; lwsync
    (Or getAndAdd() ~ lwsync ; _1: ldarx ; add ; stdcx. ; bne _1 ; hwsync for trailing fence convention)
    getAndAddAcquire() ~ _1: ldarx ; add ; stdcx. ; bne _1 ; lwsync
    getAndAddRelease() ~ lwsync ; _1: ldarx ; add ; stdcx. ; bne _1

```

■ **Figure 4** Compilation to Power

there exists various ways to use weaker compilation schemes while still ruling out thin-air reads. However, it is out of the scope of this paper and here we adopt the stronger scheme for Opaque mode to simplify the proofs. Additionally, we fix the compilation scheme for Volatile as suggested in Section 2. Note that both leading fence and trailing fence conventions ensure a `hwsync` instruction is inserted between each pair of Volatile mode accesses as long as they are used consistently (use the same convention for Volatile writes and reads). Therefore, the proof for the trailing fence convention can be carried out in a very similar way as the proof for the leading fence convention.

We start our proof by defining a *CompilesTo* relation over execution histories that relates source level executions to target level executions. Intuitively, the process of compilation can be seen as a transformation function on executions from source level to target level. With the *CompilesTo* relation, we can characterize a subset of target level executions that are constructed particularly through the compilation (following a given compilation scheme) from the source level. Note that at this step we do not check whether the resulting execution is consistent under the target level memory model, since the consistency of an execution is checked after the execution is constructed in axiomatic memory models.

► **DEFINITION 4 (COMPILATION OF AN EXECUTION).** We define the "CompilesTo" relation $\rightsquigarrow \subseteq \mathbb{H} \times \mathbb{H}$ for the compilation from Java to Power as the following: Given a Java program P_{src} , let P_{tgt} be the target-level program compiled from P_{src} using the compilation scheme in Fig. 4 (using the leading fence convention). Let H_{src} be a candidate execution history of P_{src} and H_{tgt} be a candidate execution history of P_{tgt} . We say $H_{src} \rightsquigarrow H_{tgt}$ if:

- $H_{tgt}.IW = H_{src}.IW$
- $H_{tgt}.FW = H_{src}.FW$
- $H_{tgt}.E = H_{src}.E$
- $H_{tgt}.rf = H_{src}.rf$
- $H_{tgt}.po = H_{src}.po$
- $H_{tgt}.co \subseteq H_{src}.to$
- If $i_1 \in H_{src}.E$, $i_{rmw} \in H_{src}.RMW$ and $i_{rmw} \xrightarrow{po} i_1$, then $i_{rmw} \xrightarrow{ctrl} i_1$ in H_{tgt}
- If $i_R^{\neg O} \in H_{src}.R$, $i_1 \in H_{src}.E$ and $i_R^{\neg O} \xrightarrow{po} i_1$, then $i_R \xrightarrow{ctrl} i_1$ in H_{tgt}
- If $i_1, i_2 \in H_{src}.E$ and $i_1 \xrightarrow{push} i_2$, then $i_1 \xrightarrow{sync} i_2$ for $i_1, i_2 \in H_{tgt}.E$

- If $i_1, i_2 \in H_{src}.E$ and $i_1 \xrightarrow{ra} i_2$, then $i_1 \xrightarrow{lwsync} i_2$ for $i_1, i_2 \in H_{tgt}.E$

Once we have the source level and target level execution histories, we use the memory model to check for consistency. A correct compilation, intuitively, should not introduce any new program behavior. In this context, it means there should not be any execution H_{src} that is forbidden by the source level memory model being related (by the "CompilesTo" relation) with a H_{tgt} that is allowed by the target level memory model. That is, if H_{tgt} is consistent under the target level memory model, then H_{src} should also be consistent under source level memory model. Formally, we have the following definition (recall that we use $Histories_M(P)$ to denote the set of consistent execution histories if a program P under a memory model M).

► **DEFINITION 5 (COMPILATION CORRECTNESS)**. Let P_{src} be a source program and S be a memory model that supports the source language, P_{tgt} be the target program compiled from P_{src} using a compilation scheme and T be a memory model that supports the target language. We say a compiler that compiles P_{src} to P_{tgt} is **correct** if for all $H_{tgt} \in Histories_T(P_{tgt})$ there exists a $H_{src} \in Histories_S(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

4.3 Proof of Compilation Correctness

We leverage an intermediate memory model, JAM'_{21} , to prove the compilation correctness to Power. While the complete definition of JAM'_{21} can be found in Appendix D, it is important to note that JAM'_{21} is *observationally equivalent* to JAM_{21} , which means they allow the same visible program behaviors given the same program. Intuitively, each consistent execution under JAM_{21} has a corresponding consistent execution under JAM'_{21} with the same set of events and the same observable value on each event. Formally, we give the following definitions for observational equivalence.

► **DEFINITION 6 (OBSERVATIONAL EQUIVALENCE OF EXECUTION HISTORIES)**. Given a program P , let H and H' be two execution histories of P . We say H and H' are **observationally equivalent** if:

- $H.IW = H'.IW$
- $H.FW = H'.FW$
- $H.E = H'.E$
- $H.po = H'.po$
- $H.rf = H'.rf$
- $\forall i \in H.E, H.AccessMode(i) = H'.AccessMode(i)$

► **DEFINITION 7 (OBSERVATIONAL EQUIVALENCE OF MEMORY MODELS)**. Given a program P , let M_1 and M_2 be two memory models that support the architecture of the programming language that P is written in. Let $Histories_{M_1}(P)$ be the set of all M_1 -consistent candidate executions of P ; let $Histories_{M_2}(P)$ be the set of all M_2 -consistent candidate executions of P . We say M_1 and M_2 are **observationally equivalent** if:

- (\Rightarrow) For all $H_1 \in Histories_{M_1}(P)$, there exists $H_2 \in Histories_{M_2}(P)$ such that H_1 is observationally equivalent to H_2 .
- (\Leftarrow) For all $H_2 \in Histories_{M_2}(P)$, there exists $H_1 \in Histories_{M_1}(P)$ such that H_2 is observationally equivalent to H_1 .

Then we prove the compilation correctness from JAM'_{21} to Power.

► **LEMMA 1 (JAM'_{21} TO POWER)**. Let P_{src} be a Java program, P_{tgt} be the Power program compiled from P_{src} using the compilation scheme in Fig. 4 (with the leading fence convention).

For all $H_{tgt} \in \text{Histories}_{\text{Power}}(P_{tgt})$ there exists a $H_{src} \in \text{Histories}_{\text{JAM}'}(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

Please see Appendix D for the proof.

Finally, we associate JAM_{21} with JAM'_{21} through the notion of observational equivalence and prove the compilation correctness from JAM_{21} to Power.

► **THEOREM 1 (COMPILATION CORRECTNESS TO POWER (LEADING FENCE CONVENTION))**. The compilation from Java to Power following the compilation scheme in Fig. 4 (using the leading fence convention) is correct. That is, let P_{src} be a Java program, P_{tgt} be the Power program compiled from P_{src} using the compilation scheme in Fig. 4 (using the leading fence convention). For all $H_{tgt} \in \text{Histories}_{\text{Power}}(P_{tgt})$ there exists a $H_{src} \in \text{Histories}_{\text{JAM}}(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

Please see Appendix D for the proof.

► **COROLLARY 1 (COMPILATION CORRECTNESS TO POWER (TRAILING FENCE CONVENTION))**. The compilation from Java to Power following the compilation scheme in Fig. 4 (using the trailing fence convention) is correct. That is, let P_{src} be a Java program, P_{tgt} be the Power program compiled from P_{src} using the compilation scheme in Fig. 4 (using the trailing fence convention). For all $H_{tgt} \in \text{Histories}_{\text{Power}}(P_{tgt})$ there exists a $H_{src} \in \text{Histories}_{\text{JAM}}(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

Please see Appendix D for the proof.

5 Compiler Transformations

One important aspect of compilers is the program transformations that they apply to the program. A correct compiler transformation should not introduce any new program behavior. While this is relatively simple for sequential programs, it can yield subtle issues when applying the same transformations to concurrent programs. A memory model's task is then to accommodate a set of common program transformations while still provide intuitive synchronization guarantees to the programmers. In Section 4 we show that Java and C/C++11 can use the same compilation scheme to Power (and x86, see Appendix F). However, Java has a stronger semantics for `Volatile` comparing to `seq_cst` in C/C++11 and can adopt only a strict subset of the transformations that are valid for C/C++11.

In this section, we use the set of compiler transformations detailed by [6] and compare their soundness in Java with C/C++11. We provide formal proofs for the sound transformations and counter-examples for invalid transformations. We conclude this section by discussing the implications of our results.

To prove a transformation is valid, intuitively, we show that there does not exist a H_{src} of P_{src} such that it is forbidden by JAM_{21} but the corresponding H_{tgt} of P_{tgt} is allowed.

► **DEFINITION 8 (VALID PROGRAM TRANSFORMATION)**. Let P_{src} be a Java program which has a set of candidate executions, $\text{Histories}(P_{src})$. Let $T : \mathbb{H} \rightarrow \mathbb{H}$ be a program transformation and $H_{tgt} = T(H_{src})$ for each candidate execution H_{src} of P_{src} . Then we say T is **valid** under JAM_{21} if and only if for each H_{tgt} , if H_{tgt} is JAM_{21} -consistent, then H_{src} is also JAM_{21} -consistent.

The results for Java comparing them C/C++11 [6] are summarized in Fig. 5.

Transformation		C/C++11	Java
Strengthening	[Sec. 5.1]	✓	✓
Sequentialisation	[Sec. 5.2]	✓	✓
Reordering	[Sec. 5.3]		See Fig. 6
Merging	[Sec. 5.4]		See Fig. 7
Register Promotion	[Sec. 5.5]	✓	For locations that does not have <code>Volatile</code> access

■ **Figure 5** Compiler Transformations in C/C++11 and Java

5.1 Strengthening

Strengthening transforms the access mode of accesses to stronger access modes. It is supported by JAM_{21} due to the monotonicity property of the memory model. The formal theorem is the following:

► **THEOREM 2 (STRENGTHENING)**. Let H_{tgt} an execution of P_{tgt} , which is obtained from applying Strengthening to a program P_{src} . There exists an execution H_{src} of P_{src} such that:

- $H_{src}.E = H_{tgt}.E$
- $H_{src}.po = H_{tgt}.po$
- $H_{src}.rf = H_{tgt}.rf$
- $\forall i \in H_{src}.E, H_{src}.AccessMode(i) \sqsubseteq H_{tgt}.AccessMode(i)$

If H_{tgt} is JAM_{21} -consistent, then H_{src} is JAM_{21} -consistent.

Proof. By Monotonicity of JAM_{21} , all the constraints in H_{src} are preserved in the strengthened execution H_{tgt} . Therefore, if H_{tgt} is JAM_{21} -consistent, so is H_{src} . ◀

5.2 Sequentialisation

Sequentialisation transforms two concurrent accesses into accesses in a single sequential process. It is naturally supported by JAM_{21} because sequentialisation does not remove any synchronization from the program.

► **THEOREM 3 (SEQUENTIALISATION)**. Let P_{src} be a Java program and P_{tgt} be a Java program obtained by performing a sequentialisation operation on a pair of accesses a and b . Let H_{tgt} be an execution of P_{tgt} . Then there exists an execution H_{src} of P_{src} such that

- $H_{src}.po \cup \{\langle a, b \rangle\} = H_{tgt}.po$ where $\langle a, b \rangle \notin H_{src}.po$ and $\langle b, a \rangle \notin H_{src}.po$
- $H_{src}.rf = H_{tgt}.rf$
- $H_{src}.E = H_{tgt}.E$
- $H_{src}.to = H_{tgt}.to$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{src}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$

and if H_{tgt} is JAM_{21} -consistent, then H_{src} is JAM_{21} -consistent.

Proof. Assume towards contradiction that H_{src} is not JAM_{21} -consistent. Then there are two cases: either there is a $(po | rf)^+$ cycle or a co cycle in H_{src} . Whether or not a and b are included in this cycle, adding a po edge between a and b cannot eliminate this cycle (although it might introduces new cycles). Therefore, H_{tgt} is also not JAM_{21} -consistent, contradicting to our assumption. ◀

	$\mathbf{R}_y^{m_2}$	$\mathbf{W}_y^{m_2}$	$\mathbf{RMW}_y^{m_2}$	\mathbf{F}^{m_2}
$\mathbf{R}_x^{m_1}$	$m_1 \sqsubseteq \text{Opaque}$	$m_1, m_2 \sqsubseteq \text{Opaque} \wedge (m_1 = \text{Plain} \vee m_2 = \text{Plain})$	$m_1 = \text{Plain} \wedge m_2 \sqsubseteq \text{Acquire}$	$(m_1 \sqsubseteq \text{Opaque} \wedge m_2 = \text{Release} \wedge \forall i, \mathbf{F}^{m_2} \xrightarrow{\text{po}} i \Rightarrow i \notin H.W) \vee (m_1 = \text{Acquire} \wedge m_2 = \text{Acquire}) \vee (m_1 = \text{Acquire} \wedge m_2 = \text{Release})$
$\mathbf{W}_x^{m_1}$	$m_1 \neq \text{Volatile} \vee m_2 \neq \text{Volatile}$	$m_2 \sqsubseteq \text{Opaque}$	$m_2 \sqsubseteq \text{Acquire}$	$(m_2 = \text{Acquire}) \vee (m_2 = \text{Release} \wedge \forall i, \mathbf{F}^{m_2} \xrightarrow{\text{po}} i \Rightarrow i \notin H.W) \vee (m_2 = \text{Release} \wedge \forall i, \mathbf{F}^{m_2} \xrightarrow{\text{po}} i \wedge i \in H.W \Rightarrow \text{AccessMode}(i) = \text{Release})$
$\mathbf{RMW}_x^{m_1}$	$m_1 \sqsubseteq \text{Release}$	$m_1 \sqsubseteq \text{Release} \wedge m_2 = \text{Plain}$	-	$(m_1 \sqsupseteq \text{Acquire} \wedge m_2 = \text{Acquire}) \vee (m_2 = \text{Release} \wedge \forall i, \mathbf{F}^{m_2} \xrightarrow{\text{po}} i \Rightarrow (i \in H.R \vee (i \in H.W \wedge \text{AccessMode}(i) = \text{Release})))$
\mathbf{F}^{m_1}	$(m_1 = \text{Release}) \vee (m_1 = \text{Acquire} \wedge \forall i, i \xrightarrow{\text{po}} \mathbf{F}^{m_1} \Rightarrow i \notin H.R)$	$m_1 = \text{Release} \wedge m_2 \sqsupseteq \text{Release} \vee (m_1 = \text{Acquire} \wedge \forall i, i \xrightarrow{\text{po}} \mathbf{F}^{m_1} \Rightarrow i \notin H.R)$	$m_1 = \text{Release} \wedge m_2 \sqsupseteq \text{Release} \vee (m_1 = \text{Acquire} \wedge \forall i, i \xrightarrow{\text{po}} \mathbf{F}^{m_1} \Rightarrow i \notin H.R)$	$(m_1 = \text{Release} \wedge m_2 = \text{Acquire}) \vee (m_1 = \text{Acquire} \wedge \forall i, i \xrightarrow{\text{po}} \mathbf{F}^{m_1} \Rightarrow i \notin H.R) \vee (m_2 = \text{Release} \wedge \forall i, \mathbf{F}^{m_2} \xrightarrow{\text{po}} i \Rightarrow i \notin H.W)$

■ **Figure 6** Allowed Deordering Pairs in JAM_{21}

5.3 Reordering

The operation of *reordering* can be seen as composing *deordering* with *sequentialisation*. Since we know that sequentialisation is sound in JAM_{21} , we only need to show that deordering is sound in order to show reordering is sound in JAM_{21} .

5.3.0.1 Deordering

Deordering is a transformation that turns a pair of accesses related by a **po** relation into a pair of concurrent accesses. In effect, it removes an **po** edge in the execution graph.

First, we adopt the same definition of adjacent events from [6]:

► **DEFINITION 9 (ADJACENT EVENTS)**. Two events a and b are **adjacent** in a partial order \mathbf{R} if for all c , we have:

- $c \xrightarrow{\mathbf{R}} a \Rightarrow c \xrightarrow{\mathbf{R}} b$
- $b \xrightarrow{\mathbf{R}} c \Rightarrow a \xrightarrow{\mathbf{R}} c$

For Java, the table of allowed reordering two adjacent events (with each row as the first event and column as the second event) is shown in Fig. 6 (some of the cases are different from C11 [6] and we have marked them in red). Intuitively, the sound deorderable pairs are ordered by the **po** edges that does not impose any synchronization in the program. Therefore, deordering (removing the **po** edge) does not introduce new program behavior.

To prove that JAM_{21} supports the reordering shown in this table, we need to prove each cell shown in the table is valid for JAM_{21} .

► **THEOREM 4 (DEORDERING)**. Let P_{src} be a Java program and P_{tgt} be a Java program obtained by performing a deordering operation on a pair of accesses a and b according to Fig. 6. Let H_{tgt} be an execution of P_{tgt} . Then there exists an execution H_{src} of P_{src} such that

Name	C/C++11	Java
Read-read Merging	$R^m; R^m \rightsquigarrow R^m$	$R^m \sqsubseteq \text{Acq}; R^m \sqsubseteq \text{Acq} \rightsquigarrow R^m$
Write-write Merging	$W^m; W^m \rightsquigarrow W^m$	$W^m \sqsubseteq \text{Rel}; W^m \sqsubseteq \text{Rel} \rightsquigarrow W^m$
Write/RMW-read Merging	$W^m; R^{\text{acq}} \rightsquigarrow W^m$ $W^{\text{sc}}; R^{\text{sc}} \rightsquigarrow W^{\text{sc}}$ $\text{RMW}^m; R^{m_r \sqsubseteq m} \rightsquigarrow \text{RMW}^m$	$W^m; R^m \sqsubseteq \text{Opq} \rightsquigarrow W^m$ X $\text{RMW}^m; R^m \sqsubseteq \text{Opq} \rightsquigarrow \text{RMW}^m$
Write-RMW Merging	$W^{m_w \sqsubseteq m}; \text{RMW}^m \rightsquigarrow W^{m_w}$	$W^{m_w \sqsubseteq \text{Rel}}; \text{RMW}^m \sqsubseteq \text{Vol} \rightsquigarrow W^{m_w}$
RMW-RMW Merging	$\text{RMW}^m; \text{RMW}^m \rightsquigarrow \text{RMW}^m$	$\text{RMW}^m \sqsubseteq \text{Vol}; \text{RMW}^m \sqsubseteq \text{Vol} \rightsquigarrow \text{RMW}^m$
Fence-fence Merging	$F^m; F^m \rightsquigarrow F^m$	$F^m; F^m \rightsquigarrow F^m$

■ **Figure 7** Mergable Pairs in C/C++11 [6] and Java

- $H_{src}.\text{po} = H_{tgt}.\text{po} \cup \{a, b\}$ where a and b are po-adjacent
- $H_{src}.\text{rf} = H_{tgt}.\text{rf}$
- $H_{src}.\text{E} = H_{tgt}.\text{E}$
- $H_{src}.\text{to} = H_{tgt}.\text{to}$
- $H_{src}.\text{IW} = H_{tgt}.\text{IW}$
- $\forall i \in H_{src}.\text{E}, H_{src}.\text{AccessMode}(i) = H_{tgt}.\text{AccessMode}(i)$

and if H_{tgt} is JAM_{21} -consistent, then H_{src} is JAM_{21} -consistent.

Please see Appendix G for the proof.

Reordering, as mentioned previously, can be decomposed into two steps: deordering and sequentialisation. Since we have already shown the soundness of the two transformations, the soundness of reordering follows naturally.

► **COROLLARY 2 (REORDERING)**. JAM_{21} supports the reordering transformation for pairs of adjacent accesses shown in Fig. 6.

5.4 Merging

Merging transforms two adjacent accesses into one single equivalent access to reduce the number of memory accesses in the program. We have grouped all types of merging transformations appeared in C/C++11 [6] here in one section. A summarized result of mergable pairs comparing with C/C++11 can be found in Fig. 7. The results are mostly similar except for *Volatile*. Many merging transformation are invalid for *Volatile* because they remove the cross-thread synchronization of *Volatile*.

5.4.1 Read-Read Merging

Read-read merging is sometimes done when the compiler is optimizing redundant loads in the same thread. When we are encountering two consecutive reads to the same location, the first read is unchanged but the second read becomes a local read without accessing the memory.

Let a' and b be two adjacent read accesses reading from the same write access a . $a \xrightarrow{\text{rf}} a'$ and $a \xrightarrow{\text{rf}} b$, and $a' \xrightarrow{\text{po}} b$. Assuming $\text{AccessMode}(a') = \text{AccessMode}(b)$, then

- $\forall i, a' \xrightarrow{\text{po}} i \Rightarrow b \xrightarrow{\text{po}} i$
- $\forall i, a' \xrightarrow{\text{ra}} i \Rightarrow b \xrightarrow{\text{ra}} i$
- $\forall i, a' \xrightarrow{\text{push}} i \Rightarrow b \xrightarrow{\text{push}} i$
- $\forall j, j \xrightarrow{\text{po}} b \Rightarrow j \xrightarrow{\text{po}} a'$

For executions, this corresponds to the following transformation in the execution graph: since the value of $r1$ and $r2$ are guaranteed to have the same value in P_{tgt} , we know that this corresponds to the execution of P_{src} where the two read accesses read from the same write access. Then we want to show that, if H_{tgt} is JAM_{21} -consistent, H_{src} is also JAM_{21} -consistent.

► **THEOREM 5 (READ-READ MERGING).** Let H_{tgt} be an JAM_{21} -consistent execution. Let $a \in H_{tgt}.R \setminus RMW$ and let $a' \in H_{tgt}.E$ such that $a \xrightarrow{\text{rf}} a'$. Let $b \notin H_{tgt}.E$. There exists a H_{src} such that:

- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, b \rangle \mid i \xrightarrow{\text{po}} a\} \cup \{\langle b, j \rangle \mid a \xrightarrow{\text{po}} j\}$
- $H_{src}.rf = H_{tgt}.rf \cup \{\langle a', b \rangle\}$
- $H_{src}.E = H_{tgt}.E \cup \{b\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{\text{to}} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{\text{to}} j\}$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $b \in H_{src}.R$
- $H_{src}.AccessMode(b) = H_{src}.AccessMode(a) \sqsubseteq \text{Acquire}$

and H_{src} is JAM_{21} -consistent.

Please see Appendix G for the proof.

Note that JAM_{21} does not allow read-read merging if the two read accesses are both Volatile mode reads. We provide an example of this in Appendix G.

5.4.2 Write-Write Merging

The write-write merge transformation refers to the program transformation that merges two consecutive write operations into one by removing the former one. JAM_{21} support write-write merge when the access modes of the two writes are the same and they are not Volatile mode accesses.

Let a and b be the two adjacent writes such that $a \xrightarrow{\text{po}} b$. We once again have the properties:

- $\forall i, i \xrightarrow{\text{po}} a \Rightarrow i \xrightarrow{\text{po}} b$
- $\forall j, b \xrightarrow{\text{po}} j \Rightarrow a \xrightarrow{\text{po}} j$
- $\forall i, i \xrightarrow{\text{ra}} a \Rightarrow i \xrightarrow{\text{ra}} b$

We have the following theorem.

► **THEOREM 6 (WRITE-WRITE MERGING).** Let H_{tgt} be an JAM_{21} -consistent execution. Let $b \in H_{tgt}.W \setminus RMW$ and let $a \notin H_{tgt}.E$ and $\text{loc}(a) = \text{loc}(b) \wedge \forall i \in H_{tgt}.W, \text{loc}(i) = \text{loc}(b) \Rightarrow \text{val}(a) \neq \text{val}(i)$. There exists a H_{src} such that:

- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{\text{po}} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{\text{po}} j\}$
- $H_{src}.rf = H_{tgt}.rf$

- $H_{src}.E = H_{tgt}.E \cup \{a\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{to} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{to} j\}$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $a \in H_{src}.W$
- $H_{src}.AccessMode(a) = H_{src}.AccessMode(b) \sqsubseteq \text{Release}$

and H_{src} is JAM_{21} -consistent.

Please see Appendix G for the proof.

Note that write-write merging is not valid for Volatile mode writes. We provide an example of this in Appendix G.

5.4.3 Write/RMW-read Merging

The Write/RMW-read merging refers to the program transformation that merges a write/RMW and a read into a single write/RMW and a local access.

Similarly, the transformation with an RMW operation and a read operation optimizes the latter read operation to read locally and in effect removes a memory load operation in the execution graph.

JAM_{21} only support this transformation when the read operation is (or is weaker than) Opaque mode which is different from RC11 [6]'s result for C/C++11. We provide a counterexample in Appendix G to show that write/RMW-read merging is invalid when the read is (or is stronger than) Acquire mode.

► **THEOREM 7 (WRITE/RMW-READ MERGING).** Let H_{tgt} be a JAM_{21} -consistent execution. Let $a \in H_{tgt}.W$ and $b \notin H_{tgt}.E$. There exists a H_{src} such that:

- $H_{src}.E = H_{tgt}.E \cup \{b\}$
- $b \in H_{src}.R$
- $H_{src}.loc(b) = H_{src}.loc(a)$
- $H_{src}.val(b) = H_{src}.val(a)$
- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{po} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{po} j\}$
- $H_{src}.rf = H_{tgt}.rf \cup \{\langle a, b \rangle\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{to} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{to} j\}$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $H_{src}.AccessMode(b) \sqsubseteq \text{Opaque}$

Please see Appendix G for the proof.

5.4.4 Write-RMW Merging

The write-RMW merging refers to the program transformation that merges a write and a consecutive RMW operation into a write with the value of the RMW. For example, if we have the following pattern in a program:

```
x = 1;
x.getAndSet(1, 2);
```

It can be transformed to:

```
x = 2;
```

Similar to write-write merging, JAM_{21} supports write-RMW merging when the access mode of the write is {Opaque, Release} and the access mode of the RMW is {Acquire, Release}.

► **THEOREM 8 (WRITE-RMW MERGING).** Let H_{tgt} be a JAM_{21} -consistent execution. Let $b \in H_{tgt}.W \setminus H_{tgt}.RMW$, $a \notin H_{tgt}.E$ and $v \in \mathbf{Val}$. There exists a H_{src} such that:

- $H_{src}.E = H_{tgt}.E \cup \{a\}$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $H_{src}.AccessMode(a) \in \{\text{Opaque}, \text{Release}\}$
- $H_{src}.AccessMode(b) \in \{\text{Acquire}, \text{Release}\}$
- $H_{src}.loc(b) = H_{src}.loc(a)$
- $b \in H_{src}.RMW$
- $H_{src}.val(b) = (H_{src}.val(a), v)$
- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{po} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{po} j\}$
- $H_{src}.rf = H_{tgt}.rf \cup \{\langle a, b \rangle\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{to} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{to} j\}$
- $H_{src}.IW = H_{tgt}.IW$

and H_{src} is JAM_{21} -consistent.

Please see Appendix G for the proof.

5.4.5 RMW-RMW Merging

The RMW-RMW merging transformation refers to the program transformation that merges two consecutive RMW operations into one such that it has the first RMW's (expected) read value and the second RMW's write value. For example, if we have the following pattern in a program:

```
x.getandSet(1,2);
x.getandSet(2,3);
```

then it might be transformed into:

```
x.getAndSet(1,3);
```

The RMW-RMW merging transformation is essentially the same as write-write merging and read-read merging described previously. Therefore, the set of constraints on valid access modes for merging is the intersection of the two. That is, two RMWs are mergeable if they are both Acquire mode or Release mode. For the counter-examples showing this transformation is invalid for Volatile accesses, please see the examples for write-write and read-read merging.

► **THEOREM 9 (RMW-RMW MERGING).** Let H_{tgt} be a JAM_{21} -consistent execution. Let x be a memory location and $a \in H_{tgt}.E$ with $H_{tgt}.val(a) = (v_r, v_w)$, $H_{tgt}.loc(a) = x$, and $H_{tgt}.AccessMode(a) \in \{\text{Release}, \text{Acquire}\}$. Let $b \notin H_{tgt}.E$, there exists a H_{src} such that:

- $H_{src}.E = H_{tgt}.E \cup \{b\}$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $H_{src}.val(a) = (v_r, v)$
- $H_{src}.val(b) = (v, v_w)$
- $H_{src}.loc(b) = x$
- $H_{src}.AccessMode(b) = H_{src}.AccessMode(a) \in \{\text{Release}, \text{Acquire}\}$
- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, b \rangle \mid i \xrightarrow{po} a\} \cup \{\langle b, j \rangle \mid a \xrightarrow{po} j\}$
- $H_{src}.rf = H_{tgt}.rf \cup \{\langle a, b \rangle\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, b \rangle \mid i \xrightarrow{to} a\} \cup \{\langle b, j \rangle \mid a \xrightarrow{to} j\}$

$$\cdot H_{src}.IW = H_{tgt}.IW$$

and H_{src} is JAM_{21} -consistent.

Please see Appendix G for the proof.

5.4.6 Fence-fence Merging

The Fence-fence merging refers to the program transformation that merges two consecutive fences of the same access mode into one. For example, if we have:

```
VarHandle.fullFence();
VarHandle.fullFence();
```

then it can be optimized to:

```
VarHandle.fullFence();
```

Since JAM_{21} is fence-based such that each fence is converted into an edge between memory accesses, this is trivially supported since the execution graph before and after the transformation is exactly the same.

5.5 Register Promotion for Non-shared Variable

Register Promotion promotes memory accesses of a non-shared memory location to local registers. It has the effect of removing memory accesses for thread-local variables. JAM_{21} only supports register promotion for variables without any `Volatile` accesses in the program. For non-`Volatile` accesses, since the variable is not shared across threads, it is safe to remove them without worrying about removing synchronization from the program. In contrast, `Volatile` accesses impose cross-thread synchronizations with `Volatile` accesses for other variables, so removing such accesses can potentially remove important synchronization in the program and introduce new behaviors that were previously forbidden by the memory model. We provide a counter-example in this section showing that we cannot promote `Volatile` accesses to local register accesses even if the location is only accessed by one thread.

Suppose all accesses to a memory location are in the same thread, the transformation can be seen as two steps:

1. Weakening the accesses to Plain mode accesses
2. Removing the Plain mode accesses

► **THEOREM 10 (WEAKENING FOR NON-SHARED VARIABLE).** Let H_{tgt} be a JAM_{21} -consistent execution such that, for all accesses i and j in $H_{tgt}.E$, $loc(i) = loc(j) = x \Rightarrow Tid(i) = Tid(j)$ for some memory location x . In addition, $\forall i \in H_{tgt}.E, loc(i) = x \Rightarrow AccessMode(i) = Plain$. There exists an execution H_{src} such that:

- $H_{src}.E = H_{tgt}.E$
- $H_{src}.po = H_{tgt}.po$
- $H_{src}.rf = H_{tgt}.rf$
- $H_{src}.to = H_{tgt}.to$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{src}.E, loc(i) = x \Rightarrow AccessMode(i) \in \{Release, Acquire\}$

and H_{src} is JAM_{21} -consistent.

Please see Appendix G for the proof.

► **THEOREM 11 (REMOVING PLAIN ACCESSES FOR NON-SHARED VARIABLE).** Let H_{tgt} be a JAM_{21} -consistent execution. Let x be a memory location and for all $i \in H_{tgt}.E$ such that $loc(i) = x$, $Tid(i) = t$ for some t . Let $a \notin H_{tgt}.E$. There is a H_{src} such that:

- $H_{src}.E = H_{tgt}.E \cup \{a\}$
- $H_{src}.loc(a) = x$
- $H_{src}.AccessMode(a) = \text{Plain}$
- $H_{src}.po \supset H_{tgt}.po$
- for all $i \in H_{src}.E$ such that $H_{src}.loc(i) = x$, $i \xrightarrow{po} a$ or $a \xrightarrow{po} i$
- $H_{src}.rf = H_{tgt}.rf$ if $a \in H_{src}.W \setminus RMW$, otherwise, $H_{src}.rf = H_{tgt}.rf \cup \{\langle i, a \rangle\}$ such that $(i \in H_{src}.W) \wedge (loc(i) = x) \wedge (i \xrightarrow{po} a) \wedge (\forall j \in H_{src}.E, (loc(j) = x) \wedge (j \xrightarrow{po} a) \Rightarrow (j \xrightarrow{po} i))$.
- $H_{src}.to = H_{tgt}.to$
- $H_{src}.IW = H_{tgt}.IW$

and H_{src} is JAM_{21} -consistent.

Please see Appendix G for the proof.

5.5.0.1 Counter Example

We now show a counter example for invalid register promotion on locations with Volatile accesses. Consider the following program:

```

Thread0 {
  int r1 = X.getOpaque(); // 1
  int r2 = X.getOpaque(); // 2
}

Thread1 {
  int r3 = Y.getOpaque(); // 1
  int r4 = Y.getOpaque(); // 2
}

Thread2 {
  X.setOpaque(2);
  Z.setVolatile(1);
  Y.setVolatile(1);
}

Thread3 {
  Y.setVolatile(2);
  X.setVolatile(1);
}


```

An execution with the annotated values in this program is not allowed by JAM_{21} . The execution graph before the transformation is shown in Fig. 8. First note that the Volatile access on z also has Release semantics due to the monotonicity of access modes, which yields the **ra** edge in Thread 2. The total order among **push** edges gives us two cases:


1. $Wz = 1 \xrightarrow{vvo} Wx = 1$. Since $Wx = 2 \xrightarrow{ra} Wz = 1$ and $ra \subseteq vvo$ and $vvo^+ \subseteq vo$, we have $Wx = 2 \xrightarrow{vo} Wx = 1$, which contradict with the **co** edge established by the observation from Thread 0.
2. $Wy = 2 \xrightarrow{vvo} Wy = 1$. This contradict with the **co** edge established by the observation from Thread 1.

In both cases there is a contradiction (a **co** cycle). Therefore, this execution is forbidden by JAM_{21} .

In this example, the memory location z is only accessed by Thread 2. It maybe tempting to promote z to a local register on Thread 2 to reduce the number of memory instructions, which yields the following program:



■ **Figure 8** Before Register Promotion on Volatile access (Forbidden)



■ **Figure 9** After Register Promotion on Volatile access (Allowed)

```

Thread0 {
    int r1 = X.getOpaque();
    int r2 = X.getOpaque();
}

Thread1 {
    int r3 = Y.getOpaque();
    int r4 = Y.getOpaque();
}

Thread2 {
    X.setOpaque(2);
    int z = 1;
    Y.setVolatile(1);
}

Thread3 {
    Y.setVolatile(2);
    X.setVolatile(1);
}

```

The execution graph after the transformation is shown in Fig. 9.

The annotated program behavior becomes allowed by JAM_{21} after the transformation. As the execution graph shows, since `Volatile` accesses also have cross-thread synchronization effect, we cannot simply weaken it to a `Plain` access without introducing new program behaviors.

5.6 Why are many transformations invalid for `Volatile`?

As we have shown, many local transformations are invalid for `Volatile` accesses under JAM_{21} . This is not a surprise and is intended to provide programmers a more intuitive semantics for `Volatile` accesses.

First, as we have confirmed with the author of [9], Java’s Access Modes intend equivalent semantics for `Volatile` mode and `fullFence()`. In this way, the programmers can easily understand the semantics of both once they understand `fullFence()`. To accurately capture this intention, JAM_{21} used a fence-based approach with `push` order to model `Volatile` mode. As we described in Section 3, `fullFence()` in Java has cross-thread synchronization effects. As a result, any local program transformation that removes a `Volatile` access from the execution graph may also remove its cross-thread synchronization, and might introduce new

program behavior after the transformation. Therefore, those transformations on `Volatile` accesses are mostly not allowed by JAM_{21} . On the other hand, the `sc` fence in C/C++11 [6] has slightly stronger synchronization effect than `sc` accesses so that they can be used to restore sequential consistency when inserted between every pair of accesses. Some of the transformations are allowed to apply to `sc` accesses but not to the fence version of the program.

In addition, restricting the set of possible transformations that is allowed to apply to `Volatile` variables can keep the coding process simple for programmers. From the programmers' perspective, one of the biggest challenges of developing and debugging concurrent programs comes from the compiler transformations that introduces surprising program behaviors that are not observable under sequential consistency. Therefore, restricting the set of possible transformations on `Volatile` accesses can restrict the set of surprising program behaviors that can happen when using `Volatile` mode, making the development process simpler. From this perspective, JAM_{21} provides more synchronization guarantees for `Volatile` mode than C/C++11 for `sc` mode atomic accesses.

Lastly, as we have confirmed with the author of [9], the current implementation of OpenJDK JVM does not apply those transformations on `Volatile` accesses.

6 Performance Implications

At the time of writing, the compiler bug [18] has been reported but still not resolved. The main argument against fixing the bug by inserting the missing fence instruction is that it may slow down the performance significantly. In this section, we argue that this is not the case.

The reason we only translated our `volatile-non-sc` example to Power instructions is that we only expect changes in the implementation of compilers targeting Power architectures. There is no need to change the Java compilers for x86 [15] and ARMv8 [13] all thanks to a property called *write atomicity*. Write atomicity, or *multicopy atomicity*, ensures that, when a write issued by a thread becomes observable by any other thread, it is observable by all other threads in the system. The issue that we demonstrate in this paper is caused by a write operation becoming visible to some threads before some other threads. Therefore, this violation of sequential consistency may only be observed when compiling to non-multicopy atomic architectures. If the underlying architecture ensures multicopy atomicity, then we can be sure that all writes are committed in a broadcast style and Release-Acquire semantics is sufficient. Since x86 [15] and ARMv8 [13] are multicopy atomic, we do not expect the incorrect program behavior to appear on those architectures. Therefore, no change is needed in compilers targeting multicopy-atomic architectures. In fact, we give a correctness proof for x86 in Appendix F to concretely show that the current compilation scheme to x86 is correct with respect to the x86-TSO memory model. Furthermore, the fence instruction that compilers use to compile to ARMv7 is the `DMB SY` instruction [8], which captures the same effects of a `fullFence()`. The only change that needs to be made is when compiling to Power instructions. This change might slow down some programs. However, relative to all other major factors that affect the performance of Java programs, we expect the impact by this change in compilers to be small.

Furthermore, symmetric to "leading fence" scheme, the "trailing fence" scheme is also valid. A correct compiler may choose to either of the schemes. Usually one may wish to choose the "trailing fence" scheme for better performance. In this case, comparing to the original compilation scheme, the fix only changes the compilation scheme for each `Volatile`

read:

1. Remove the `hwsync` in front of the `lwz` instruction
2. Change the `lwsync` following the `lwz` instruction to `hwsync`

It is easy to see that this fix only requires, in effect, moving the `hwsync` instructions that were originally inserted before the `lwz` instruction, but does not add more. In addition, it removes the `lwsync` instructions. Therefore, we do not expect this change to the compilation scheme to have much performance impact as argued in the discussions in the bug report [18].

On the other hand, the impact of this change for compiler optimizations is unclear. That is, whether this revised compilation scheme disables some of the compiler optimizations is still a question. However, since C/C++11 compilers has long adopted this compilation scheme and performance has always been the first priority in their implementations, the possibility of disabling optimisations is unlikely. We leave a detailed empirical study for future work.

7 Related Work

7.1 Sequential Consistency Issue in C/C++11

A similar but different issue in C/C++11 memory model for atomic operations with sequentially consistent memory order was pointed out by Manerkar, et al. [11] and Lahav, et al. [6]. In particular, when using the "trailing fence" convention for compiling to Power and ARMv7 on GCC, the intended sequentially consistent semantics for certain atomic accesses can be lost due to the different placement of fences in the programs. In other words, the previous C/C++11 memory model was not able to support the two existing compilation schemes on GCC. On the other hand, *JAM*₁₉ did not have the same problem. Since *JAM*₁₉ defined the semantics of Volatile mode in terms of `push` orders, which emulates the effect of a full fence, it already supports and aligned with the existing compilation scheme found on OpenJDK JVMs.

The problem, however, was that the existing compilation scheme does not give sufficient synchronization to some programs with all accesses marked as `Volatile`. Since *JAM*₁₉ models the problematic compilation scheme, it is necessary to repair the problem for both the compiler and the formal model.

7.2 Using Volatile to Restore Sequential Consistency in Java

Due to the complexity of the original Java Memory Model (JMM) [12], a class of bugs caused by missing “`volatile`” annotations on certain shared variables, called *missing-annotation bugs*, is found across real-world Java applications [10]. Aiming to improve the safety guarantees of the Java language, volatile-by-default JVM was proposed and developed by [10] to advocate the idea that variables should have `volatile` semantics by default and relaxed semantics by choice. Following their idea, the correctness of volatile (or Volatile mode, as they are equivalent) semantics become especially important. After all, if we cannot restore sequential consistency by annotating every variable as `volatile` (or use Volatile mode for every access), then volatile-by-default JVM would not be able to ensure intuitive program behaviors either. As of today, we are not aware of any `volatile`-by-default JVM for versions of Java after JDK9. Thus, we suggest that researchers carefully ensure the correctness of the `volatile` (or Volatile mode) implementations when implementing such JVM for Java versions after JDK9.

7.3 Memory Fairness and Compiler Transformations

Recently a declarative definition of *memory fairness* was proposed for axiomatic relaxed memory models [5]. As an improvement to the existing definition of *thread fairness*, the declarative memory fairness property can be easily integrated into axiomatic models with the No-Thin-Air restriction and can be used to prove the termination of concurrent algorithms. We noticed that the original *JAM* model [3] was published before this definition was proposed and therefore did not make any assertions regarding memory fairness. We leave it as our future work to verify whether memory fairness preserves the correctness of the compiler transformations and the compilation schemes.

8 Conclusion

In this paper, we have demonstrated that Java can use a compilation scheme that is similar to C/C++11. On the other hand, one should not simply compile Java's Access Modes the same way as C/C++11 compiles atomic memory orders since the formal memory models supports different compiler optimizations. In the future, we hope the bug can be resolved soon and the examples in this paper can be added to the Java Concurrency Stress Tests *jstress* [17] tool suite to aid in maintaining the correctness of the OpenJDK HotSpot implementations.

References

- 1 Jade Alglave, Luc Maranget, and Michael Tautschnig. Herding cats: Modelling, simulation, testing, and data mining for weak memory. *ACM Trans. Program. Lang. Syst.*, 36(2), July 2014. doi:10.1145/2627752.
- 2 ARM ARM. Architecture reference manual-armv8, for armv8-a architecture profile. *ARM Limited, Dec*, 2017.
- 3 John Bender and Jens Palsberg. A formalization of java's concurrent access modes. *Proc. ACM Program. Lang.*, 3(OOPSLA), October 2019. doi:10.1145/3360568.
- 4 Peter Sewell Jaroslav Sevcik. C/c++11 mappings to processors. Technical report, University of Cambridge, 10 2016. URL: <https://www.cl.cam.ac.uk/~pes20/cpp/cpp0xmappings.html>.
- 5 Ori Lahav, Egor Namakonov, Jonas Oberhauser, Anton Podkopaev, and Viktor Vafeiadis. Making weak memory models fair. *arXiv preprint arXiv:2012.01067*, 2020.
- 6 Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. Repairing sequential consistency in c/c++11. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017*, page 618–632, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3062341.3062352.
- 7 L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Comput.*, 28(9):690–691, September 1979. doi:10.1109/TC.1979.1675439.
- 8 Doug Lea. The jsr-133 cookbook for compiler writers. <http://gee.cs.oswego.edu/dl/jmm/cookbook.html>, 2011. Last modified: Tue Mar 22 07:11:36 2011.
- 9 Doug Lea. Using jdk 9 memory order modes. <http://gee.cs.oswego.edu/dl/html/j9mm.html>, 2018. Last Updated: Fri Nov 16 08:46:48 2018.
- 10 Lun Liu, Todd Millstein, and Madanlal Musuvathi. A volatile-by-default jvm for server applications. *Proc. ACM Program. Lang.*, 1(OOPSLA), October 2017. doi:10.1145/3133873.
- 11 Yatin A Manerkar, Caroline Trippel, Daniel Lustig, Michael Pellauer, and Margaret Martonosi. Counterexamples and proof loophole for the c/c++ to power and armv7 trailing-sync compiler mappings. *arXiv preprint arXiv:1611.01507*, 2016.
- 12 Jeremy Manson, William Pugh, and Sarita V. Adve. The java memory model. *SIGPLAN Not.*, 40(1):378–391, January 2005. doi:10.1145/1047659.1040336.

- 13 Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. Simplifying arm concurrency: Multicopy-atomic axiomatic and operational models for armv8. *Proc. ACM Program. Lang.*, 2(POPL), December 2017. doi:10.1145/3158107.
- 14 Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. Understanding power multiprocessors. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '11, page 175–186, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/1993498.1993520.
- 15 Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. X86-tso: A rigorous and usable programmer’s model for x86 multiprocessors. *Commun. ACM*, 53(7):89–97, July 2010. doi:10.1145/1785414.1785443.
- 16 Dennis Shasha and Marc Snir. Efficient and correct execution of parallel programs that share memory. *ACM Trans. Program. Lang. Syst.*, 10(2):282–312, April 1988. URL: <http://doi.acm.org/10.1145/42190.42277>, doi:10.1145/42190.42277.
- 17 Aleksey Shipilev. jcstress - the java concurrency stress tests. <https://wiki.openjdk.java.net/display/CodeTools/jcstress>, 2017. Last Updated: Wed Dec 05 13:55 2018.
- 18 Aleksey Shipilev. [JDK-8262877] PPC sequential consistency problem: volatile stores are too weak. Technical report, OpenJDK Bug System, 03 2021. URL: <https://bugs.openjdk.java.net/browse/JDK-8262877>.
- 19 Conrad Watt, Christopher Pulte, Anton Podkopaev, Guillaume Barbier, Stephen Dolan, Shaked Flur, Jean Pichon-Pharabod, and Shu-yu Guo. Repairing and mechanising the javascript relaxed memory model. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2020, page 346–361, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3385412.3385973.

A the Full JAM_{21} Model

```

let opq = O | RA | V
let rel = W & (RA | V)
let acq = R & (RA | V)
let f_rel = REL | V
let f_acq = ACQ | V
let vol = V
let fence = F

(* volatile accesses extend push order *)
let svo = po;[fence & f_rel];po;[W] | [R];po;[fence & f_acq];po
let spush = po;[fence & vol];po

(* release acquire ordering *)
let ra = po;[rel] | [acq];po

(* intra thread volatile ordering *)
let volint = [vol];po;[vol] (* OLD: po;[vol & R] | [vol & W];po *)

(* intrathread ordering constraints *)
let into = svo | spush | ra | volint

(* cross thread push ordering extended with volatile memory accesses *)
let push = spush | volint
with pushto from linearisations(domain(push), ((W * FW) & loc & ~id) | rf | po)

(* extend ra visibility *)
let vvo = rf | svo | ra | push | pushto;push
let vo = vvo+ | po-loc

include "filters.cat" (* includes WW filter *)
let WWco(rel) = WW(rel) & loc & ~id
(* final writes are co-after everything *)
let cofw = WWco((W * FW))

(* jam coherence *)
let coww = WWco(vo)
let cowl = WWco(vo;inrf)
let corw = WWco(vo;po)
let corr = [opq] ; WWco(rf;po;inrf) ; [opq]
let coint = loc & IW*(W\IW)

include "cross.cat"
let co0 = loc & (IW * (W \ IW)|(W \ FW) * FW)
with cormwtotal from generate_orders(RMW, co0)

let rec co-jom = coww | cowl | corw | corr | cormwtotal
  | WWco((rf;[RMW])^-1;co-jom) | coint | cofw

acyclic (po | rf) ; [opq] as no-thin-air
acyclic co-jom as coherence

```

B the Full JAM'_{21} Model

```

include "filters.cat"
include "cross.cat"
let WWco(rel) = WW(rel) & loc & ~id
let co0 = loc & (IW * (W \ IW) | (W \ FW) * FW)
with cormwttotal from generate_orders(RMW, co0)

let opq = 0 | RA | V
let rel = W & (RA | V)
let acq = R & (RA | V)
let f_rel = REL | V
let f_acq = ACQ | V
let vol = V
let fence = F

(* volatile accesses extend push order *)
let svo = po; [fence & f_rel]; po; [W] | [R]; po; [fence & f_acq]; po
let spush = po; [fence & vol]; po

(* release acquire ordering *)
let ra = po; [rel] | [acq]; po

(* intra thread volatile ordering *)
let volint = [vol]; po; [vol] (* OLD: po; [vol & R] | [vol & W]; po *)

(* intrathread ordering constraints *)
let into = svo | spush | ra | volint
let push = spush | volint

let rec co-jom = cown | cownr | corw | corr | cormwttotal
  | WWco((rf; [RMW])^-1; co-jom) | coint | cofw

and fr-jom = rf^-1 ; co-jom
and fr-jom-e = fr-jom & ext
and co-jom-e = co-jom & ext
and chapo = rfe | fr-jom-e | co-jom-e | (fr-jom-e ; rfe) | (co-jom-e ; rfe)

(* extend ra visibility *)
and vvo = rf | svo | ra | push | push ; chapo ; push
and vo = vvo+ | po-loc

and cofw = WWco((W * FW))
and cown = WWco(vo)
and cownr = WWco(vo; invrf)
and corw = WWco(vo; po)
and corr = [opq] ; WWco(rf; po; invrf) ; [opq]
and coint = loc & IW*(W\IW)

acyclic (po | rf) ; [opq] as no-thin-air
acyclic co-jom as coherence

```

C The Power Memory Model in Herd7

```

PPC
(* Model for Power *)
include "cos.cat" (* Used to compute the coherence order*)

(* Uniproc *)
acyclic po-loc | rf | fr | co as scperlocation

(* Atomic *)
empty rmw & (fre;coe) as atomic

(* Utilities *)
let dd = addr | data
let rdw = po-loc & (fre;rfe)
let detour = po-loc & (coe ; rfe)
let addrpo = addr;po

(*****)
(* ppo *)
(*****)

let sync = try fencerel(SYNC) with 0
let lwsync = try fencerel(LWSYNC) with 0
let eieio = try fencerel(EIEIO) with 0
let isync = try fencerel(ISYNC) with 0
show sync,lwsync,eieio

(* Dependencies *)
show data,addr
let ctrlisync = try ctrlfence(ISYNC) with 0
show ctrlisync
show isync ctrlisync as isync
show ctrl ctrlisync as ctrl
show isync,ctrlisync

(* Initial value *)
let ci0 = ctrlisync | detour
let ii0 = dd | rfi | rdw
let cc0 = dd | po-loc | ctrl | addrpo
let ic0 = 0

(* Fixpoint from i -> c in instructions and transitivity *)
let rec ci = ci0 | (ci;ii) | (cc;ci)
and ii = ii0 | ci | (ic;ci) | (ii;ii)
and cc = cc0 | ci | (ci;ic) | (cc;cc)
and ic = ic0 | ii | cc | (ic;cc) | (ii ; ic) (* | ci inclus dans ii et cc *)

let ppo =
  let ppoR = ii & (R * R)
  and ppoW = ic & (R * W) in
  ppoR | ppoW

(* fences *)

let lwsync = lwsync (W * R)
let eieio = eieio & (W * W)

(* All arm barriers are strong *)
let strong = sync
let light = lwsync|eieio

let fence = strong|light

(* happens before *)
let hb = ppo | fence | rfe

```

23:30 Compiling Volatile Correctly in Java

```
acyclic hb as thinair

(* prop *)
let hbstar = hb*
let propbase = (fence|(rfe;fence));hbstar

let chapo = rfe|fre|coe|(fre;rfe)|(coe;rfe)

let prop = propbase & (W * W) | (chapo? ; propbase*; strong; hbstar)

acyclic co|prop as propagation
irreflexive fre;prop;hbstar as observation

let xx = po & (X * X)
acyclic co | xx as scXX
```

D A Proof of Compilation Correctness to Power

D.1 The JAM'_{21} Model

The motivation of JAM'_{21} is to enable simpler compilation proofs. JAM_{21} enforces a total order among `fullFence()`s, which introduces complexity when proving compilation correctness. Therefore, we introduce an intermediate memory model JAM'_{21} that is observationally equivalent to the JAM_{21} . Thus, we start by defining and proving the observational equivalence of the two models. Then we use JAM'_{21} to prove the correctness of the compilation schemes.

The JAM'_{21} model is the same as JAM_{21} except for the semantics of full fences. Instead of having a total order on full fences, JAM'_{21} only enforces order when there is a communication edge. The full semantics of JAM'_{21} can be found in Appendix B. Here we only include the updated portion. The definition for `chapo` is newly added⁶. The cross-thread synchronization effect of `fullFence()`s is then defined as `push; chapo; push` (instead of `pushto; push` as before):

$$\begin{aligned} \text{chapo} &\triangleq \text{rfe} \mid \text{fre} \mid \text{coe} \mid (\text{fre} ; \text{rfe}) \mid (\text{coe} ; \text{rfe}) \\ \text{vvo} &\triangleq \dots \mid \text{push} ; \text{chapo} ; \text{push} \end{aligned}$$

The rest of JAM'_{21} are the same as JAM_{21} .

► **DEFINITION 10** (JAM'_{21} CONSISTENCY). An execution history H is JAM'_{21} -**consistent** if it is trace coherent and satisfies the following two requirements:

1. **NO-THIN-AIR**: `po` \mid `rf` is acyclic. $\text{acyclic}(\overrightarrow{-\text{po} \mid \text{rf}})$
2. **COHERENCE**: `co-jom` is acyclic, $\text{acyclic}(\overrightarrow{-\text{co-jom}})$

We say such an execution history H is **allowed** by JAM'_{21} . Otherwise, it is **forbidden**.

► **LEMMA 2**. Observational equivalence is a transitive relation. That is, if H and H' are observationally equivalent, and H' and H'' are observationally equivalent, then H and H'' are observationally equivalent.

Proof. The transitivity follows directly from the transitivity of set equivalence in the definition. ◀

► **LEMMA 3**. (\Rightarrow) Given a program P , for any JAM_{21} -consistent execution $H \in \text{Histories}_{JAM}(P)$, there exists a $H' \in \text{Histories}_{JAM'}$ such that H' is observationally equivalent to H .

Proof. Given a program P , it's obvious that there exists an H' that is observationally equivalent to H . We prove that $H' \in \text{Histories}_{JAM'}(P)$. That is, let H' be a candidate execution of P and H' is observationally equivalent to H . We show that H' is JAM'_{21} -consistent. Since the only difference between JAM_{21} and JAM'_{21} is at the effect of full fences, we focus on this part in our proof. In particular, we first show that, if $i_1 \xrightarrow{\text{push}} i_3$, $i_2 \xrightarrow{\text{push}} i_4$, and $i_3 \xrightarrow{\text{chapo}} i_2$, then it must be that $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$. We can prove this by analyzing five cases:

⁶ The name `chapo` comes directly from the Power Memory Model in the Herd [1] repository. We use the same name here so that the readers can easily relate them

1. $i_3 \xrightarrow{\text{rf}} i_2$: then we have $i_1 \xrightarrow{\text{push}} i_3 \xrightarrow{\text{rf}} i_2 \xrightarrow{\text{push}} i_4$, which is equivalent to $i_1 \xrightarrow{\text{vvo}} i_3 \xrightarrow{\text{vvo}} i_2 \xrightarrow{\text{vvo}} i_4$, which means $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$ (as it'd create a **vo** cycle in the latter case).
2. $i_3 \xrightarrow{\text{co}} i_2$: then it cannot be $i_2 \xrightarrow{\text{vo}} i_3$ on the right side as it immediately gives us a coherence cycle by **coww**. So it must be $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$.
3. $i_3 \xrightarrow{\text{fr}} i_2$: then there exists a write event W such that $W \xrightarrow{\text{rf}} i_3$ and $W \xrightarrow{\text{co}} i_2$. If we have $i_2 \xrightarrow{\text{vo}} i_3$ then we would have $i_2 \xrightarrow{\text{co}} W$ by **cowr**, which gives us a coherence cycle. Therefore it must be that $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$.
4. $i_3 \xrightarrow{\text{co}} W_1 \xrightarrow{\text{rf}} i_2$: if we have $i_2 \xrightarrow{\text{vo}} i_3$ then, because **rf** is also a visibility order, we have $W_1 \xrightarrow{\text{vo}} i_3$. By **coww** we get $W_1 \xrightarrow{\text{co}} i_3$, contradicting with the earlier assumption that $i_3 \xrightarrow{\text{co}} W_1$. So it must be that $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$.
5. $i_3 \xrightarrow{\text{fr}} W_1 \xrightarrow{\text{rf}} i_2$: then it means there is W_2 such that $W_2 \xrightarrow{\text{rf}} i_3$ and $W_2 \xrightarrow{\text{co}} W_1$. If we have $i_2 \xrightarrow{\text{vo}} i_3$, since **rf** is also a visibility order, we have $W_1 \xrightarrow{\text{vo}} i_3$. By **cowr**, we have $W_1 \xrightarrow{\text{co}} W_2$, contradicting with the assumption of $W_2 \xrightarrow{\text{co}} W_1$ earlier. So it must be $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$.

Thus we have shown that if we have $i_1 \xrightarrow{\text{push}} i_3$, $i_2 \xrightarrow{\text{push}} i_4$, and $i_3 \xrightarrow{\text{chapo}} i_2$, then it must be that $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$.

We now show that H' is JAM'_{21} -consistent. The No-Thin-Air requirement is automatically fulfilled since $H.E = H'.E$, $H.\text{po} = H'.\text{po}$ and $H.\text{rf} = H'.\text{rf}$.

Previously, we have shown that $i_1 \xrightarrow{\text{push}} i_3$, $i_2 \xrightarrow{\text{push}} i_4$, and $i_3 \xrightarrow{\text{chapo}} i_2$ implies $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$ in H . Since H is JAM_{21} -consistent, then **co-jom** is acyclic with either $i_1 \xrightarrow{\text{vo}} i_4$ or $i_2 \xrightarrow{\text{vo}} i_3$. Now we have two cases:

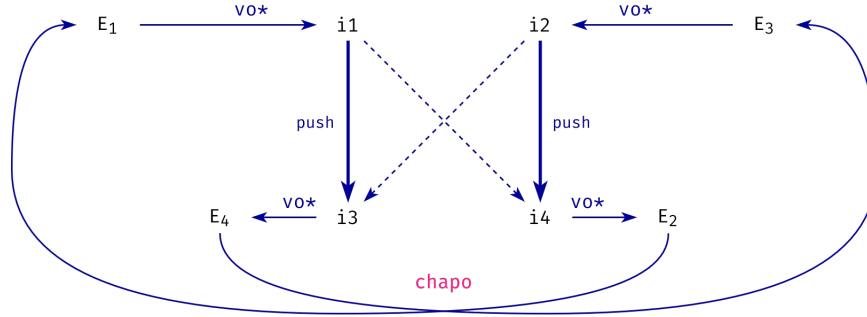
1. No communication: then we do not have any extra **vo** edge we can use to infer in H' either. Since **co-jom** is acyclic in H , and $H.\text{vo} = H'.\text{vo}$, **co-jom** is acyclic in H' , too.
2. With communication **chapo**: suppose $i_3 \xrightarrow{\text{chapo}} i_2$ (the other direction is symmetrically the same), then it must be that $i_1 \xrightarrow{\text{vo}} i_4$ and not $i_2 \xrightarrow{\text{vo}} i_3$ in H . Since H is JAM_{21} -consistent, we know that $i_1 \xrightarrow{\text{vo}} i_4$ cannot lead to any **co-jom** cycle. In H' , we use (**vo-5'**) to infer that $i_1 \xrightarrow{\text{vo}} i_4$. Since other portions of H' satisfies the above conditions, we can infer that $i_1 \xrightarrow{\text{vo}} i_4$ cannot lead to any **co-jom** cycle in H' either.

Since neither case leads to a **co-jom** cycle in H' , we can conclude that H' is JAM'_{21} -consistent and hence $H' \in \text{Histories}_{JAM'}(P)$ ◀

► **LEMMA 4.** Let **fullfence-vo** order be the **vo** order derived using the rule (**vo-5**) of JAM_{21} . In JAM_{21} , for any **co-jom** cycle derived from the **fullfence-vo** orders, there is a **vo***; **fullfence-vo**; **vo***; **chapo** cycle.

Proof. Let i_1, i_2, i_3 and i_4 be four events in an execution history H such that $i_1 \xrightarrow{\text{push}} i_3$ and $i_2 \xrightarrow{\text{push}} i_4$. By (**vo-5**), we derive the condition $(i_1 \xrightarrow{\text{fullfence-vo}} i_4) \vee (i_2 \xrightarrow{\text{fullfence-vo}} i_3)$. Suppose that H is not consistent due to this condition under JAM_{21} , i.e., following either side of the disjunction we can derive a **co-jom** cycle. We analyze one side of the disjunction since the other side of the disjunction symmetrically follow the same reasoning. We analyze each possible rule to derive a coherence cycle. Since **fullfence-vo** is included in **vo** order, we only need to analyze the cases where **vo** appears:

- **coww**. If we derived the coherence cycle from **coww** rule, then it means there exists W_1 and W_2 such that $W_1 \xrightarrow{\text{vo}} W_2$ and $W_2 \xrightarrow{\text{co}} W_1$. The fact that $i_1 \xrightarrow{\text{vo}} i_4$ "enables" us (we can only use rule (**vo-7**) here as other visibility rules imply program structures that does not



■ **Figure 10** $i_1 \xrightarrow{\text{push}} i_3 \xrightarrow{\text{vo}^*} E_4 \xrightarrow{\text{chapo}} E_3 \xrightarrow{\text{vo}^*} i_2 \xrightarrow{\text{push}} i_4 \xrightarrow{\text{vo}^*} E_2 \xrightarrow{\text{chapo}} E_1$ cycle

require $i_1 \xrightarrow{\text{vo}} i_4$) to derive this contradiction implies the following structure: $W_1 \xrightarrow{\text{vo}^*} i_1 \xrightarrow{\text{vo}} i_4 \xrightarrow{\text{vo}^*} W_2 \xrightarrow{\text{co}} W_1$, revealing a cycle of vo^* ; **fullfence-vo**; vo^* ; **chapo**.

- **cowr**. If we derived the coherence cycle from **cowr** rule, then it means there exists R_1 , W_1 , and W_2 such that $W_1 \xrightarrow{\text{rf}} R_1$, $W_2 \xrightarrow{\text{vo}} R_1$, and $W_1 \xrightarrow{\text{co}} W_2$. The fact that $i_1 \xrightarrow{\text{vo}} i_4$ "enables" us to derive this contradiction implies the following structure: $W_2 \xrightarrow{\text{vo}^*} i_1 \xrightarrow{\text{vo}} i_4 \xrightarrow{\text{vo}^*} R_1$. Because $W_1 \xrightarrow{\text{co}} W_2$ and $W_1 \xrightarrow{\text{rf}} R_1$, we have $R_1 \xrightarrow{\text{fr}} W_2$. We now have a cycle $W_2 \xrightarrow{\text{vo}^*} i_1 \xrightarrow{\text{vo}} i_4 \xrightarrow{\text{vo}^*} R_1 \xrightarrow{\text{fr}} W_2$, which is a cycle of vo^* ; **fullfence-vo**; vo^* ; **chapo**.
- **corw**. If we derived the coherence cycle from **corw** rule, then it means there exists R_1 , W_1 , and W_2 , such that $W_1 \xrightarrow{\text{rf}} R_1$, $R_1 \xrightarrow{\text{vo}} W_2$, and $W_2 \xrightarrow{\text{co}} W_1$. The fact that $i_1 \xrightarrow{\text{vo}} i_4$ "enables" us to derive this contradiction implies the following structure: $R_1 \xrightarrow{\text{vo}^*} i_1 \xrightarrow{\text{vo}} i_4 \xrightarrow{\text{vo}^*} W_2 \xrightarrow{\text{co}} W_1 \xrightarrow{\text{rf}} R_1$, which is a cycle of vo^* ; **fullfence-vo**; vo^* ; **chapo**.

◀

► **LEMMA 5.** (\Leftarrow) Given a program P , for any JAM'_{21} -consistent execution $H' \in \text{Histories}_{JAM'}(P)$, there exists an execution history $H \in \text{Histories}_{JAM}(P)$ such that H is observationally equivalent to H' .

Proof. Given a program P , it's obvious that there exists an H that is observationally equivalent to H' . We prove that $H \in \text{Histories}_{JAM}(P)$. That is, let H be a candidate execution of P and H is observationally equivalent to H' . We show that H is JAM_{21} -consistent. To help the reader better understand this, consider Fig. 10. Suppose H is an execution history that is forbidden by the rules of JAM_{21} , we show that its corresponding H' is also forbidden by JAM'_{21} . Since the only difference between JAM_{21} and JAM'_{21} is at the effects of full fences, we only analyze that part. In other words, H is forbidden by JAM_{21} precisely due to the total order of full fences. Let i_1 , i_2 , i_3 , and i_4 be events in H such that $i_1 \xrightarrow{\text{push}} i_3$ and $i_2 \xrightarrow{\text{push}} i_4$. By Lemma 4, we can generalize the structure and infer that there are E_1 , E_2 , E_3 , and E_4 such that $E_1 \xrightarrow{\text{vo}^*} i_1$, $E_3 \xrightarrow{\text{vo}^*} i_2$, $i_4 \xrightarrow{\text{vo}^*} E_2$, and $i_3 \xrightarrow{\text{vo}^*} E_4$. In addition, we also have $E_2 \xrightarrow{\text{chapo}} E_1$ and $E_4 \xrightarrow{\text{chapo}} E_3$. Because H is forbidden under JAM_{21} , it means we have two cycles, $i_1 \xrightarrow{\text{vo}} i_4 \xrightarrow{\text{vo}^*} E_2 \xrightarrow{\text{chapo}} E_1 \xrightarrow{\text{vo}^*} i_1$ and $i_2 \xrightarrow{\text{vo}} i_3 \xrightarrow{\text{vo}^*} E_4 \xrightarrow{\text{chapo}} E_3 \xrightarrow{\text{vo}^*} i_2$, so that no matter which side of the disjunction we choose, we always end up with a contradiction. In H' , on the other hand, we do not have $i_1 \xrightarrow{\text{vo}} i_4$ or $i_2 \xrightarrow{\text{vo}} i_3$. However, despite the absence of the two edges, we now have a larger cycle: $i_1 \xrightarrow{\text{push}} i_3 \xrightarrow{\text{vo}^*} E_4 \xrightarrow{\text{chapo}} E_3 \xrightarrow{\text{vo}^*} i_2 \xrightarrow{\text{push}} i_4 \xrightarrow{\text{vo}^*} E_2 \xrightarrow{\text{chapo}} E_1$, which forms a vo cycle by (vo -5') in JAM'_{21} . Therefore, execution history H' is forbidden under JAM'_{21} .

as well, which contradicts to our previous assumption. Thus, since $H' \in \text{Histories}_{JAM'}(P)$ implies that H' is JAM'_{21} -consistent, $H \in \text{Histories}_{JAM}(P)$. ◀

Essentially, the replacement of the (vo-5) rule give no actual effect in forbidding executions. In JAM_{21} , we look for two vo cycles to forbid an execution, whereas in JAM'_{21} we combine the two cycles into one to forbid the execution.

▶ **THEOREM 12 (OBSERVATIONAL EQUIVALENCE OF JAM_{21} AND JAM'_{21}).** JAM'_{21} is observationally equivalent to JAM_{21} .

Proof. Given a program P , let $\text{Histories}_{JAM}(P)$ be the set of JAM_{21} -consistent candidate executions of P and $\text{Histories}_{JAM'}(P)$ be the set of JAM'_{21} -consistent candidate executions of P . By Lemma 3, we know that for all $H \in \text{Histories}_{JAM}(P)$, there exists an $H' \in \text{Histories}_{JAM'}(P)$ such that H' and H are observationally equivalent. Similarly, by Lemma 5, we know that for all $H' \in \text{Histories}_{JAM'}(P)$, there exists an $H \in \text{Histories}_{JAM}(P)$ such that H and H' are observationally equivalent. Combining together, we can conclude that JAM_{21} and JAM'_{21} are observationally equivalent. ◀

▶ **COROLLARY 3.** JAM'_{21} satisfies the same important properties (Theorem 14, Theorem 15, Theorem 16, Theorem 17 and Corollary 4) in Section. H.

Proof. Since the definitions in JAM'_{21} are the same as JAM_{21} except for the semantics of fullFences, JAM'_{21} automatically satisfies Theorem. 14, Theorem. 15, Theorem 16, and Theorem. 17. By Theorem. 12, JAM_{21} and JAM'_{21} allow the same set of execution histories up to observational equivalence. Therefore, JAM'_{21} also satisfies Corollary. 4. ◀

D.2 Compilation to Power

▶ **LEMMA 1 (JAM'_{21} TO POWER).** Let P_{src} be a Java program, P_{tgt} be the Power program compiled from P_{src} using the compilation scheme in Fig. 4 (with the leading fence convention). For all $H_{tgt} \in \text{Histories}_{Power}(P_{tgt})$ there exists a $H_{src} \in \text{Histories}_{JAM'}(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

Proof. It is obvious that there exists a candidate execution history H_{src} of P_{src} such that $H_{src} \rightsquigarrow H_{tgt}$. We show that $H_{src} \in \text{Histories}_{JAM'}(P)$. That is, H_{src} is JAM'_{21} -consistent. In order to be consistent under JAM'_{21} , we need H_{src} to satisfy two requirements:

1. **NO-THIN-AIR Requirement:** (po | rfi) is acyclic. The intra-thread rfi order cannot contradicting the po given that H_{tgt} is Power-consistent. Therefore, we need to show that (po | rfe) is acyclic. Note that since the only inter-thread order is rfe. If there is a cycle in (po | rfe), then the head of each thread is a R^{Opq} and the last event in each thread participating in this cycle is a W^{Opq} , where $R^{Opq} \xrightarrow{po} W^{Opq}$ in H_{src} . Using a compiler that follows the compilation scheme, this translates to $R \xrightarrow{ctrl} W$ in H_{tgt} . Further we can infer that $R \xrightarrow{hb} W$ in H_{tgt} . Power ensures that (hb | rfi) is acyclic. Therefore, if there is a cycle of (po | rfi) in H_{src} , H_{tgt} would not be consistent under Power's memory model, contradicting to our previous assumption. For readers who do not care about this guarantee, getOpaque() can be directly compiled to a lwz instruction.
2. **COHERENCE Requirement:** co-jom is acyclic. We now prove that co-jom is acyclic in H_{src} . In order to show this, we show that co-jom is a partial order of co in H_{tgt} . In other words, co is a linear extension of co-jom. We prove this by assuming the opposite and deriving a contradiction.

Suppose that there exists $i_1 \xrightarrow{\text{co-jom}} i_2$ in H_{src} but $i_2 \xrightarrow{\text{co}} i_1$ in H_{tgt} . We analyze each of the possible cases where we can derive a **co-jom** order.

- **coinit**. This automatically gives us a contradiction since $H_{src}.\text{IW} = H_{tgt}.\text{IW}$.
- **cofw**. This automatically gives us a contradiction since $H_{src}.\text{FW} = H_{tgt}.\text{FW}$.
- **corr**. This implies that there exist R_1 and R_2 such that, $R_1 \xrightarrow{\text{po}} R_2$, $i_1 \xrightarrow{\text{rf}} R_1$, and $i_2 \xrightarrow{\text{rf}} R_2$. We also have $i_2 \xrightarrow{\text{co}} i_1$. Now the SC-per-location requirement of Power is violated, contradicting with our previous assumption that H_{tgt} is Power-consistent.
- **coww**. This implies that $i_1 \xrightarrow{\text{vo}} i_2$ but $i_2 \xrightarrow{\text{co}} i_1$.
 - $i_1 \xrightarrow{\text{po-loc}} i_2$. $i_2 \xrightarrow{\text{co}} i_1$ would violate the SC-Per-Location Requirement of Power, making H_{tgt} inconsistent, contradicting to our previous assumption.
 - $i_1 \xrightarrow{\text{svo}} i_2$ or $i_1 \xrightarrow{\text{ra}} i_2$ or $i_1 \xrightarrow{\text{push}} i_2$. In H_{tgt} , this means $i_1 \xrightarrow{\text{lwsync}} i_2$ or $i_1 \xrightarrow{\text{sync}} i_2$. Note that all three cases of them are included in program order with access to the same location. Therefore, $i_2 \xrightarrow{\text{co}} i_1$ would violate the SC-Per-Location Requirement of Power, making H_{tgt} inconsistent, contradicting to our previous assumption.
 - $i_1 \xrightarrow{\text{push}} E_1 \xrightarrow{\text{chapo}} E_2 \xrightarrow{\text{push}} i_2$ for some E_1 and E_2 . Note that we have $E_1 \xrightarrow{\text{chapo}} E_2 \xrightarrow{\text{sync}} i_2$ and $i_2 \xrightarrow{\text{co}} i_1 \xrightarrow{\text{sync}} E_1$. They form a propagation (**prop**) cycle between i_2 and E_1 , which makes H_{tgt} not Power-consistent, giving us a contradiction.
 - $i_1 \xrightarrow{\text{vvo}} i_3 \xrightarrow{\text{vvo}} i_2$. This corresponds to the inductive case where the visibility order is formed by two visibility orders through another event, i_3 . Note that all the cases of **vo** orders produce propagation (**prop**) orders. Therefore, we get a violation of the Propagation requirement if $i_2 \xrightarrow{\text{co}} i_1$ in Power, contradicting to our previous assumption that H_{tgt} is Power-consistent.
- **cowr**. This implies that there exists R such that $i_2 \xrightarrow{\text{rf}} R$ and $i_1 \xrightarrow{\text{vo}} R$.
 - $i_1 \xrightarrow{\text{po-loc}} i_2$. $i_2 \xrightarrow{\text{co}} i_1$ would violate the SC-Per-Location Requirement of Power, making H_{tgt} inconsistent, contradicting to our previous assumption.
 - $i_1 \xrightarrow{\text{svo}} R$ or $i_1 \xrightarrow{\text{ra}} R$ or $i_1 \xrightarrow{\text{push}} R$. In H_{tgt} , this means $i_1 \xrightarrow{\text{lwsync}} i_2$ or $i_1 \xrightarrow{\text{sync}} i_2$. Note that all three cases of them are included in program order with access to the same location. Therefore, $i_2 \xrightarrow{\text{co}} i_1$ would violate the SC-Per-Location Requirement of Power, making H_{tgt} inconsistent, contradicting to our previous assumption.
 - $i_1 \xrightarrow{\text{push}} E_1 \xrightarrow{\text{chapo}} E_2 \xrightarrow{\text{push}} R$ for some E_1 and E_2 . Because $i_2 \xrightarrow{\text{rf}} R$ and $i_2 \xrightarrow{\text{co}} i_1$, we have $R \xrightarrow{\text{fr}} i_1$. We now have $E_1 \xrightarrow{\text{chapo}} E_2 \xrightarrow{\text{sync}} R$ and $R \xrightarrow{\text{fr}} i_1 \xrightarrow{\text{sync}} E_1$. Both of them form a propagation order between R and E_1 and they contradict with each other.
 - $i_1 \xrightarrow{\text{vvo}} i_3 \xrightarrow{\text{vvo}} R$. This corresponds to the inductive case where the visibility order is formed by two visibility orders through another event, i_3 . Note that all the cases of **vvo** orders produce propagation orders. Therefore, we get $R \xrightarrow{\text{fr}} i_1 \xrightarrow{\text{prop}} R$, which violates the observation requirement in Power.
- **corw**. This implies that there exists R such that $R \xrightarrow{\text{po}} i_2$ and $i_1 \xrightarrow{\text{rf}} R$. $i_2 \xrightarrow{\text{co}} i_1$ would cause a violation of SC-Per-Location requirement in Power.
- **cormwexcl**. This implies that i_1 is a *RMW* operation and there exists i_3 such that $i_3 \xrightarrow{\text{rf}} i_1$ and $i_3 \xrightarrow{\text{co}} i_2$. Having $i_2 \xrightarrow{\text{co}} i_1$ immediately violates the atomicity requirement of Power.
- **cormwtotal**. $i_1 \xrightarrow{\text{co}} i_2$ because $H_{tgt}.\text{co} \subseteq H_{src}.\text{to}$. Therefore, having $i_2 \xrightarrow{\text{co}} i_1$ in this case would yield a cycle in **co**, violating the propagation requirement in Power.

Thus, we have shown that `co-jom` is a partial order of `co` in H_{tgt} and the coherence requirement is automatically fulfilled because `co` is acyclic. Hence, $H_{src} \in \text{Histories}_{JAM'}(P)$. \blacktriangleleft

► **THEOREM 1 (COMPILATION CORRECTNESS TO POWER (LEADING FENCE CONVENTION))**. The compilation from Java to Power following the compilation scheme in Fig. 4 (using the leading fence convention) is correct. That is, let P_{src} be a Java program, P_{tgt} be the Power program compiled from P_{src} using the compilation scheme in Fig. 4 (using the leading fence convention). For all $H_{tgt} \in \text{Histories}_{Power}(P_{tgt})$ there exists a $H_{src} \in \text{Histories}_{JAM}(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

Proof. By Lemma 1, we know that there exists an $H'_{src} \in \text{Histories}_{JAM'}(P_{src})$ such that $H'_{src} \rightsquigarrow H_{tgt}$. Therefore, by definition of the \rightsquigarrow relation,

- H_{tgt} is observationally equivalent to H'_{src}
- $H_{tgt}.\text{co} \subseteq H'_{src}.\text{to}$
- If $RMW, i_1 \in H'_{src}.\text{E}$ and $RMW \xrightarrow{\text{po}} i_1$, then $RMW \xrightarrow{\text{ctrl}} i_1$ in H_{tgt}
- If $R^{\text{Opq}}, i_1 \in H'_{src}.\text{E}$ and $R^{\text{Opq}} \xrightarrow{\text{po}} i_1$, then $R \xrightarrow{\text{ctrl}} i_1$ in H_{tgt}
- If $i_1, i_2 \in H'_{src}.\text{E}$ and $i_1 \xrightarrow{\text{push}} i_2$, then $i_1 \xrightarrow{\text{sync}} i_2$ for $i_1, i_2 \in H_{tgt}.\text{E}$
- If $i_1, i_2 \in H'_{src}.\text{E}$ and $i_1 \xrightarrow{\text{ra}} i_2$, then $i_1 \xrightarrow{\text{lwsync}} i_2$ for $i_1, i_2 \in H_{tgt}.\text{E}$

By Theorem 12, we know that for all H'_{src} , there exists an $H_{src} \in \text{Histories}_{JAM}(P)$ such that H_{src} is observationally equivalent to H'_{src} . By Lemma 2, H_{src} is observationally equivalent to H_{tgt} . Furthermore,

- $H_{tgt}.\text{co} \subseteq H_{src}.\text{to}$ because $H_{src}.\text{to} = H'_{src}.\text{to}$.
- If $RMW, i_1 \in H_{src}.\text{E}$ and $RMW \xrightarrow{\text{po}} i_1$, then $RMW \xrightarrow{\text{ctrl}} i_1$ in H_{tgt} because $H_{src}.\text{E} = H'_{src}.\text{E}$ and $H_{src}.\text{po} = H'_{src}.\text{po}$.
- If $R^{\text{Opq}}, i_1 \in H_{src}.\text{E}$ and $R^{\text{Opq}} \xrightarrow{\text{po}} i_1$, then $R \xrightarrow{\text{ctrl}} i_1$ in H_{tgt} because $H_{src}.\text{E} = H'_{src}.\text{E}$ and $H_{src}.\text{po} = H'_{src}.\text{po}$.
- If $i_1, i_2 \in H_{src}.\text{E}$ and $i_1 \xrightarrow{\text{push}} i_2$, then $i_1 \xrightarrow{\text{sync}} i_2$ for $i_1, i_2 \in H_{tgt}.\text{E}$ because $\forall i \in H_{src}.\text{E}, H_{src}.\text{AccessMode}(i) = H'_{src}.\text{AccessMode}(i)$ and $H_{src}.\text{po} = H'_{src}.\text{po}$, which means $H_{src}.\text{push} = H'_{src}.\text{push}$.
- If $i_1, i_2 \in H_{src}.\text{E}$ and $i_1 \xrightarrow{\text{ra}} i_2$, then $i_1 \xrightarrow{\text{lwsync}} i_2$ for $i_1, i_2 \in H_{tgt}.\text{E}$ because $\forall i \in H_{src}.\text{E}, H_{src}.\text{AccessMode}(i) = H'_{src}.\text{AccessMode}(i)$ and $H_{src}.\text{po} = H'_{src}.\text{po}$, which means $H_{src}.\text{ra} = H'_{src}.\text{ra}$.

Therefore, we have shown that for all $H_{tgt} \in \text{Histories}_{Power}(P)$, there exists an $H_{src} \in \text{Histories}_{JAM}(P)$ such that $H_{src} \rightsquigarrow H_{tgt}$. That is, the compilation scheme shown in Fig. 4 is correct. \blacktriangleleft

► **COROLLARY 1 (COMPILATION CORRECTNESS TO POWER (TRAILING FENCE CONVENTION))**. The compilation from Java to Power following the compilation scheme in Fig. 4 (using the trailing fence convention) is correct. That is, let P_{src} be a Java program, P_{tgt} be the Power program compiled from P_{src} using the compilation scheme in Fig. 4 (using the trailing fence convention). For all $H_{tgt} \in \text{Histories}_{Power}(P_{tgt})$ there exists a $H_{src} \in \text{Histories}_{JAM}(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

Proof. It is obvious that all the properties described in Definition 4 still hold with the trailing fence convention. Most importantly, the property that transforms `push` in the source level execution to `sync` in the target level execution is preserved as long as there is a `hwsync` instruction inserted between every `Volatile` accesses. The trailing fence convention, if used consistently, clearly satisfy this property. Then the rest of Definition 4 is unchanged since leading/trailing fence convention only concerns the compilation schemes for `Volatile` accesses. Therefore, the rest of the proof for the correctness of the trailing fence convention can be naturally derived similarly. ◀

E The x86 TSO Model in Herd7

```
X86 TSO
include "x86fences.cat"
include "filters.cat"
include "cos.cat"

(* Uniproc check *)
let com = rf | fr | co
acyclic po-loc | com

(* Atomic *)
empty rmw & (fre;coe)

(* GHB *)
#ppo
let po_ghb = WW(po) | RM(po)

#mfence
let mfence = try fencerel(MFENCE) with 0
let lfence = try fencerel(LFENCE) with 0
let sfence = try fencerel(SFENCE) with 0

show data,addr,ctrl

#implied barriers
let poWR = WR(po)
let i1 = MA(poWR)
let i2 = AM(poWR)
let implied = i1 | i2

let ghb = mfence | implied | po_ghb | rfe | fr | co
show implied
acyclic ghb as tso
```

F Compilation to x86-TSO

In this section we show that the current compilation scheme to x86-TSO is correct with respect to the TSO memory model.

F.0.1 The x86-TSO Model

We use the x86-TSO model defined in Herd7 [1], and the full model can be found in Appendix E):

► DEFINITION 11. An execution history H is **TSO-consistent** if it is trace coherent and satisfies the following three requirements:

1. **SC-PER-LOCATION**: `po-loc | com` is acyclic
2. **ATOMICITY**: `rmw & (fre ; coe)` is empty
3. **GLOBAL HAPPENS-BEFORE**: `ghb` is acyclic

We say such execution history H is **allowed** by TSO. Otherwise, it is **forbidden**.

F.0.2 Compilation Scheme

We use the following compilation scheme⁷:

```

    getOpaque() ~> mov
    setOpaque() ~> mov
    getAcquire() ~> mov
    setRelease() ~> mov
    getVolatile() ~> mov
    setVolatile() ~> mov ; mfence
    AcquireFence() ~> NoOp
    ReleaseFence() ~> NoOp
    fullFence() ~> mfence
    getAndAdd() ~> lock xaddl
    getAndAddAcquire() ~> lock xaddl
    getAndAddRelease() ~> lock xaddl
  
```

F.0.3 Proof of Compilation Correctness

► DEFINITION 12 (COMPILATION OF AN EXECUTION). We define the "CompilesTo" relation $\rightsquigarrow \subseteq \mathbb{H} \times \mathbb{H}$ for the compilation from Java to x86 as the followings: Given a Java program P_{src} and a memory model J that supports Java, let P_{tgt} be the target-level program compiled from P_{src} using the compilation scheme to x86 as shown above. Let H_{src} be a candidate execution history of P_{src} and H_{tgt} be a candidate execution history of P_{tgt} . We say $H_{src} \rightsquigarrow H_{tgt}$ if:

1. H_{tgt} is observationally equivalent to H_{src}
2. $H_{tgt}.co \subseteq H_{src}.to$
3. If $i_1, i_2 \in H_{src}.E$ and $i_1 \xrightarrow{\text{push}} i_2$ and i_1 is a write, then $i_1 \xrightarrow{\text{po}} i_3 \xrightarrow{\text{po}} i_2$ for $i_1, i_2 \in H_{tgt}.E$ and $i_3 \in H_{tgt}.F$ where i_3 is an event stem from an `mfence` instruction.

⁷ Note that instead of using `mfence` instruction for full fences, HotSpot uses a read-modify-write instruction to emulate the synchronization effect of it. According to the definition of TSO, the synchronization effect of a RMW event is exactly the same as an `mfence` event. Both of them produce a `ghb` order before and after the event. Therefore, we keep the simplicity of the proof here by using the `mfence` instruction.

Note that this definition does not say anything about whether an execution graph is consistent under a memory model.

► **LEMMA 6** (JAM'_{21} TO x86-TSO). Let P_{src} be a Java program, P_{tgt} be the x86 program compiled from P_{src} using the compilation scheme to x86 as shown above. For all $H_{tgt} \in \text{Histories}_{TSO}(P_{tgt})$ there exists a $H_{src} \in \text{Histories}_{JAM'}(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

Proof. It is obvious that there exists an H_{src} such that $H_{src} \rightsquigarrow H_{tgt}$. We show that $H_{src} \in \text{Histories}_{JAM'}(P)$. That is, we show that H_{src} is consistent under JAM'_{21} by showing that it fulfills the two requirements of JAM'_{21} .

1. **NO-THIN-AIR** Requirement. ($\text{po}|\text{rf}$) is acyclic. The rfi order is included in the po . Therefore, we need to show that ($\text{po}|\text{rfe}$) is acyclic. Note that since the only inter-thread order is rfe . If there is a cycle in ($\text{po}|\text{rfe}$), then the head of each thread is a R^{Opq} and the last event in each thread participating in this cycle is a W^{Opq} , where $R^{\text{Opq}} \xrightarrow{\text{po}} W^{\text{Opq}}$ in H_{src} . Using a compiler that follows the compilation scheme, this translates to $R \xrightarrow{\text{po}} W$ in H_{tgt} . Further we can infer that $R \xrightarrow{\text{ghb}} W$ in H_{tgt} . x86-TSO ensures that ghb is acyclic. Therefore, if there is a cycle of ($\text{po}|\text{rf}$) in H_{src} , H_{tgt} would not be consistent under x86-TSO's memory model, contradicting to our previous assumption.
2. **COHERENCE** Requirement. In order to show that H_{src} fulfills the coherence requirement, we need to show that co in H_{tgt} is a linear extension of co-jom in H_{src} . We prove this by analyzing each case for co-jom . That is, if $i_1 \xrightarrow{\text{co-jom}} i_2$ but $i_2 \xrightarrow{\text{co}} i_1$, then H_{tgt} is inconsistent under x86-TSO.

- **coinit.** This follows naturally as $H_{src}.\text{IW} = H_{tgt}.\text{IW}$.
- **cofw.** This follows naturally as $H_{src}.\text{FW} = H_{tgt}.\text{FW}$.
- **corr.** This implies that there exists R_1 and R_2 such that $i_1 \xrightarrow{\text{rf}} R_1$, $i_2 \xrightarrow{\text{rf}} R_2$, and $R_1 \xrightarrow{\text{po}} R_2$. From $i_2 \xrightarrow{\text{co}} i_1$ we can infer that $R_2 \xrightarrow{\text{fr}} i_1$. Note that $R_1 \xrightarrow{\text{ghb}} R_2$ in this case since a po order from a read event is preserved in TSO. Now we have a ghb cycle $R_2 \xrightarrow{\text{fr}} i_1 \xrightarrow{\text{rf}} R_1 \xrightarrow{\text{ghb}} R_2$, contradicting to our previous assumption that H_{tgt} is consistent under x86-TSO.
- **covw.** This implies that $i_1 \xrightarrow{\text{vo}} i_2$. We analyze each case of vo order.
 - $i_1 \xrightarrow{\text{po-loc}} i_2$. $i_2 \xrightarrow{\text{co}} i_1$ would violate the SC-Per-Location Requirement of Power, making H_{tgt} inconsistent, contradicting to our previous assumption.
 - $i_1 \xrightarrow{\text{ra}} i_2$ or $i_1 \xrightarrow{\text{push}} i_2$ or $i_1 \xrightarrow{\text{svo}} i_2$. Note that all three cases are included in $i_1 \xrightarrow{\text{po}} i_2$ to the same location. Thus, if $i_2 \xrightarrow{\text{co}} i_1$, then there would be a cycle of ($\text{po-loc}|\text{com}$) in H_{tgt} , contradicting to our previous assumption that H_{tgt} is consistent under TSO.
 - $i_1 \xrightarrow{\text{push}} E_1 \xrightarrow{\text{chapo}} E_2 \xrightarrow{\text{push}} i_2$ for some E_1 and E_2 in H_{src} . Since i_1 is a write, we can infer that $i_1 \xrightarrow{\text{po}} i_{\text{LOCK}} \xrightarrow{\text{po}} E_1$ in H_{tgt} , where i_{LOCK} is a RMW event. According to x86-TSO, we can further infer that $i_1 \xrightarrow{\text{ghb}} i_{\text{LOCK}} \xrightarrow{\text{ghb}} E_1$ in H_{tgt} , which can be simplified to $i_1 \xrightarrow{\text{ghb}} E_1$. Similarly, we can infer that $E_2 \xrightarrow{\text{ghb}} i_2$ (if E_2 is a read then the po order is included in ghb ; otherwise, there must be a RMW event between E_2 and i_2 , which yields a ghb order too). Now, since the communication edges induced by chapo are also included in ghb in H_{tgt} , $i_2 \xrightarrow{\text{co}} i_1$ would directly produce a ghb cycle, contradicting with our previous assumptions.
 - $i_1 \xrightarrow{\text{vvo}} E \xrightarrow{\text{vvo}} i_2$ for some E in H_{src} . Note that vo orders in H_{src} only produce ghb orders in H_{tgt} . Therefore, $i_2 \xrightarrow{\text{co}} i_1$ would always result in a ghb cycle in H_{tgt} , contradicting to the previous assumption.

- **cowr**. This implies that $i_1 \xrightarrow{\text{vo}} R$ and $i_2 \xrightarrow{\text{rf}} R$ for some R in H_{src} . With $i_2 \xrightarrow{\text{co}} i_1$ we can infer that $R \xrightarrow{\text{fr}} i_1$. As in previous cases, we observe that **vo** only produce **ghb** in H_{tgt} . Therefore, having $i_2 \xrightarrow{\text{co}} i_1$ would result in a **ghb** cycle in H_{tgt} .
- **corw**. This implies that there is R in H_{src} such that $R \xrightarrow{\text{po}} i_2$ and $i_1 \xrightarrow{\text{vo}} R$. Note that **po** order from a Read access is included in **ghb** in H_{tgt} . As in previous cases, we observe that **vo** only produce **ghb** in H_{tgt} . Therefore, having $i_2 \xrightarrow{\text{co}} i_1$ would result in a **ghb** cycle, contradicting to our previous assumption.
- **cormwexcl**. This implies that $W \xrightarrow{\text{rf}} R$, $R \xrightarrow{\text{rmw}} i_1$, and $W \xrightarrow{\text{co}} i_2$ for some W and R in H_{src} . We also have $R \xrightarrow{\text{fr}} i_2$. Now having $i_2 \xrightarrow{\text{co}} i_1$ would violate the Atomicity Requirement of TSO, contradicting to our previous assumptions.
- **cormwtotal**. Since there is no other restrictions on the total order among **RMW** operations except that it has to be compatible with the **rf** and intra-thread orders, it automatically becomes a subset of **co** in H_{tgt} . Therefore, having $i_2 \xrightarrow{\text{co}} i_1$ in this case would yield a cycle in **co**, producing a **ghb** cycle in H_{tgt} .

Thus we have shown that **co** is a linear extension of **co-jom**. As a consequence, **co-jom** is guaranteed to be acyclic is H_{tgt} is consistent under TSO.

Thus H_{src} is JAM'_{21} -consistent. That is, $H_{src} \in \text{Histories}_{JAM'}(P)$. ◀

► **THEOREM 13 (COMPILATION CORRECTNESS TO x86-TSO)**. Let P_{src} be a Java program, P_{tgt} be the x86 program compiled from P_{src} using the compilation scheme to x86 as shown above. For all $H_{tgt} \in \text{Histories}_{TSO}(P_{tgt})$ there exists a $H_{src} \in \text{Histories}_{JAM}(P_{src})$ such that $H_{src} \rightsquigarrow H_{tgt}$.

Proof. By Lemma 6, we know that there exists an $H'_{src} \in \text{Histories}_{JAM'}(P_{src})$ such that $H'_{src} \rightsquigarrow H_{tgt}$. Therefore, by definition of the \rightsquigarrow relation,

1. H_{tgt} is observationally equivalent to H'_{src}
2. $H_{tgt}.\text{co} \subseteq H'_{src}.\text{to}$
3. If $i_1, i_2 \in H'_{src}.\text{E}$ and $i_1 \xrightarrow{\text{push}} i_2$ and i_1 is a write, then $i_1 \xrightarrow{\text{po}} i_3 \xrightarrow{\text{po}} i_2$ for $i_1, i_2 \in H_{tgt}.\text{E}$ and $i_3 \in H_{tgt}.\text{F}$ where i_3 is an event stem from an **mfence** instruction.

By Theorem 12, we know that for all H'_{src} , there exists an $H_{src} \in \text{Histories}_{JAM}(P)$ such that H_{src} is observationally equivalent to H'_{src} . By Lemma 2, H_{src} is observationally equivalent to H_{tgt} . Furthermore,

1. $H_{tgt}.\text{co} \subseteq H_{src}.\text{to}$ because $H_{src}.\text{to} = H'_{src}.\text{to}$.
2. If $i_1, i_2 \in H_{src}.\text{E}$ and $i_1 \xrightarrow{\text{push}} i_2$ and i_1 is a write, then $i_1 \xrightarrow{\text{po}} i_3 \xrightarrow{\text{po}} i_2$ for $i_1, i_2 \in H_{tgt}.\text{E}$ and $i_3 \in H_{tgt}.\text{F}$ where i_3 is an event stem from an **mfence** instruction, because $\forall i \in H_{src}.\text{E}, H_{src}.\text{AccessMode}(i) = H'_{src}.\text{AccessMode}(i)$ and $H_{src}.\text{po} = H'_{src}.\text{po}$, which means $H_{src}.\text{push} = H'_{src}.\text{push}$.

Therefore, we have shown that for all $H_{tgt} \in \text{Histories}_{TSO}(P)$, there exists an $H_{src} \in \text{Histories}_{JAM}(P)$ such that $H_{src} \rightsquigarrow H_{tgt}$. That is, the compilation scheme to x86 is correct. ◀

G

 Program Transformations

G.1 Deordering and Reordering

► **THEOREM 4 (DEORDERING)**. Let P_{src} be a Java program and P_{tgt} be a Java program obtained by performing a deordering operation on a pair of accesses a and b according to Fig. 6. Let H_{tgt} be an execution of P_{tgt} . Then there exists an execution H_{src} of P_{src} such that

- $H_{src}.po = H_{tgt}.po \cup \{(a, b)\}$ where a and b are po-adjacent
- $H_{src}.rf = H_{tgt}.rf$
- $H_{src}.E = H_{tgt}.E$
- $H_{src}.to = H_{tgt}.to$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{src}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$

and if H_{tgt} is JAM_{21} -consistent, then H_{src} is JAM_{21} -consistent.

Proof. Note that for the **COHERENCE** requirement, only three kinds of edges contributes to **co-jam**: **vo**, **rf**, and **po** to the same location. Since here we are considering deorderable pairs, which are pairs of accesses to different locations related by **po** in H_{src} , we only need to consider whether deordering them would affect the set of **vo** in the execution. We can analyze this case by case.

- $(R_x * R_y)$. The **NO-THIN-AIR** requirement is fulfilled automatically in H_{src} since we are deordering a pair of reads. Since $po \cap (R_x^{\text{Opaque}} * R_y) \not\subseteq vo$, it follows that H_{src} is also JAM_{21} -consistent.
- $(R_x * W_y)$. H_{src} fulfills **COHERENCE** since the **po** edge between two accesses whose access mode is weaker or equal to **Opaque** mode does not contribute to any new **vo** edge. In addition, H_{src} fulfills the **NO-THIN-AIR** requirement because one of the accesses is in **Plain** mode where as **NO-THIN-AIR** only requires the acyclicity of $po \cup rf$ among **Opaque** mode accesses.
- $(R_x * RMW_y)$. First note that the **Volatile** mode for **RMWs** include the effect of **Release Mode**. H_{src} fulfills **COHERENCE** since the **po** edge between R_x and RMW_y does not contribute to any new **vo** edge. H_{src} fulfills **NO-THIN-AIR** because R_x is **Plain** mode.
- $(R_x * F)$. It's easy to see that the **po** edge added in H_{src} does not contribute to any new **vo** edge therefore the **COHERENCE** is fulfilled. Since we are deordering a read and a fence, H_{src} fulfills **NO-THIN-AIR** automatically.
- $(W_x * R_y)$. It's easy to see that the **po** edge added in H_{src} does not contribute to any new **vo** edge therefore the **COHERENCE** is fulfilled. Since we are deordering a write and a read, H_{src} fulfills **NO-THIN-AIR** automatically. The only situation the two accesses cannot be deordered is when they are both in **Volatile** mode because the **po** between two **Volatile** accesses can be derived into a **vo** order.
- $(W_x * W_y)$. It's easy to see that the **po** edge added in H_{src} does not contribute to any new **vo** edge therefore the **COHERENCE** is fulfilled. Since we are deordering a write and a write, H_{src} fulfills **NO-THIN-AIR** automatically. Here we need the second write W_y to be weaker than **release mode** to ensure that the **po** between the two accesses does not contribute to the **vo** order.
- $(W_x * RMW_y)$. It's easy to see that the **po** edge added in H_{src} does not contribute to any new **vo** edge therefore the **COHERENCE** is fulfilled. Since we are deordering a write and a read-modify-write, H_{src} fulfills **NO-THIN-AIR** automatically. Since RMW_y is both in the

set of reads and in the set of writes of the execution graph H_{src} , we take the intersection of previous cases. In addition, **rel** and **acq** do not subsume each other, so it is safe for o_2 to be **acq** mode.

- $(W_x * F)$. It's easy to see that the **po** edge added in H_{src} does not contribute to any new **vo** edge therefore the **COHERENCE** is fulfilled. Since we are deordering a write and a fence, H_{src} fulfills **NO-THIN-AIR** automatically. Here we are basically avoiding the situation when W_x and F can form any **svo**. In addition, there are also situations where W_x and F have a **svo** if the writes that follows the fence are already in **rel** modes. Since **svo** and **ra** are considered equivalently in terms of their memory order effect in the JAM_{21} model, the **svo** they form is redundant in the presence of all the **ra**.
- $(RMW_x * R_y)$. It's easy to see that the **po** edge added in H_{src} does not contribute to any new **vo** edge therefore the **COHERENCE** is fulfilled. Since we are deordering a read-modify-write and a read, H_{src} fulfills **NO-THIN-AIR** automatically. Here we want to avoid the RMW_x to have an access mode stronger or equal to **acq** mode because it'd create an **ra** edge which is considered as a **vo** edge.
- $(RMW_x * W_y)$. It's easy to see that the **po** edge added in H_{src} does not contribute to any new **vo** edge therefore the **COHERENCE** is fulfilled. For the **NO-THIN-AIR** requirement, since W_y is in **Plain** mode, it does not contribute to any **po** \cup **rf** cycle among **Opaque** accesses.
- $(RMW_x * RMW_y)$. They cannot be deordered due to the **NO-THIN-AIR** requirement. (Note that RMW operations are atomic by definition, so there is no **Plain** mode or **Opaque** mode for RMW).
- $(RMW_x * F)$. Similar to the previous cases.
- Deordering with fence. Similar to the previous cases.

◀

► **COROLLARY 2 (REORDERING)**. JAM_{21} supports the reordering transformation for pairs of adjacent accesses shown in Fig. 6.

Proof. Let a and b be a pair of such memory events and $a \xrightarrow{-po} b$ in H_{src} . By Theorem 4, we know that removing the **po** edge between a and b does not introduce new program behavior. Let H' be the execution graph after the deordering transformation. By Theorem 3, we know that adding a **po** edge from b to a in H' does not introduce new program behavior either. Therefore, reordering of access pairs in Fig. 6 is supported by JAM_{21} . ◀

G.2 Merging

G.2.1 Read-read Merging

► **THEOREM 5 (READ-READ MERGING)**. Let H_{tgt} be an JAM_{21} -consistent execution. Let $a \in H_{tgt}.R \setminus RMW$ and let $a' \in H_{tgt}.E$ such that $a \xrightarrow{rf} a'$. Let $b \notin H_{tgt}.E$. There exists a H_{src} such that:

- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, b \rangle \mid i \xrightarrow{-po} a\} \cup \{\langle b, j \rangle \mid a \xrightarrow{-po} j\}$
- $H_{src}.rf = H_{tgt}.rf \cup \{\langle a', b \rangle\}$
- $H_{src}.E = H_{tgt}.E \cup \{b\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{to} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{to} j\}$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $b \in H_{src}.R$

- $H_{src}.AccessMode(b) = H_{src}.AccessMode(a) \sqsubseteq \text{Acquire}$

and H_{src} is JAM_{21} -consistent.

Proof. We show that H_{src} fulfills the two requirements needed to be JAM_{21} -consistent.

- Suppose H_{src} violates the **NO-THIN-AIR** requirement, then there is a $(\text{po} \cup \text{rf})^+$ cycle involving b . If we have $a' \xrightarrow{\text{rf}} b \xrightarrow{-(\text{po}|\text{rf})^+} a'$, then $a' \xrightarrow{\text{rf}} a \xrightarrow{-(\text{po}|\text{rf})^+} a'$. If we have $a \xrightarrow{\text{po}} b \xrightarrow{-(\text{po}|\text{rf})^+} a$, then $a \xrightarrow{-(\text{po}|\text{rf})^+} a$. In both of the cases, H_{tgt} is inconsistent, which contradicts with our previous assumption. Therefore, the **NO-THIN-AIR** requirement is fulfilled by H_{src} .
- Suppose H_{src} violates the **COHERENCE** requirement, then there is a **co** cycle. Note that $AccessMode(b) = AccessMode(a) \sqsubseteq \text{Acq}$. In addition, for all events i , if $b \xrightarrow{\text{vo}} i$, then $a \xrightarrow{\text{vo}} i$ and for all events j , if $j \xrightarrow{\text{vo}} b$, then $j \xrightarrow{\text{vo}} a$. Therefore, for any coherence cycle derived from the edges from and to b , there is also a coherence cycle derived from the edges from and to a . If H_{src} has a **co** cycle, H_{tgt} also has a **co** cycle, which contradicts with our previous assumption. ◀

G.2.1.1 Counter Example

Here, we give an example showing that read-read merging is not allowed by JAM_{21} if the read accesses are both Volatile mode. Consider the following program:

```

Thread0 {
  int r1 = X.getOpaque(); // 1
  int r2 = X.getOpaque(); // 2
}

Thread1 {
  int r3 = Y.getOpaque(); // 1
  int r4 = Y.getOpaque(); // 2
}

Thread2 {
  X.setOpaque(2);
}

Thread3 {
  int r5 = X.getVolatile(); // 2
  int r6 = X.getVolatile(); // 2
  Y.setRelease(1);
}

Thread4 {
  Y.setVolatile(2);
  X.setVolatile(1);
}

```

Applying the read-read merging transformation to this program yields:

```

Thread0 {
  int r1 = X.getOpaque(); // 1
  int r2 = X.getOpaque(); // 2
}

Thread1 {
  int r3 = Y.getOpaque(); // 1
  int r4 = Y.getOpaque(); // 2
}

Thread2 {
  X.setOpaque(2);
}


Thread3 {
  int r5 = X.getVolatile(); // 2
  int r6 = r5
  Y.setRelease(1);
}

Thread4 {
  Y.setVolatile(2);
  X.setVolatile(1);
}


```

The execution graphs with the annotated read values is shown in Fig. 11 and Fig. 12.

For the two read accesses of x on Thread 3, one may think it's OK to merge them into one. However, since they are Volatile accesses, they also impose a **push** edge which is totally



■ **Figure 11** Execution Graph before read-read merge on Volatile (Forbidden)



■ **Figure 12** Execution Graph after read-read merge on Volatile (Allowed)

ordered with other **push** edges. Merging the two reads removes the synchronization provided by the **push** edge, introducing the program behavior shown in Fig. 12.

G.2.2 Write-write Merging


► **THEOREM 6 (WRITE-WRITE MERGING)**. Let H_{tgt} be an JAM_{21} -consistent execution. Let $b \in H_{tgt}.W \setminus RMW$ and let $a \notin H_{tgt}.E$ and $loc(a) = loc(b) \wedge \forall i \in H_{tgt}.W, loc(i) = loc(b) \Rightarrow val(a) \neq val(i)$. There exists a H_{src} such that:

- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{po} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{po} j\}$
- $H_{src}.rf = H_{tgt}.rf$
- $H_{src}.E = H_{tgt}.E \cup \{a\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{to} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{to} j\}$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $a \in H_{src}.W$
- $H_{src}.AccessMode(a) = H_{src}.AccessMode(b) \sqsubseteq \text{Release}$

and H_{src} is JAM_{21} -consistent.

Proof. We show that H_{src} fulfills the two requirements to be JAM_{21} -consistent.

- **NO-THIN-AIR.** Note that $a \xrightarrow{po} b \xrightarrow{-(po|rf)+} a$ implies that $b \xrightarrow{-(po|rf)+} b$. Therefore, if H_{src} violates **NO-THIN-AIR**, H_{tgt} also violates **NO-THIN-AIR**, contradicting to our previous assumption.



■ **Figure 13** Execution graph before write-write merge on Volatile (Forbidden)

- **COHERENCE.** First note that there is no extra **rf** edge from a and $\forall i, (i \xrightarrow{\text{vo}} a \Rightarrow i \xrightarrow{\text{vo}} b) \wedge (a \xrightarrow{\text{vo}} i \Rightarrow b \xrightarrow{\text{vo}} i)$ (because a and b have the same access mode and they are not in Volatile mode). Therefore, any **co** cycle derived from a , we can derive the same **co** cycle with b . While $a \xrightarrow{\text{po}} b$ implies that $a \xrightarrow{\text{vo}} b$, since there is no **rf** edge from a , it cannot contribute to any extra **co** cycle. Therefore, if there is a **co** cycle in H_{src} , then it implies that there is a **co** cycle in H_{tgt} , contradicting to our previous assumption.



G.2.2.1 Counter Example

We now provide a counter-example showing write-write merge is not valid for Volatile mode writes. Consider the following example program:

```

Thread0 {
    int r1 = X.getOpaque(); // 2
    int r2 = X.getOpaque(); // 3
}

Thread1 {
    int r3 = Y.getOpaque(); // 1
    int r4 = Y.getOpaque(); // 2
}

Thread2 {
    Y.setOpaque(2);
    X.setVolatile(1);
    X.setVolatile(2);
}

Thread3 {
    X.setVolatile(3);
    Y.setVolatile(1);
}
    
```

The execution graph of the program before the transformation is shown in Fig. 13. Applying write-write merging transformation to Thread 2, we have:

```


Thread0 {
    int r1 = X.getOpaque(); // 2
    int r2 = X.getOpaque(); // 3
}

Thread1 {
    int r3 = Y.getOpaque(); // 1
    int r4 = Y.getOpaque(); // 2
}

Thread2 {
    Y.setOpaque(2);
    X.setVolatile(2);
}

Thread3 {
    X.setVolatile(3);
    Y.setVolatile(1);
}
    
```

The execution graph after the transformation is shown in Fig. 14. After removing the write access in Volatile mode, the cross-thread synchronization effect between Thread 2 and Thread 3 is also removed, introducing the new behavior in the figure.



■ **Figure 14** Execution graph after write-write merge on Volatile (Allowed)

G.2.3 Write/RMW-read Merging

► **THEOREM 7 (WRITE/RMW-READ MERGING).** Let H_{tgt} be a JAM_{21} -consistent execution. Let $a \in H_{tgt}.W$ and $b \notin H_{tgt}.E$. There exists a H_{src} such that:

- $H_{src}.E = H_{tgt}.E \cup \{b\}$
- $b \in H_{src}.R$
- $H_{src}.loc(b) = H_{src}.loc(a)$
- $H_{src}.val(b) = H_{src}.val(a)$
- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{po} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{po} j\}$
- $H_{src}.rf = H_{tgt}.rf \cup \{\langle a, b \rangle\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{to} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{to} j\}$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $H_{src}.AccessMode(b) \sqsubseteq \text{Opaque}$

Proof. We show that H_{src} fulfills the two requirements to be JAM_{21} -consistent.

- **NO-THIN-AIR.** First note that, by the well-formedness of **rf** order, a is the only access in the execution graph that has a **rf** edge to b . Therefore, $a \xrightarrow{rf} b \xrightarrow{-(po|rf)+} a$ implies that $a \xrightarrow{-(po|rf)+} a$, which means there is also a $(po|rf)+$ cycle in H_{tgt} , contradicting to our previous assumption.
- **COHERENCE.** Since $AccessMode(b) = \text{Opaque}$, there is no out-going cross-thread edge from b and for all event i such that $b \xrightarrow{vo} i$, we have $a \xrightarrow{vo} i$ (similarly, for all event j such that $j \xrightarrow{vo} b$, we have $j \xrightarrow{vo} a$). Since $a \xrightarrow{rf} b$ is intra-thread, for any **co** edge derived from $a \xrightarrow{rf} b$ using the **corr** rule, it implies that there is a read access R and write access W such that $a \xrightarrow{rf} b \xrightarrow{po} R$ and $W \xrightarrow{rf} R$ we can derive the same **co** edge using the **cowr** rule with $a \xrightarrow{vo} R$ and $W \xrightarrow{rf} R$. Similarly, for any **co** edge derived from $a \xrightarrow{rf} b$ using the **coww** rule, it implies that there is a write access W such that $W \xrightarrow{vo} b$. Then $W \xrightarrow{vo} a$ as well. Using the **coww** rule we can derive the same **co** edge. Thus, if there is any **co** cycle in H_{src} , the same **co** cycle also appear in H_{tgt} , contradicting to our previous assumption.

◀

G.2.3.1 Counter Example

Here we show that write/RMW-read merging is not valid if the read is or is stronger than Acquire mode. Consider the following example:

```

Thread0 {
  int r1 = X.getOpaque(); // 1
  int r2 = X.getOpaque(); // 2
}

Thread1 {
  int r3 = Y.getOpaque(); // 1
  int r4 = Y.getOpaque(); // 2
}

Thread2 {
  Y.setOpaque(1);
}

Thread3 {
  X.setRelease(2);
  int r7 = X.getAcquire(); // 2
  int r5 = Z.getVolatile(); // 0
  int r6 = Y.getVolatile(); // 1
}

Thread4 {
  Y.setVolatile(2);
  X.setVolatile(1);
}

```

The execution graph can be found in Fig. 15. The execution is forbidden. Indeed, there are two possible cases:

1. $R_z^V = 0 \xrightarrow{\text{vvo}} W_x^V = 1$. Since $\text{rf} \subseteq \text{vvo}$ and $\text{ra} \subseteq \text{vvo}$, we can infer that $W_x^{\text{rel}} = 2 \xrightarrow{\text{vvo}} R_z^V = 0 \xrightarrow{\text{vvo}} W_x^V = 1$. Using the coww rule, we can infer that $W_x^{\text{rel}} = 2 \xrightarrow{\text{co}} W_x^V = 1$, which contradicts with the co edge we inferred using the corr rule and Thread 0.
2. $W_y^V = 2 \xrightarrow{\text{vvo}} R_y^V = 1$. This immediately contradicts with the co edge we derived using the corr rule with Thread 1.

Applying the transformation, we have:

```

Thread0 {
  int r1 = X.getOpaque(); // 1
  int r2 = X.getOpaque(); // 2
}

Thread1 {
  int r3 = Y.getOpaque(); // 1
  int r4 = Y.getOpaque(); // 2
}

Thread2 {
  Y.setOpaque(1);
}

Thread3 {
  X.setRelease(2);
  int r7 = 2;
  int r5 = Z.getVolatile(); // 0
  int r6 = Y.getVolatile(); // 1
}

Thread4 {
  Y.setVolatile(2);
  X.setVolatile(1);
}


```

The execution graph is shown in Fig. 16. Due to the removal of the rf and ra edge, the previously forbidden behavior is introduced after the transformation.

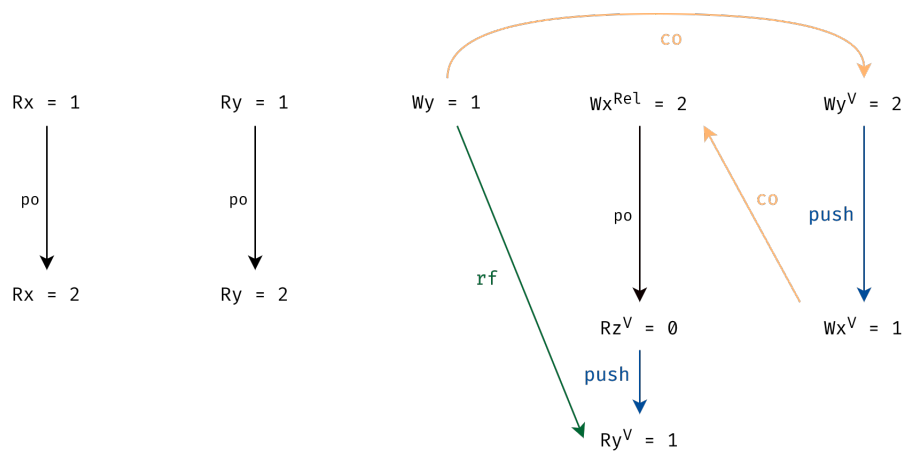
G.2.4 Write-RMW Merging

► **THEOREM 8 (WRITE-RMW MERGING).** Let H_{tgt} be a JAM_{21} -consistent execution. Let $b \in H_{tgt}.W \setminus H_{tgt}.RMW$, $a \notin H_{tgt}.E$ and $v \in \text{Val}$. There exists a H_{src} such that:

- $H_{src}.E = H_{tgt}.E \cup \{a\}$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $H_{src}.AccessMode(a) \in \{\text{Opaque}, \text{Release}\}$
- $H_{src}.AccessMode(b) \in \{\text{Acquire}, \text{Release}\}$
- $H_{src}.loc(b) = H_{src}.loc(a)$
- $b \in H_{src}.RMW$
- $H_{src}.val(b) = (H_{src}.val(a), v)$
- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{\text{po}} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{\text{po}} j\}$
- $H_{src}.rf = H_{tgt}.rf \cup \{\langle a, b \rangle\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, a \rangle \mid i \xrightarrow{\text{to}} b\} \cup \{\langle a, j \rangle \mid b \xrightarrow{\text{to}} j\}$



■ **Figure 15** Execution Graph before Write-read Merge Transformation (Forbidden)



■ **Figure 16** Execution Graph after Write-read Merge Transformation (Allowed)

$$\cdot H_{src}.IW = H_{tgt}.IW$$

and H_{src} is JAM_{21} -consistent.

Proof. Most parts of the proof is similar to the proof for write-write merging except for the case where there is a **co** cycle in H_{src} due to the total coherence order among RMW operations. Suppose H_{src} violates **COHERENCE** by having a **co** cycle built from the **cormwtotal** rule. That is, we have a RMW operation i such that:

- If $b \xrightarrow{\text{cormwtotal}} i$, then $i \xrightarrow{\text{co}} b$
- If $i \xrightarrow{\text{cormwtotal}} b$, then $b \xrightarrow{\text{co}} i$

Note that we cannot use existing **co** orders to derive other orders than $\xrightarrow{\text{cormwexcl}}$ orders (which is also a **co** order). If the **co** between i and b are not $\xrightarrow{\text{cormwexcl}}$ edges, then they co-exists in one execution. Now we have $i \xrightarrow{\text{co}} b \xrightarrow{\text{co}} i$. If the **co** between i and b are $\xrightarrow{\text{cormwexcl}}$ edges, then there exist two RMW operations j and k , such that, $b \xrightarrow{\text{rf}} j$, $i \xrightarrow{\text{rf}} k$, $i \xrightarrow{\text{co}} j$ and $b \xrightarrow{\text{co}} k$. Note that there is still a total order among i, j, k in H_{tgt} . Now we have either $j \xrightarrow{\text{cormwtotal}} k$ or $k \xrightarrow{\text{cormwtotal}} j$. Each case yields a contradiction by the **coermwexcl** rule. Therefore, if there is a **co** cycle in H_{src} , H_{tgt} is also forbidden, which contradicts to our previous assumption. \blacktriangleleft

G.2.5 RMW-RMW Merging

► **THEOREM 9 (RMW-RMW MERGING).** Let H_{tgt} be a JAM_{21} -consistent execution. Let x be a memory location and $a \in H_{tgt}.E$ with $H_{tgt}.val(a) = (v_r, v_w)$, $H_{tgt}.loc(a) = x$, and $H_{tgt}.AccessMode(a) \in \{\text{Release, Acquire}\}$. Let $b \notin H_{tgt}.E$, there exists a H_{src} such that:

- $H_{src}.E = H_{tgt}.E \cup \{b\}$
- $\forall i \in H_{tgt}.E, H_{src}.AccessMode(i) = H_{tgt}.AccessMode(i)$
- $H_{src}.val(a) = (v_r, v)$
- $H_{src}.val(b) = (v, v_w)$
- $H_{src}.loc(b) = x$
- $H_{src}.AccessMode(b) = H_{src}.AccessMode(a) \in \{\text{Release, Acquire}\}$
- $H_{src}.po = H_{tgt}.po \cup \{\langle a, b \rangle\} \cup \{\langle i, b \rangle \mid i \xrightarrow{\text{po}} a\} \cup \{\langle b, j \rangle \mid a \xrightarrow{\text{po}} j\}$
- $H_{src}.rf = H_{tgt}.rf \cup \{\langle a, b \rangle\}$
- $H_{src}.to = H_{tgt}.to \cup \{\langle a, b \rangle\} \cup \{\langle i, b \rangle \mid i \xrightarrow{\text{to}} a\} \cup \{\langle b, j \rangle \mid a \xrightarrow{\text{to}} j\}$
- $H_{src}.IW = H_{tgt}.IW$

and H_{src} is JAM_{21} -consistent.

Proof. We show that H_{src} fulfills the two requirements of JAM_{21} -consistency.

- **NO-THIN-AIR.** Suppose H_{src} violates this requirement and has a **(po|rf)+** cycle. Since $H_{src}.val(a) = (v_r, v)$ and $H_{src}.val(b) = (v, v_w)$, if $a \xrightarrow{\text{po}} b \xrightarrow{\text{(po|rf)+}} a$ in H_{src} , it implies that $a \xrightarrow{\text{(po|rf)+}} a$ in H_{tgt} , contradicting to our previous assumption.
- **COHERENCE.** First note that there is only one **rf** edge from a in H_{src} and that is $a \xrightarrow{\text{rf}} b$. In addition, for all event i such that $i \xrightarrow{\text{vo}} b$ in H_{src} , $i \xrightarrow{\text{vo}} a$ in H_{tgt} . For all j such that $b \xrightarrow{\text{vo}} j$ in H_{src} , $a \xrightarrow{\text{vo}} j$ in H_{tgt} . Therefore, if there is a **co** cycle in H_{src} , there is also a **co** cycle in H_{tgt} , contradicting to our previous assumption. \blacktriangleleft

G.3 Register Promotion for non-shared Variable

► **THEOREM 10 (WEAKENING FOR NON-SHARED VARIABLE)**. Let H_{tgt} be a JAM_{21} -consistent execution such that, for all accesses i and j in $H_{tgt}.E$, $loc(i) = loc(j) = x \Rightarrow Tid(i) = Tid(j)$ for some memory location x . In addition, $\forall i \in H_{tgt}.E, loc(i) = x \Rightarrow AccessMode(i) = Plain$. There exists an execution H_{src} such that:

- $H_{src}.E = H_{tgt}.E$
- $H_{src}.po = H_{tgt}.po$
- $H_{src}.rf = H_{tgt}.rf$
- $H_{src}.to = H_{tgt}.to$
- $H_{src}.IW = H_{tgt}.IW$
- $\forall i \in H_{src}.E, loc(i) = x \Rightarrow AccessMode(i) \in \{Release, Acquire\}$

and H_{src} is JAM_{21} -consistent.

Proof. We show that H_{src} fulfills the two requirements of JAM_{21} -consistency.

1. **NO-THIN-AIR.** Note that there is no cross-thread **rf** edge from or to accesses of location x . Therefore, since $H_{src}.po = H_{tgt}.po$ and $H_{src}.rf = H_{tgt}.rf$, if there is a $(po|rf)^+$ cycle in H_{src} , there is a $(po|rf)^+$ cycle in H_{tgt} , contradicting to our previous assumption.
2. **COHERENCE.** Note that the transformation is equivalent to removing all the **ra** edges that involve accesses to x . Therefore, $H_{src}.vo = H_{tgt}.vo \setminus \{(a, b) \mid (loc(a) = x \wedge a \xrightarrow{ra} b \wedge AccessMode(b) \neq Release) \vee (loc(b) = x \wedge a \xrightarrow{ra} b \wedge AccessMode(a) \neq Acquire)\}$. For accesses i and j such that $loc(i) \neq x$ and $loc(j) \neq x$, if $i \xrightarrow{vo} j$ in H_{src} , $i \xrightarrow{vo} j$ in H_{tgt} . In addition, since x is not shared across different threads, all accesses to location x are related by **po**. Since all accesses to x have an access mode of either **Release** or **Acquire**, there is no cross-thread **vo** edges or **rf** edges from or to these accesses. Therefore, for all memory location $y \neq x$, $H_{src}.vo \upharpoonright_y = H_{tgt}.vo \upharpoonright_y$. Suppose H_{src} violates this requirement by having a **co** cycle:
 - If there is a **co** cycle with accesses to location x . Since $H_{src}.po-loc = H_{tgt}.po-loc$ and $po-loc \subseteq vo$, then there is also a **co** cycle with accesses to location x in H_{tgt} , contradicting to our previous assumption.
 - If there is a **co** cycle with accesses to other locations. Since for all memory location $y \neq x$, $H_{src}.vo \upharpoonright_y = H_{tgt}.vo \upharpoonright_y$ and $H_{src}.rf = H_{tgt}.rf$, it implies there is also a **co** cycle in H_{tgt} , contradicting to our previous assumption.

◀

► **THEOREM 11 (REMOVING PLAIN ACCESSES FOR NON-SHARED VARIABLE)**. Let H_{tgt} be a JAM_{21} -consistent execution. Let x be a memory location and for all $i \in H_{tgt}.E$ such that $loc(i) = x$, $Tid(i) = t$ for some t . Let $a \notin H_{tgt}.E$. There is a H_{src} such that:

- $H_{src}.E = H_{tgt}.E \cup \{a\}$
- $H_{src}.loc(a) = x$
- $H_{src}.AccessMode(a) = Plain$
- $H_{src}.po \supset H_{tgt}.po$
- for all $i \in H_{src}.E$ such that $H_{src}.loc(i) = x$, $i \xrightarrow{po} a$ or $a \xrightarrow{po} i$
- $H_{src}.rf = H_{tgt}.rf$ if $a \in H_{src}.W \setminus RMW$, otherwise, $H_{src}.rf = H_{tgt}.rf \cup \{(i, a)\}$ such that $(i \in H_{src}.W) \wedge (loc(i) = x) \wedge (i \xrightarrow{po} a) \wedge (\forall j \in H_{src}.E, (loc(j) = x) \wedge (j \xrightarrow{po} a) \Rightarrow (j \xrightarrow{po} i))$.
- $H_{src}.to = H_{tgt}.to$

23:52 Compiling Volatile Correctly in Java

$$\cdot H_{src}.IW = H_{tgt}.IW$$

and H_{src} is JAM_{21} -consistent.

Proof. It is clear that H_{src} does not violate **NO-THIN-AIR** and there is no **co** cycle for accesses to location x . For **COHERENCE**, note that for all memory location $y \neq x$, $H_{src}.vo \upharpoonright_y = H_{tgt}.vo \upharpoonright_y$ and $H_{src}.rf = H_{tgt}.rf$, it implies that if there is a **co** cycle in H_{src} there is also a **co** cycle in H_{tgt} , contradicting to our previous assumption. ◀

H Key Properties of the JAM_{21} Model

In this section, we show some key properties of JAM_{21} . First we show that the prior theorems of JAM_{19} still hold for JAM_{21} in Section H.1. Then in Section H.2, we prove that when all accesses in program order are **push** ordered then the semantics of executions is sequentially consistent. As a corollary when all accesses are **Volatile**, which implies a **push** order, then the executions are sequentially consistent. We have defined and proved these theorems in Coq. The Coq source code is included in our supplementary materials.

H.1 Prior Theorems

The JAM_{21} model preserves the two main theoretical results of [3], namely the monotonicity of access modes and the causal-acquire reads. We recount each theorem here briefly beginning with the monotonicity of access modes. In addition, we prove the DRF-SC theorem for JAM_{21} .

We use the reflexive ordering of the access modes as $\text{Plain} \sqsubseteq \text{Opaque} \sqsubseteq \text{ReleaseAcquire} \sqsubseteq \text{Volatile}$ and extend it to accesses $l_{m_1} \sqsubseteq l_{m_2}$, $l_{m_1} := n_1 \sqsubseteq l_{m_2} := n_2$, $\text{RMW}(l, n_1) \sqsubseteq \text{RMW}(l, n_2)$ whenever $m_1 \sqsubseteq m_2$. We treat read-modify-write (RMW) events as always having the same order. We extend the order to histories by matching identifiers and ordering the accesses.

$$H_1 \sqsubseteq H_2 \triangleq \forall i a_1 a_2, H_1(\text{is}(i, a_1)) \wedge H_2(\text{is}(i, a_2)) \Rightarrow a_1 \sqsubseteq a_2$$

We adopt the same notion of "well-formedness" from [3] for a given history H , i.e., *trace coherence*⁸.

► DEFINITION 13 (TRACE COHERENCE). An execution history H is trace coherent if:

- Each memory location is initialized by an initial write. For each event $i \in H.E \setminus \{H.IW \cup H.F\}$, there exists an initial write event $iw \in H.IW$ such that $H.loc(i) = H.loc(iw)$ and $iw \xrightarrow{\text{to}} i$.
- Reads-from edges are well-formed. For all $r \in H.R$, there exists a unique write w such that $H.loc(w) = H.loc(r)$, $H.val(w) = H.val(r)$, and $w \xrightarrow{\text{rf}} r$.
- There exists a total trace order to for all $e \in H.E$ such that to is compatible with po , rf , ra , svo , and push .

When the po , rf , and to relations of two histories H_1 and H_2 have the following relationships: $H_2.\text{po} \subseteq H_1.\text{po}$, $H_2.\text{to} \subseteq H_1.\text{to}$, $H_2.\text{rf} \subseteq H_1.\text{rf}$, then we say they *match*.

► THEOREM 14 (MONOTONICITY). [coq/Monotonicity.v, monotonicity]
For two histories H_1 and H_2 , suppose that both match, both are trace coherent, and $H_2 \sqsubseteq H_1$. Further suppose that $\text{acyclic}(\xrightarrow{\text{co}}_{H_1})$ and that there are no specified visibility orders or **push** orders in H_2 , then $\text{acyclic}(\xrightarrow{\text{co}}_{H_2})$

A version of DRF-SC theorem was proved in [3]. However, the theorem was different from the standard DRF-SC theorem.

- It did not use the conventional definition of data race with the "happens-before" order. Instead, [3] defined a **sync** order that captures the synchronizations between events and defined the notion of "data-race-free" using **sync**.

⁸ We have omitted some of the details of trace coherence that are related to the internals of the modeling language as they are irrelevant here

- It used a stronger assumption than the standard DRF-SC theorem. In particular, given a program P , the standard DRF-SC theorem assumes only the SC-consistent executions of P are data race free. On the other hand, the DRF-SC theorem proved in [3] assumes all executions of P are data race free. A similar theorem was also proved in [19], called a "model-agnostic" definition of DRF-SC.

Here, we first prove the standard DRF-SC theorem (DRF-SC) with a weaker assumption than [3], and then prove the "model-agnostic" DRF-SC theorem (EXECUTION-DRF), both using "happens-before" (**hb**).

We require the following standard definitions including the traditional notion of sequential consistency (SC-consistency) [16]:

$$\begin{aligned} i_1 \xrightarrow{\text{fr}} i_2 &\triangleq \exists i_3, i_3 \xrightarrow{\text{rf}} i_1 \wedge i_3 \xrightarrow{\text{co}} i_2 \\ i_1 \xrightarrow{\text{com}} i_2 &\triangleq i_1 \xrightarrow{\text{co}} i_2 \vee i_1 \xrightarrow{\text{rf}} i_2 \vee i_1 \xrightarrow{\text{fr}} i_2 \\ i_1 \xrightarrow{\text{sc}} i_2 &\triangleq i_1 \xrightarrow{\text{po}} i_2 \vee i_1 \xrightarrow{\text{com}} i_2 \end{aligned}$$

An execution H is SC-consistent if $\text{acyclic}(\xrightarrow{\text{sc}}_H)$.

We also require the notion of *happens-before* (**hb**) defined using the *synchronizes-with* (**sw**) order:

$$\begin{aligned} i_1 \xrightarrow{\text{sw}} i_2 &\triangleq (i_1 \xrightarrow{\text{rf}} i_2 \wedge \text{AccessMode}(i_1) = \text{Release} \wedge \text{AccessMode}(i_2) = \text{Acquire}) \\ &\quad \vee (\exists i_3 i_4 \in \mathbf{F}, \text{AccessMode}(i_3) = \text{Release} \wedge \text{AccessMode}(i_4) = \text{Acquire} \\ &\quad \wedge i_3 \xrightarrow{\text{po}} i_1 \xrightarrow{\text{rf}} i_2 \xrightarrow{\text{po}} i_4) \\ i_1 \xrightarrow{\text{hb}} i_2 &\triangleq i_1 \xrightarrow{\text{po}} i_2 \vee i_1 \xrightarrow{\text{sw}} i_2 \vee \exists i_3, i_1 \xrightarrow{\text{hb}} i_3 \xrightarrow{\text{hb}} i_2 \end{aligned}$$

► DEFINITION 14. Two memory accesses i_1 and i_2 are **conflicting** in an execution H if:

- $i_1, i_2 \in H.E$
- $H.loc(i_1) = H.loc(i_2)$
- At least one of i_1 and i_2 is a write

► DEFINITION 15. Two memory accesses i_1 and i_2 form a **data race** if:

- i_1 and i_2 are conflicting
- $\neg(i_1 \xrightarrow{\text{hb}} i_2 \vee i_2 \xrightarrow{\text{hb}} i_1)$

We say they form a **volatile-race** if both i_1 and i_2 are Volatile mode accesses.

Finally, our DRF-SC theorem is stated as the following:

► THEOREM 15 (DRF-SC). Given a program P , if all its SC-consistent executions are data-race-free or only have volatile-races, then the set of all *JAM*-consistent executions of P coincide with the set of SC-consistent executions.

Please see Appendix I for the proof.

We also provide the "model-agnostic" [19] version of the DRF-SC theorem:

► THEOREM 16 (EXECUTION-DRF). Any *JAM*-consistent execution that is data race free or only has volatile-races is SC-consistent.

Please see Appendix J for the proof.

Finally, we demonstrate the revised semantics preserves causality with acquire reads.

► THEOREM 17 (CAUSAL ACQUIRE-READS). [`coq/ReleaseAcquire.v`, `acq_causality`] If H is trace coherent and all reads in H are acquire-reads, then $\text{acyclic}(\xrightarrow{\text{po|rf}})$.

H.2 Volatile implies SC

Here we demonstrate that when all accesses are volatile, the program will have SC semantics.

To begin, we note that both full fences and **Volatile** pairs result in **push** orders in the formalism of [3]. That is, either a full fences or a **volint** edge implies a **push** edge. Our approach is to prove that when all program order accesses are **push** ordered then the semantics is SC. Thus, SC semantics follows as a corollary when all accesses are **volint**.

Recall from [3] that visibility order is acyclic. Intuitively, ordering induced by synchronization should not admit cycles.

► LEMMA 7 (ACYCLIC VISIBILITY). [coq/Truncate.v, vop_irreflex]
If H is trace coherent then, $\text{acyclic}(\text{--vvo--})$.

Next we show that the communication relation is not contradicted by visibility. Since the **com** relationship is composed from reads and coherence relationships, both of which encode the ordering of effects, we expect that visibility should not contradict such an ordering.

► LEMMA 8 (COMMUNICATION WRITE NOT-VISIBLE). [coq/SC/Volatile.v, coms_vo_contra]

If H is trace coherent, all accesses are executed, $i_1 \text{--com--}^* i_2$ and i_2 is a write then $\neg(i_2 \text{--vvo--}^+ i_1)$

Next we will establish that, when two pairs of **push** ordered accesses are connected by a possibly empty sequence of **com** edges, the first access of the first pair has been executed before the first access of the second pair. Intuitively, whenever there is a full fence between these two pairs of accesses then the order in which those fences executed must be consistent with the direction of the **com** relation.

► LEMMA 9 (PUSH TRACE-ORDERED). [coq/SC/Volatile.v, svo_comp_svo_to]

If H is trace coherent, $\text{acyclic}(\text{--co--})$, all accesses are executed, all accesses are **push** ordered, $i_1 \text{--push--} i_2 \text{--com--}^* i_3 \text{--push--} i_4$, then $i_1 \text{--to--} i_3$.

Proof. First, note that it is decidable whether i_3 is a write. We will begin by considering the case where it is a write. Since --to-- is total we consider each case for i_1 and i_3 . First, if $i_1 \text{--to--} i_3$ we are done. Second, for $i_1 = i_3$ we will demonstrate a contradiction. By assumption we have $i_1 \text{--push--} i_2$, then by substitution we have $i_3 \text{--push--} i_2$. By the definition of **push** we have $i_3 \text{--vo--} i_2$. By assumption we have $i_2 \text{--com--}^* i_3$. By Lemma 8 we have $\neg(i_3 \text{--vo--} i_2)$ and a contradiction. Finally, for $i_3 \text{--to--} i_1$ we will also demonstrate a contradiction. By assumption we have both $i_1 \text{--push--} i_2$ and $i_3 \text{--push--} i_4$. Then by the definition of --push-- and $i_3 \text{--to--} i_1$ we have $i_3 \text{--vo--} i_2$. As before we have $i_2 \text{--com--}^* i_3$. By Lemma 8 we have $\neg(i_3 \text{--vo--} i_2)$ and a contradiction.

Now consider the case where i_3 is not a write, then it must be a read and there exists some write i_w such that $i_w \text{--rf--} i_3$. It can be shown that $i_2 \text{--com--}^* i_w$. Thus we have $i_1 \text{--push--} i_2 \text{--com--}^* i_w \text{--rf--} i_3 \text{--push--} i_4$. Note that, because Lemma 8 applies to $i_2 \text{--com--}^* i_w$ we have that $\neg(i_2 \text{--vvo--}^+ i_w)$ and we must derive a contradiction by showing $i_w \text{--vvo--}^+ i_2$. For $i_1 = i_3$ we have $i_3 \text{--vvo--} i_2$ as before. Since $i_w \text{--rf--} i_3$ by the definition of --vvo-- we have $i_w \text{--vvo--} i_3$ and $i_w \text{--vvo--} i_2 \text{--vvo--} i_3$ as required. For $i_3 \text{--to--} i_1$ again we have that $i_3 \text{--vvo--} i_2$ and in turn we have $i_w \text{--vvo--} i_2 \text{--vvo--} i_3$ as required. ◀

Now we can demonstrate that when all accesses are **push** ordered the program semantics is SC. The key idea is that any cycle in the **sc** relation (i.e. a non-SC execution) will have at least one program order edge and at least one **com** edge. Thus we can show that the program order edge will appear twice in the cycle and then use Push Trace-ordered inductively to show that such a **push** order would have to execute before itself, thereby deriving a contradiction.

► THEOREM 18 (ALL PUSH SC). [coq/SC/Volatile.v, push_sc]
 If H is trace coherent, $\text{acyclic}(\xrightarrow{\text{co}})$, all accesses are executed, and all accesses are **push** ordered then $\text{acyclic}(\xrightarrow{\text{sc}})$.

Proof. We assume $i_1 \xrightarrow{\text{sc}} i_1$ and derive a contradiction. Observe that $i_1 \xrightarrow{\text{sc}} i_1$ must include at least one $\xrightarrow{\text{com}}$ edge because $\xrightarrow{\text{po}}$ is acyclic. Further observe that it must also include at least one program order edge because $\xrightarrow{\text{com}}$ is also acyclic. Thus there exists some access i_2 such that we can rearrange to obtain a sequence of program order and communication edges, $i_2(\xrightarrow{\text{po}}\xrightarrow{\text{com}}^+)^+i_2$. We proceed by induction on the length of this sequence. In the base case there exists some i_3 such that $i_2 \xrightarrow{\text{po}} i_3 \xrightarrow{\text{com}}^+ i_2$ which wraps around to give $i_2 \xrightarrow{\text{po}} i_3 \xrightarrow{\text{com}}^+ i_2 \xrightarrow{\text{po}} i_3$. Then Push Trace-ordered applies to give $i_2 \xrightarrow{\text{to}} i_2$, but this is a contradiction since the trace order is total. In the inductive case we use the same argument and connect the trace order from the inductive hypothesis to give a contradiction. ◀

From Theorem 18 we can derive two corollaries. The first shows that when all accesses are **Volatile** the semantics is SC. The second shows that when all accesses have full fences between them, represented by **spush** in the model, the semantics is SC. The structure of the model and our definition for **Volatile** accesses shines through here as both results follow directly from a single result about the behavior of full fences.

► COROLLARY 4 (ALL VOLATILE SC). [coq/SC/Volatile.v, volatile_sc]
 If H is trace coherent, $\text{acyclic}(\xrightarrow{\text{co}})$, all accesses are executed, and all accesses are **Volatile** mode accesses then $\text{acyclic}(\xrightarrow{\text{sc}})$.

Proof. If all accesses are volatile then any two program order accesses are **push** ordered and we can appeal to Theorem 18. ◀

► COROLLARY 5 (ALL SPECIFIED PUSH SC). [coq/SC/Volatile.v, spush_sc]
 If H is trace coherent, $\text{acyclic}(\xrightarrow{\text{co}})$, all accesses are executed, and all program ordered (**po**) accesses are ordered by specified push order (**spush**) then $\text{acyclic}(\xrightarrow{\text{sc}})$.

Proof. If any two program order accesses have a specified push order then they are similarly **push** ordered and we can again appeal to Theorem 18. ◀

I The Standard DRF-SC Theorem

► **THEOREM 15 (DRF-SC).** Given a program P , if all its SC-consistent executions are data-race-free or only have volatile-races, then the set of all JAM -consistent executions of P coincide with the set of SC-consistent executions.

Proof. Let P be a program, and suppose all its SC-consistent executions only has Volatile-races. We want to show that P has no weak behavior. Toward contradiction, let's assume there exists an execution H of P such that H is JAM -consistent but not SC-consistent.

► **DEFINITION 16.** An execution H' is called a *prefix* of an execution H if H' is obtained by restricting H to a set of events E such that:

1. the set of initialization events $E_0 \in E$
2. for any event $b \in E$, if there is $a \xrightarrow{\text{po}} b$ or $a \xrightarrow{\text{rf}} b$ in H , then $a \in E$. (Closed with respect to $(H.\text{po} \cup H.\text{rf})$)

1.0.0.1 Claim 1

Any prefix of a JAM -consistent execution is JAM -consistent.

1.0.0.2 Claim 2

Any prefix of an SC-consistent execution is SC-consistent.

Proof. The above two claims are true because a prefix consists of a subset of edges and events of the original execution. If there is a cycle that violates the requirements of the memory models, then the same cycle is present in the original execution graph. Therefore, since we assumed H is JAM -consistent, any prefix of H is also JAM -consistent. ◀

1.0.0.3 Notations

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 $\{H.\text{AccessMode}(a), H.\text{AccessMode}(b)\} \neq \{\text{Volatile}\}$ and $\langle a, b \rangle, \langle b, a \rangle \notin (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$.

$$\Pi(E) = \{\langle a, b \rangle \in E \times E \mid \{H.\text{AccessMode}(a), H.\text{AccessMode}(b)\} \neq \{\text{Volatile}\}, \\ \langle a, b \rangle, \langle b, a \rangle \notin (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+\}$$

Let a_1, \dots, a_n be an enumeration of events ordered by trace orders (recall that trace order is a total order among the events in an execution that is compatible with $(H.\text{po} \cup H.\text{rf})^+$).

Let E_i denotes the subset of events $E_0 \cup \{a_1, \dots, a_i\}$ and H_i be the execution restrict to E_i . This is easy to see that each H_i is a prefix to H because the trace order is compatible with $(H.\text{po} \cup H.\text{rf})^+$. Therefore, H_i is also JAM -consistent by **Claim 1**.

1.0.0.4 Claim 3

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

Proof. Suppose that $\Pi(E_i) = \emptyset$. Then, for every conflicting pair $\langle a, b \rangle$, either $H.\text{AccessMode}(a) = H.\text{AccessMode}(b) = \text{Volatile}$ or $\langle a, b \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$ or $\langle b, a \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$.

1. For any $a \xrightarrow{\text{rf}} b$ in H_i .
 - If $H.\text{AccessMode}(a) = H.\text{AccessMode}(b) = \text{Volatile}$, then $\langle a, b \rangle \in H.\text{rf} \upharpoonright_{\text{Volatile}}$.
 - If $\langle a, b \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$, then $\langle a, b \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$.
 - If $\langle b, a \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$, then there is a $(\text{po} \cup \text{rf})^+$ cycle between a and b . By the **NO-THIN-AIR** requirement, H_i is not *JAM*-consistent. Contradicting to our previous assumption. Therefore it is impossible to have $\langle b, a \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$.

Thus, $H_i.\text{rf} \subseteq (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$

2. For any $a \xrightarrow{\text{co}} b$ in H_i ,
 - If $H.\text{AccessMode}(a) = H.\text{AccessMode}(b) = \text{Volatile}$, then $\langle a, b \rangle \in H.\text{co} \upharpoonright_{\text{Volatile}}$.
 - If $\langle a, b \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$, then $\langle a, b \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$.
 - If $\langle b, a \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$, then the the domains of the $H.\text{rf}$ on the path from b to a has access mode equal to *Volatile* which includes *Release* semantics. Similarly, the ranges of the $H.\text{rf}$ have access mode equal to *Volatile* which includes *Acquire* semantics. Therefore, $\text{po} \subseteq \text{ra}$ on this path. That is, we have $\langle b, a \rangle \in (H.\text{ra} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+ \subseteq H.\text{vvo}^+$. By *coww*, we have $\langle b, a \rangle \in H.\text{co}$. With $a \xrightarrow{\text{co}} b$, we now have a *co* cycle, contradicting to the earlier assumption that H_i is *JAM*-consistent. Therefore, it is impossible that $\langle b, a \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$.

Thus, we have $H_i.\text{co} \subseteq (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+ \cup H.\text{co} \upharpoonright_{\text{Volatile}}$.

3. For any $a \xrightarrow{\text{fr}} b$ in H_i ,
 - If $H.\text{AccessMode}(a) = H.\text{AccessMode}(b) = \text{Volatile}$, then $\langle a, b \rangle \in H.\text{fr} \upharpoonright_{\text{Volatile}}$.
 - If $\langle a, b \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$, then $\langle a, b \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$.
 - If $\langle b, a \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$, then the the domains of the $H.\text{rf}$ on the path from b to a has access mode equal to *Volatile* which includes *Release* semantics. Similarly, the ranges of the $H.\text{rf}$ have access mode equal to *Volatile* which includes *Acquire* semantics. Therefore, $\text{po} \subseteq \text{ra}$ on this path. That is, we have $\langle b, a \rangle \in (H.\text{ra} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+ \subseteq H.\text{vvo}^+$. Expanding the definition of *fr*, there exists a write event i such that $\langle i, a \rangle \in H_i.\text{rf}$ and $\langle i, b \rangle \in H_i.\text{co}$. By *cowr*, we have $\langle b, i \rangle \in H_i.\text{co}$. Now we have a *co* cycle, contradicting with the earlier assumption that H_i is *JAM*-consistent. Therefore, it is impossible that $\langle b, a \rangle \in (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+$.

Thus, we have $H_i.\text{fr} \subseteq (H.\text{po} \cup H.\text{rf} \upharpoonright_{\text{Volatile}})^+ \cup H.\text{fr} \upharpoonright_{\text{Volatile}}$.

Therefore, we have $H_i.\text{po} \cup \text{rf} \cup \text{fr} \cup \text{co} \subseteq H.\text{po} \cup \text{rf} \upharpoonright_{\text{Volatile}} \cup \text{fr} \upharpoonright_{\text{Volatile}} \cup \text{co} \upharpoonright_{\text{Volatile}}$. Since H_i is *JAM*-consistent, any prefix of H_i is *JAM*-consistent as well (by **Claim 1**). By our previous lemma that *Volatile* \Rightarrow *SC*, any prefix of H_i with all events marked as *Volatile* are *SC*-consistent. If H_i is not *SC*-consistent and there is a cycle of $H.\text{po} \cup \text{rf} \upharpoonright_{\text{Volatile}} \cup \text{fr} \upharpoonright_{\text{Volatile}} \cup \text{co} \upharpoonright_{\text{Volatile}}$, then there exists a prefix of H_i containing all the events in this cycle and hence not *SC*-consistent. This contradicts with our previous assumption. Thus, H_i is *SC*-consistent. \blacktriangleleft

Now, continuing our proof, since H is not *SC*-consistent, we know that $\Pi(E) \neq \emptyset$.

1.0.0.5 Convention

In the rest of the proof, we treat an Read-modify-write (RMW) event as two separate events ordered by the *rmw* order. That is, each RMW event in H consists of two events i_1 and i_2 such that $\langle i_1, i_2 \rangle \in H.\text{rmw}$, where $i_1 \in H.\text{R}$ and $i_2 \in H.\text{W}$.

Let $k = \min\{i \mid \Pi(E_i) \neq \emptyset\}$. Then it is clear that H_{k-1} is the maximum *SC*-consistent prefix of H . That is, $\Pi(E_{k-1}) = \emptyset$ and H_{k-1} is *SC*-consistent. We also have $\Pi(E_k) \neq \emptyset$ and H_k is

not SC-consistent. That is, there exists $j < k$ such that $\{H.AccessMode(a_j), H.AccessMode(a_k)\} \neq \{Volatile\}$, and $\langle a_j, a_k \rangle, \langle a_k, a_j \rangle \notin H.po \cup H.rf \upharpoonright_{Volatile}$.

I.0.0.6 Claim 4

Let $B = \{b \mid \langle b, a_k \rangle \in H_k.po\}$, then $\langle a_j, b \rangle \notin (H.po \cup H.rf)^+$.

Proof. Since $\Pi(E_{k-1}) = \emptyset$, we have $H_{k-1}.rf \subseteq (H.po \cup H.rf \upharpoonright_{Volatile})^+$. If $\langle a_j, b \rangle \in (H.po \cup H.rf)^+$, then we have $\langle a_j, b \rangle \in (H.po \cup H.rf \upharpoonright_{Volatile})^+$. From that we have $\langle a_j, a_k \rangle \in (H.po \cup H.rf \upharpoonright_{Volatile})^+$ since $\langle b, a_k \rangle \in H.po$, which contradicts to our previous assumption. Therefore, $\langle a_j, b \rangle \notin (H.po \cup H.rf)^+$. ◀

1. $a_k \in H.W$.

I.0.0.7 Claim 5

$\langle a_j, a_k \rangle$ forms a data race.

Proof. Since H_k is closed under $(H.po \cup H.rf)$ and a_k is the last event in the total trace order in H_k , there is no outgoing edge from a_k . Therefore, $\langle a_k, a_j \rangle \notin (H.po \cup H.rf)^+$. In addition, Since a_k is a write, we cannot have $\langle a_j, a_k \rangle \in H.rf$. Therefore, if $\langle a_j, a_k \rangle \in H.hb$, then it would imply that $\langle a_j, b \rangle \in (H.po \cup H.rf)^+$ for some b such that $\langle b, a_k \rangle \in H.po$, which contradicts with **Claim 4**. Therefore, $\langle a_j, a_k \rangle \notin (H.po \cup H.rf)^+$ and $\langle a_j, a_k \rangle$ forms a race in H_k . ◀

I.0.0.8 Claim 6

H_k is not SC-consistent.

Proof. Given that $\langle a_j, a_k \rangle$ forms a data race in H_k and $\{H.AccessMode(a_j), H.AccessMode(a_k)\} \neq \{Volatile\}$, this follows from the assumption. ◀

I.0.0.9 Claim 7

There does not exist a read event b such that $\langle b, a_k \rangle \in H.rmw$.

Proof. Suppose toward contradiction that there is b such $\langle b, a_k \rangle \in H.rmw$. Since H_k is not SC-consistent, there is cycle of $(H_k.po \cup H_k.rf \cup H_k.fr \cup H_k.co)^+$. Since H_{k-1} is SC-consistent, it must be that a_k is part of the cycle in H_k . That is, there is a $(H_k.po \cup H_k.rf \cup H_k.fr \cup H_k.co)$ edge from a_k . Additionally, it cannot be a **po** or **rf** because H_k is a closed prefix of H . Since a_k is a write, it cannot be **fr** either. Therefore, there is some c in H_{k-1} such that $\langle a_k, c \rangle \in H.co$ and $\langle c, a_k \rangle \in (H.po \cup H.rf \cup H.fr \cup H.co)^+$. Extending the edges, we have $\langle c, d \rangle \in (H.po \cup H.rf \cup H.fr \cup H.co)^*$, and $\langle d, a_k \rangle \in (H.co \cup H.fr \cup H.po)$. We analyze each case below.

- $\langle d, a_k \rangle \in H.co$. Then we know that $c, d \in H_{k-1}$ are writes to the same location. Since H_{k-1} is SC-consistent, and $\langle c, d \rangle \in (H_{k-1}.po \cup H_{k-1}.rf \cup H_{k-1}.fr \cup H_{k-1}.co)^*$, we have $\langle c, d \rangle \in H_{k-1}.co^*$. If $c = d$, then we have a $H_k.co$ cycle between c and a_k , making H_k not *JAM*-consistent. Therefore, we have $\langle c, d \rangle \in H_{k-1}.co$. But we also have $\langle d, a_k \rangle, \langle a_k, c \rangle \in H_k.co$. Together, they yield a **co** cycle, contradicting with our earlier assumption that H_k is *JAM*-consistent.

- $\langle d, a_k \rangle \in H.\mathbf{fr}$. Then we know d is a read. Given that $\langle a_k, c \rangle \in H.\mathbf{co}$, we can infer that $\langle d, c \rangle \in H_{k-1}.\mathbf{fr}$. But we also have $\langle c, d \rangle \in (H_{k-1}.\mathbf{po} \cup H_{k-1}.\mathbf{rf} \cup H_{k-1}.\mathbf{fr} \cup H_{k-1}.\mathbf{co})^*$, which forms a cycle between c and d , contradicting to the assumption that H_{k-1} is SC-consistent.
- $\langle d, a_k \rangle \in H.\mathbf{po}$. Then we have $\langle d, b \rangle \in H_{k-1}.\mathbf{po}^?$. Since $\langle b, a_k \rangle \in H_k.\mathbf{rmw}$, by `corrmwexcl`, we know that there exists some e such that $\langle e, b \rangle \in H_{k-1}.\mathbf{rf}$ and $\langle e, c \rangle \in H_{k-1}.\mathbf{co}$. Therefore, $\langle b, c \rangle \in H_{k-1}.\mathbf{fr}$. However, we also have $\langle c, d \rangle \in (H_{k-1}.\mathbf{po} \cup H_{k-1}.\mathbf{rf} \cup H_{k-1}.\mathbf{fr} \cup H_{k-1}.\mathbf{co})^*$ and $\langle d, b \rangle \in H_{k-1}.\mathbf{po}^?$. Now we have a cycle contradicting to the assumption that H_{k-1} is SC-consistent. ◀

Now let H'_k be a transformation of H_k such that for all $\langle a_k, c \rangle \in H_k.\mathbf{co}$, we have $\langle c, a_k \rangle \in H'_k.\mathbf{co}$. Since there is no b such that $\langle b, a_k \rangle \in H_k.\mathbf{rmw}$, transforming the `co` edges in H_k in this way won't affect any other `fr` or `rf` edges in H_k . As a result, H'_k is the same as H_k except for the `co` edges. Moreover, H'_k is SC-consistent because we have established that the only way to form a cycle in H_k is to have $\langle a_k, c \rangle \in H_k.\mathbf{co}$ for some c and H'_k essentially breaks the cycle by flipping the `co` arrows. However, we still have $\langle a_j, a_k \rangle$ forming a race in H'_k and $\{H'_k.\mathit{AccessMode}(a_j), H'_k.\mathit{AccessMode}(a_k)\} \neq \{\mathit{Volatile}\}$.

I.0.0.10 Claim 8

All prefixes of SC-consistent executions of P are data race free.

Proof. This follows from the fact that prefixes are closed under `po` \cup `rf`. ◀

I.0.0.11 Claim 9

H'_k is a prefix of some SC-consistent execution of P .

Proof. Since H'_k is SC-consistent and closed under `po` \cup `rf`, we can construct an SC-consistent execution H' such that for all event $i \in H'.E \setminus H'_k.E$ and $e \in H'_k.E$, we have $\langle e, i \rangle \in H'.\mathbf{to}$ (trace order). ◀

It is clear that the fact H'_k has a data race $\langle a_j, a_k \rangle$ contradicts with **Claim 8** and **Claim 9**.

2. $a_k \in H.R.$

Then we know that a_j is a write.

Let $E = \{a \in H.E \mid \langle a, a_k \rangle \in (H.\mathbf{po} \cup H_{k-1}.\mathbf{rf})^*\} \cup \{a \in H.E \mid \langle a, a_j \rangle \in (H.\mathbf{po} \cup H_{k-1}.\mathbf{rf})^*\}$. Note that there does not exist any $a \in E$ such that $\langle a_j, a \rangle \in (H.\mathbf{po} \cup H_{k-1}.\mathbf{rf})^*$ and $\langle a, a_k \rangle \in (H.\mathbf{po} \cup H_{k-1}.\mathbf{rf})^*$ since H_{k-1} is closed under $(H.\mathbf{po} \cup H.\mathbf{rf})$. Let H' be the restriction of H_k to the events in E .

I.0.0.12 Claim 10

$\langle a_j, a_k \rangle$ forms a data race in H'

Proof. Since a_k is a read event, there is no out-going `rf` edge from a_k . Because H' itself is closed under $H.\mathbf{po} \cup H.\mathbf{rf}$, we have $\langle a_k, a_j \rangle \notin (H'.\mathbf{po} \cup H'.\mathbf{rf})^+$. In addition, there is no out-going edge from a_j in H' , so $\langle a_j, a_k \rangle \notin (H'.\mathbf{po} \cup H'.\mathbf{rf})^+$. ◀

Now, we want to construct an SC-consistent execution that contains the race $\langle a_j, a_k \rangle$ to show a contradiction.

Let x be the location a_k accesses and c be the last write to x according to $H'.\text{co}$.

- $c \neq a_j$. Let d be the write to x such that $\langle d, a_k \rangle \in H'.\text{rf}$. We transform H' so that a_k reads from c instead of d . That is, we remove $\langle d, a_k \rangle$ from $H'.\text{rf}$, change the value of a_k to the value of c , and add $\langle c, a_k \rangle$ to $H'.\text{rf}$. The immediate consequence of this transformation is, for all e such that $\langle d, e \rangle \in H'.\text{co}$, the **fr** edge from a_k to e are also removed. Since a_k is a read event, there cannot be **rf** or **co** edge going out from a_k . As a result, we cannot form a $(\text{po} \cup \text{rf} \cup \text{co} \cup \text{fr})^+$ cycle with a_k and the resulting execution is SC-consistent. Then the fact that $\langle a_j, a_k \rangle$ is a data race gives us a contradiction.
- $c = a_j$. That is, a_k forms a race with the last write to the same location. Let d be the write to x such that $\langle d, a_k \rangle \in H'.\text{rf}$. We transform H' so that a_k reads from the write that immediately **co** ordered before c . Let e be that write. We remove $\langle d, a_k \rangle$ from $H'.\text{rf}$, change the value of a_k to the value of e , and add $\langle e, a_k \rangle$ to $H'.\text{rf}$. Since $\langle e, c \rangle \in H'.\text{co}$, we have $\langle a_k, c \rangle \in H'.\text{fr}$. However, since $c = a_j$ and $\langle a_j, a_k \rangle$ forms a race, $\langle c, a_k \rangle \notin (H'.\text{po} \cup H'.\text{rf})^+$. That is, there is no path from c to a_k in H' . As a result, we cannot form any cycle with the added **fr** edge from a_k and the execution after the transformation is SC-consistent. Since we still have $\langle c, a_k \rangle$ forming a race, we now have a contradiction.

◀

J A Proof of "Model-agnostic" DRF-SC

► **THEOREM 16 (EXECUTION-DRF).** Any *JAM*-consistent execution that is data race free or only has volatile-races is SC-consistent.

Proof. Let P be a program. We first consider the case without any Volatile-race. That is, all conflicting pairs of accesses are ordered by the happens-before (**hb**) order in some SC-consistent execution of P . We prove that there does not exist any *JAM*-consistent execution of P that is not SC-consistent. Suppose toward a contradiction that there exists an execution H of P such that H is *JAM*-consistent, race-free, but not SC-consistent. That is, H has a $(\text{po} \cup \text{rf} \cup \text{fr} \cup \text{co})^+$ cycle. In addition, we assume that $\text{co-jam} \subseteq \text{co}$.

First note that each of the communication edges in the cycle of H are also pairs of conflicting accesses. Indeed, they are defined between accesses to the same location and at least one of the accesses is a write event. Then, by our assumption that H data-race-free, they are also ordered by the **hb** order. In addition, for all conflicting accesses i_1 and i_2 in H :

- $i_1 \xrightarrow{\text{rf}} i_2 \Rightarrow i_1 \xrightarrow{\text{hb}} i_2$
- $i_1 \xrightarrow{\text{fr}} i_2 \Rightarrow i_1 \xrightarrow{\text{hb}} i_2$
- $i_1 \xrightarrow{\text{co}} i_2 \Rightarrow i_1 \xrightarrow{\text{hb}} i_2$

because the other direction can immediately lead to contradictions.

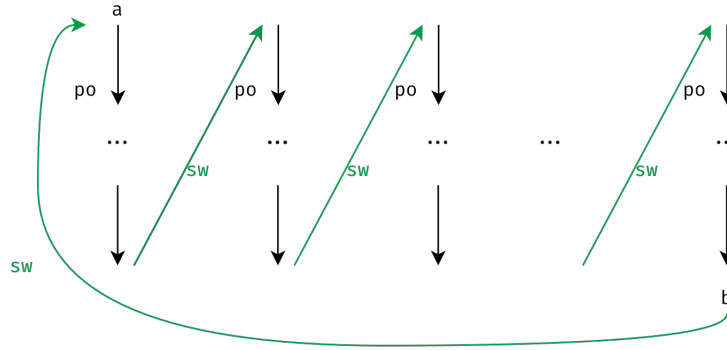
Then the cycle in H a cycle of $(\text{po} \cup \text{hb})^+$. We extend the definition of **hb**, we have a cycle of $(\text{po} \cup (\text{po} \cup \text{sw})^+)^+$, which simplifies to a $(\text{po} \cup \text{sw})^+$ cycle. Graphically the cycle has a shape shown in Figure 17.

By the definition of **sw** order, the domain of each **sync** edge is a Release mode write (or a release fence followed by a write) and the range of **sw** is an Acquire mode read (or a read followed by an acquire fence). Therefore, we know that the "head" of each thread in this cycle is an Acquire read and the "end" of each thread in this cycle is a Release write. By the

semantics of Release-Acquire mode, all of the `po` order to a Release write and all of the `po` order from an Acquire read is preserved and captured in the `ra` order, which is a subset of `vo`. In addition, $\text{sw} \subseteq \text{vo}$. As a result, the cycle in H is actually a `vo` cycle. However, we assumed that H is *JAM*-consistent and by the previous lemma by Bender et al. [3], the `vo` order is acyclic in all *JAM*-consistent executions. Thus a contradiction.

We now consider the case where there are Volatile-races in the execution. We prove this by incrementally inserting a pair of Volatile-race into an execution that is data-race-free and *JAM*-consistent and prove that it does not introduce any weak behavior to the execution. Let H' be such an execution. As we just have shown, H' is SC-consistent. We would like to insert a pair of Volatile-race $\langle a, b \rangle$ into H' . Let T_1 be the thread where a is inserted and T_2 be the thread where b is inserted. By definition of data race, $T_1 \neq T_2$. We have three possible cases:

- T_1 and T_2 are not connected by any $(\text{po} \cup \text{sw})^+$ edge before inserting $\langle a, b \rangle$ in H' . That is, there does not exist any $(\text{po} \cup \text{sw})^+$ path from T_1 to T_2 . Then inserting $\langle a, b \rangle$ into the execution cannot form any cycle in $(\text{po} \cup \text{rf} \cup \text{fr} \cup \text{co})^+$ since it can only add at most one edge to the graph.
- There is a $(\text{po} \cup \text{sw})^+$ path from T_1 to T_2 . First note that H' before inserting $\langle a, b \rangle$ has a similar shape to H in Figure 17 except that it does not have a cycle. That is, if two threads are connected, then they must be connected by at least an `sw` edge. This implies there is a release write W^{Rel} on T_1 , an acquire read R^{Acq} on T_2 , and $W^{Rel} \xrightarrow{\text{sw}} \xrightarrow{(\text{po}|\text{sw})^*} R^{Acq}$. Due to this structure, the only way to insert $\langle a, b \rangle$ and build a cycle of $(\text{po} \cup \text{rf} \cup \text{fr} \cup \text{co})^+$ is to insert a before W^{Rel} and insert b after R^{Acq} . Note that this implies $a \xrightarrow{\text{ra}} W^{Rel} \xrightarrow{\text{sw}} \xrightarrow{(\text{po}|\text{sw})^*} R^{Acq} \xrightarrow{\text{ra}} b$. So depending on what type of access a and b are, we have three cases:
 - a is a Volatile write and b is a Volatile read and $b \xrightarrow{\text{fr}} a$. By definition of `fr`, there exists a write i such that $i \xrightarrow{\text{rf}} b$ and $i \xrightarrow{\text{co}} a$. But $a \xrightarrow{\text{ra}} W^{Rel} \xrightarrow{\text{sw}} \xrightarrow{(\text{po}|\text{sw})^*} R^{Acq} \xrightarrow{\text{ra}} b$ implies that $a \xrightarrow{\text{vo}} b$, and by the `cowr` coherence rule, we have $a \xrightarrow{\text{co}} i$. Now we have a `co` cycle, contradicting to our assumption that H'' is *JAM*-consistency.
 - a is a Volatile read and b is a Volatile write and $b \xrightarrow{\text{rf}} a$. Again, we have $a \xrightarrow{\text{vo}} b$ and $b \xrightarrow{\text{rf}} a$. Since $\text{rf} \subseteq \text{vo}$, we get a `vo` cycle, contradicting to our previous assumption of *JAM*-consistency.
 - a is a Volatile write and b is a Volatile write and $b \xrightarrow{\text{co}} a$. Again, we have $a \xrightarrow{\text{vo}} b$ and $b \xrightarrow{\text{co}} a$. By `coww`, we have $a \xrightarrow{\text{co}} b$ and $b \xrightarrow{\text{co}} a$, a coherence cycle contradicting to our previous assumption of *JAM*-consistency.
- There is a $(\text{po} \cup \text{sw})^+$ path from T_2 to T_1 . Symmetrical to the previous case. ◀



■ **Figure 17** The Shape of the Cycle in H

K Validation

In this section, we explain our implementation of Java architecture for Herd7 [1] and experimental results with the JAM_{21} model. The source code of our Java architecture implementation will become available for artifact evaluation.

K.1 Methods supported by Java Architecture for Herd7

The list of supported methods in our implementation of Java in Herd7 [1] can be found in the following table. We provide a description for each one of the method and its corresponding action in Herd7.

Method	Memory Action	Description
<code>getM()</code>	R(M)	Read operation with access mode specified by M, where M can be omitted (Plain mode), Opaque , Acquire , or Volatile .
<code>setM(val)</code>	W(M)	Write operation with access mode specified by M, where M can be omitted (Plain mode), Opaque , Acquire , or Volatile , the value written is specified by <code>val</code> , which can be either an integer or a local variable.
<code>compareAndExchangeM(expect, dest)</code>	RMW(M)	An atomic compare and update operation with access mode specified by M, where M can be omitted (Volatile mode), Acquire , or Release .
<code>getAndOpM(val)</code>	RMW(M)	A numeric or bitwise atomic update operation with modifying operation specified by <code>Op</code> and access mode specified by M, where <code>Op</code> can be Add , And , Or , or Xor ; and M can be omitted, Acquire , or Release .
<code>fullFence()</code>	F(Volatile)	A full fence.
<code>releaseFence()</code>	F(Release)	A release fence.
<code>acquireFence()</code>	F(Acquire)	An acquire fence.

■ **Table 1** Methods supported by the Java Architecture

Name	JAM_{19}	JAM_{21}
volatile-non-sc.4	Sometimes	Never
volatile-non-sc.5	Sometimes	Never

■ **Figure 18** volatile-non-sc Experimental Results

K.2 Experimental Results

In this section, we show the experimental results of running the same set of litmus tests as in JAM_{19} and compare their outcomes. Three types of results can be yielded by Herd7 at the end, **Always**, **Sometimes**, and **Never**. **Always** and **Sometimes** means the behavior specified in the litmus test is allowed, whereas **Never** means it is forbidden.

Fig. 18 shows the experimental results of running `volatile-non-sc.4` example and its 5-thread version with the JAM_{19} and the JAM_{21} model. As we expected, the update to the JAM_{21} model fixes the issue we addressed earlier in the paper and the executions changed from **Sometimes** to **Forbidden**.


Fig. 19 shows the rest of the experimental results in details. Note that not all litmus tests used for JAM_{19} [3] are translatable to Java. We marked those non-translatable tests as “N/A” in the tables (since Java does not have the notion of address dependency). The result agrees with our expectation that most of the litmus tests yield the same results as JAM_{19} , except those that are related to the inconsistency issue (highlighted using bold font). We discuss each of the exceptions.

The execution graphs of `IRIW-acq-sc` are shown in Fig. 20. Our experimental results show that this execution is forbidden under JAM_{19} but is allowed under JAM_{21} . Under the JAM_{19} model, because the definition of `volint` includes orders from any instruction to a Volatile read program ordered after the instruction, we have $(c) \xrightarrow{\text{volint}} (d)$ and $(e) \xrightarrow{\text{volint}} (f)$. Between the two threads (P3 and P4), there is the visibility order $(c) \xrightarrow{\text{vo}} (f)$, or $(e) \xrightarrow{\text{vo}} (d)$. Both cases can produce the contradictory result that one of the threads observes the non-initialization write before the initialization write, i.e., a coherence cycle. Therefore, this execution is forbidden under the JAM_{19} model. In JAM_{21} , the two `volint` orders are no longer present in the execution graph because the new definition of `volint` requires both of the memory accesses to be Volatile. As a result, the execution becomes allowed under the JAM_{21} model. To see why allowing this execution is an improvement, note that the JAM_{19} model only captures the “leading fence” convention that `fullFence()` are inserted before Volatile accesses. On the other hand, if the compiler follows the “trailing fence” convention, there would not be `fullFence()`s in P3 and P4. In that case, the execution is allowed. In order to accommodate both conventions, the JAM_{21} model relaxes to allow this execution.

The execution graph of `Z6.U` are shown in Fig. 21. Due to the original problematic encoding of Volatile writes, `volint` includes orders from Volatile writes to any program ordered later memory accesses. Therefore, $(a) \xrightarrow{\text{volint}} (b)$, $(c) \xrightarrow{\text{volint}} (d)$, and $(e) \xrightarrow{\text{volint}} (f)$. Similar to the previous example, there are two possible visibility orders, $(a) \xrightarrow{\text{vo}} (f)$ or $(e) \xrightarrow{\text{vo}} (b)$. The former case leads to the derivation of $(a) \xrightarrow{\text{co}} (Wx=0)$, which contradicts the assumption that all initialization writes to variable `x` are ordered before all non-initialization writes to `x`. The latter case leads to a contradiction as well. Because (d) reads the value written by (e) and $(c) \xrightarrow{\text{volint}} (d)$, we can infer that $(c) \xrightarrow{\text{co}} (e)$. If $(e) \xrightarrow{\text{vo}} (b)$, then $(e) \xrightarrow{\text{vo}} (c)$. By the coww rule, $(e) \xrightarrow{\text{co}} (c)$. This leads to a coherence `co` cycle between (c) and (e) . The JAM_{21} model relaxes the `volint` edges in P1 and P2 in order to accommodate both the leading fence convention and the trailing fence convention. If the compiler follows the convention of inserting `fullFence()` before the Volatile accesses, there is only $(a) \xrightarrow{\text{ra}} (b)$

Name	JAM ₁₉	JAM ₂₁	Name	JAM ₁₉	JAM ₂₁
WRC+addrs	Never	N/A	2+2W	Never	Never
LB+data+data-wsi	Never	N/A	IRIW-acq-sc	Never	Sometimes
W+RR	Never	Never	RWC+syncs	Never	Never
totalco	Never	Never	W+RWC	Never	Never
PPOCA	Sometimes	N/A	Z6.U	Never	Sometimes
IRIW	Sometimes	Sometimes	IRIW-sc-rlx-acq	Never	Sometimes
IRIW+addrs	Sometimes	N/A	cppmem_iriw_relacq	Sometimes	Sometimes
IRIW+poaas+LL	Sometimes	Sometimes	cppmem_sc_atomics	Never	Never
IRIW+poaps+LL	Sometimes	Sometimes	iriw_sc	Never	Never
MP+dmb.sy+addr-ws-rf-addr	Sometimes	N/A	mp_fences	Never	Never
WW+RR+WW+RR+wsilp+poaa+wsilp+poaa	Sometimes	Sometimes	mp_relacq	Never	Never
LB	Never	Never	mp_relacq_rs	Sometimes	Sometimes
			mp_relaxed	Sometimes	Sometimes
			mp_sc	Never	Never
			4.SB	Sometimes	Sometimes
			6.SB	timeout	timeout
			6.SB+prefetch	timeout	timeout
			CoRWR	Never	Never
			SB+SC	Sometimes	Sometimes
			SB+mfences	Never	Never
			SB+rfi-pos	Sometimes	Sometimes
			SB	Sometimes	Sometimes
			X000	Sometimes	Sometimes
			X001	Sometimes	Sometimes
			X002	Sometimes	Sometimes
			X003	Sometimes	Sometimes
			X004	Sometimes	Sometimes
			X005	Sometimes	Sometimes
			X006	Sometimes	Sometimes
			iriw-internal	Sometimes	Sometimes
			iriw	Sometimes	Sometimes
			podrw000	Sometimes	Sometimes
			podrw001	Sometimes	Sometimes
			x86-2+2W	Sometimes	Sometimes
a1	Sometimes	Sometimes			
a1_reorder	Sometimes	Sometimes			
a3	Sometimes	Sometimes			
a3_reorder	Sometimes	Sometimes			
a3v2	Sometimes	Sometimes			
a4	Never	Never			
a4_reorder	Sometimes	Sometimes			
arfna	Never	Never			
arfna_transformed	Never	Never			
b	Never	Never			
b_reorder	Sometimes	Sometimes			
c	Never	Never			
c_p	Never	Never			
c_p_reorder	Never	Never			
c_pq	Never	Never			
c_pq_reorder	Never	Never			
c_q	Never	Never			
c_q_reorder	Never	Never			
c_reorder	Never	Never			
cyc	Never	Never			
cyc_na	Sometimes	Sometimes			
fig1	Always	Always			
fig6	timed out	time out			
fig6_translated	timed out	time out			
lb	Never	Never			
linearisation	Never	Never			
linearisation2	Never	Never			
roachmotel	Never	Never			
roachmotel2	Never	Never			
rseq_weak	Sometimes	Sometimes			
rseq_weak2	Always	Always			
seq	Never	Never			
seq2	Never	Never			
strengthen	Never	Never			
strengthen2	Never	Never			


■ **Figure 19** Litmus Test Comparisons



(a) Before: Forbidden

(b) After: Allowed


Figure 20 IRIW-acq-sc



(a) Before: Forbidden

(b) After: Allowed

Figure 21 Z6.U



(a) Before: Forbidden

(b) After: Allowed

Figure 22 IRIW-seq-rlx

Name	Power [14]
volatile-non-sc.4.ppc	Sometimes
volatile-non-sc.5.ppc	Sometimes

■ **Figure 23** volatile-non-sc on Power with the incorrect compilation scheme

in P1 and no synchronization between the two instructions in P2. Thus this execution is allowed under the JAM_{21} model.

Lastly, the execution graphs of IRIW-seq-rlx are shown in Fig. 22. Originally, due to the old encoding of Volatile writes, (a) $\xrightarrow{\text{volint}}$ (b) and (c) $\xrightarrow{\text{volint}}$ (d). Two possible visibility orders can be inferred, either (a) $\xrightarrow{\text{vo}}$ (d) or (c) $\xrightarrow{\text{vo}}$ (b). The former case leads to the conclusion that (a) $\xrightarrow{\text{co}}$ (Wx=0) because (a) $\xrightarrow{\text{vo}}$ (d) $\xrightarrow{\text{rf}}$ (g) $\xrightarrow{\text{ra}}$ (h) and (Wx=0) $\xrightarrow{\text{rf}}$ (h). Similarly, the latter case leads to the conclusion that (c) $\xrightarrow{\text{co}}$ (Wy=0) because (c) $\xrightarrow{\text{vo}}$ (b) $\xrightarrow{\text{rf}}$ (e) $\xrightarrow{\text{ra}}$ (f) and (Wy=0) $\xrightarrow{\text{rf}}$ (f). Each of the two conclusions contradicts the assumption that initialization writes are coherence `co` ordered before non-initialization writes. Therefore this execution is forbidden by the JAM_{19} model. In the JAM_{21} model, we relax the `volint` order in P1 and P2 to include the situation of which the compiler inserts the `fullFence()` before Volatile accesses. Thus, under the new JAM_{21} model, this execution is allowed.

In summary, the JAM_{21} model has two main differences comparing to the JAM_{19} model. First, under the JAM_{21} model, when all memory accesses use Volatile mode, the execution is guaranteed to be sequentially consistent, whereas the old JAM_{19} model has the inconsistency issue we pointed out earlier. Second, when mixing Volatile and other access modes in a program, the new JAM_{21} model accommodates both the "leading fence" convention and the "trailing fence" convention so that the compiler is free to choose either one to implement.

K.3 Compilation to Power

We translated the volatile-non-sc.4 and the volatile-non-sc.5 example to Power instructions according to the original compilation scheme:

The source code of the litmus tests in Power instructions can be found in Appendix L. Fig. 23 shows the results of running the litmus tests with Power instructions on Herd7 using Power's memory model. Both of the executions are allowed under Power's memory model, which confirms the problem we addressed in this paper. The executions becomes forbidden if we change the `lwsync` instruction in the program to `hwsync`.

L Source Code of litmus tests

In this section we provide the source code of the two examples that demonstrate the inconsistency issue we addressed in the paper. In addition, we include the same tests translated to Power instructions.

L.1 volatile-non-sc.4.litmus

```

Java volatile-non-sc.4
{
  x = 0; y = 0;
  0:X=x; 0:Y=y;
  1:X=x; 1:Y=y;
  2:X=x; 2:Y=y;
  3:X=x; 3:Y=y;
}

Thread0 {
  Y.setVolatile(2);
  int r0 = X.getVolatile();
}

Thread1 {
  X.setVolatile(1);
}

Thread2 {
  int r0 = X.getVolatile();
  Y.setVolatile(1);
}

Thread3 {
  int r0 = Y.getVolatile();
  int r1 = Y.getVolatile();
}

exists
(0:r0=0 /\ 2:r0=1 /\ 3:r0=1 /\ 3:r1=2)

```

L.2 volatile-non-sc.5.litmus

```

Java volatile-non-sc.5
{
  x = 0;
  y = 0;
  z = 0;
  0:X=x;0:Y=y;0:Z=z;
  1:X=x;1:Y=y;1:Z=z;
  2:X=x;2:Y=y;2:Z=z;
  3:X=x;3:Y=y;3:Z=z;
  4:X=x;4:Y=y;4:Z=z;
}

Thread0 {
  X.setVolatile(1);
  int r1 = Y.getVolatile();
}

Thread1 {
  Y.setVolatile(1);
}

Thread2 {
  int r1 = Y.getVolatile();
  Z.setVolatile(1);
}

Thread3 {
  Z.setVolatile(2);
  int r1 = X.getVolatile();
}

Thread4 {
  int r1 = Z.getVolatile();
  int r2 = Z.getVolatile();
}

exists
(0:r1 = 0 /\ 2:r1 = 1 /\ 3:r1 = 0
 /\ 4:r1 = 1 /\ 4:r2 = 2)

```

L.3 volatile-non-sc.4.ppc.litmus

```

PPC volatile-non-sc.4.ppc
{
  0:r1=x; 0:r2=y;
  1:r2=y;
  2:r1=x; 2:r2=y;
  3:r1=x;
}

P0      | P1      | P2      | P3      ;
li r3,2 | li r3,1 | li r3,1 | sync    ;
lwsync  | lwsync  | sync    | lwz r3,0(r1) ;
stw r3,0(r1) | stw r3,0(r2) | lwz r4,0(r2) | lwsync  ;
sync    | sync    | sync    | sync    ;
lwz r4,0(r2) |          | stw r3,0(r1) | lwz r4,0(r1) ;
lwsync  |          | sync    | lwsync  ;

exists
(0:r4=0 /\ 2:r4=1 /\ 3:r3=1 /\ 3:r4=2)

```

L.4 volatile-non-sc.5.ppc.litmus

```

PPC volatile-non-sc.5.ppc
{
  0:r1=x; 0:r2=y;
  1:r2=y;
  2:r2=y; 2:r3=z;
  3:r1=x; 3:r3=z;
  4:r3=z;
}

P0      | P1      | P2      | P3      | P4      ;
li r4,1 | li r4,1 | li r4,1 | li r4,2 | sync    ;
lwsync  | lwsync  | sync    | lwsync  | lwz r4, 0(r3);
stw r4,0(r1) | stw r4,0(r2) | lwz r5, 0(r2) | stw r4, 0(r3) | lwsync  ;
sync    | sync    | lwsync  | sync    | sync    ;
lwz r5,0(r2) |          | stw r4,0(r3) | lwz r5, 0(r1) | lwz r5, 0(r3);
lwsync  |          | sync    | lwsync  | lwsync  ;

exists
(0:r5 = 0 /\ 2:r5 = 1 /\ 3:r5 = 0 /\ 4:r4 = 1 /\ 4:r5 = 2)

```

M Full Trace and Litmus Test the example in Section. 2

The litmus test of the example of Fig. 2 is shown below. We labeled each memory instruction (in blue) in the litmus test for better readability of the trace. We obtained the trace by running the ppcm tool by [14] in the online interactive mode.

```

PPC volatile-non-sc.4.ppc
{
0:r1=x; 0:r2=y;
1:r2=y;
2:r1=x; 2:r2=y;
3:r1=x;
}

P0          | P1          | P2          | P3          ;
li r3,2     | li r3,1     | li r3,1     | r: sync     ;
a: lwsync   | f: lwsync   | m: sync     | s: lwz r3,0(r1) ;
b: stw r3,0(r1) | g: stw r3,0(r2) | n: lwz r4,0(r2) | t: sync     ;
c: sync     | h: sync     | o: lwsync   | t16: lwz r4,0(r1) ;
d: lwz r4,0(r2) | p: stw r3,0(r1) | t17: sync   ;
e: sync     | q: sync     |             ;

exists
(0:r4=0 /\ 2:r4=1 /\ 3:r3=1 /\ 3:r4=2)

```

One of the traces to show that this execution is allowed

```

(0:0) Commit reg or branch: li r3,2
(1:6) Commit reg or branch: li r3,1
(2:10) Commit reg or branch: li r3,1
(0:1) Commit barrier: lwsync: a:lwsync
(1:) Barrier propagate to thread: a:lwsync to Thread 1
(2:) Barrier propagate to thread: a:lwsync to Thread 2
(3:) Barrier propagate to thread: a:lwsync to Thread 3
(1:7) Commit barrier: sync: f:Sync
(0:) Barrier propagate to thread: f:Sync to Thread 0
(2:) Barrier propagate to thread: f:Sync to Thread 2
(3:) Barrier propagate to thread: f:Sync to Thread 3
Acknowledge sync: Sync f:Sync
(2:11) Commit barrier: sync: m:Sync
(0:) Barrier propagate to thread: m:Sync to Thread 0
(1:) Barrier propagate to thread: m:Sync to Thread 1
(3:) Barrier propagate to thread: m:Sync to Thread 3
Acknowledge sync: Sync m:Sync
(3:16) Commit barrier: sync: r:Sync
(0:) Barrier propagate to thread: r:Sync to Thread 0
(1:) Barrier propagate to thread: r:Sync to Thread 1
(2:) Barrier propagate to thread: r:Sync to Thread 2
Acknowledge sync: Sync r:Sync

(1:8) Commit write: stw r3,0(r2): g:W y=1 i:W x=0,j:W y=0
Write reaching coherence point: g:W y=1
(2:) Write propagate to thread: g:W y=1 to Thread 2
(2:12) Read from storage subsystem: lwz r4,0(r2) (from g:W y=1)
(2:12) Commit read: lwz r4,0(r2): n:R y=1
(2:13) Commit barrier: lwsync: o:Lwsync
(2:14) Commit write: stw r3,0(r1): p:W x=1 g:W y=1,i:W x=0
Write reaching coherence point: p:W x=1
(3:) Write propagate to thread: g:W y=1 to Thread 3
(3:) Barrier propagate to thread: o:Lwsync to Thread 3
(3:) Write propagate to thread: p:W x=1 to Thread 3
(3:17) Read from storage subsystem: lwz r3,0(r1) (from p:W x=1)
(3:17) Commit read: lwz r3,0(r1): s:R x=1
(0:2) Commit write: stw r3,0(r1): b:W x=2 i:W x=0,j:W y=0
Write reaching coherence point: b:W x=2

```

```

(3:) Write propagate to thread: b:W x=2 to Thread 3
(0:3) Commit barrier: sync: c:Sync
(3:) Barrier propagate to thread: c:Sync to Thread 3
(1:) Barrier propagate to thread: o:Lwsync to Thread 1
(1:) Write propagate to thread: p:W x=1 to Thread 1
(1:) Write propagate to thread: b:W x=2 to Thread 1
(1:) Barrier propagate to thread: c:Sync to Thread 1
(2:) Write propagate to thread: b:W x=2 to Thread 2
(2:) Barrier propagate to thread: c:Sync to Thread 2
Acknowledge sync: Sync c:Sync
(0:4) Read from storage subsystem: lwz r4,0(r2) (from j:W y=0)
(0:4) Commit read: lwz r4,0(r2): d:R y=0
(0:) Write propagate to thread: g:W y=1 to Thread 0
(0:) Barrier propagate to thread: o:Lwsync to Thread 0
(3:18) Commit barrier: sync: t:Sync
(0:) Barrier propagate to thread: t:Sync to Thread 0
(1:) Barrier propagate to thread: t:Sync to Thread 1
(2:) Barrier propagate to thread: t:Sync to Thread 2
Acknowledge sync: Sync t:Sync
(3:19) Read from storage subsystem: lwz r4,0(r1) (from b:W x=2)
(3:19) Commit read: lwz r4,0(r1): t16:R x=2

(0:5) Commit barrier: sync: e:Sync
(1:) Barrier propagate to thread: e:Sync to Thread 1
(2:) Barrier propagate to thread: e:Sync to Thread 2
(3:) Barrier propagate to thread: e:Sync to Thread 3
Acknowledge sync: Sync e:Sync
(1:9) Commit barrier: sync: h:Sync
(0:) Barrier propagate to thread: h:Sync to Thread 0
(2:) Barrier propagate to thread: h:Sync to Thread 2
(3:) Barrier propagate to thread: h:Sync to Thread 3
Acknowledge sync: Sync h:Sync
(2:15) Commit barrier: sync: q:Sync
(0:) Barrier propagate to thread: q:Sync to Thread 0
(1:) Barrier propagate to thread: q:Sync to Thread 1
(3:) Barrier propagate to thread: q:Sync to Thread 3
Acknowledge sync: Sync q:Sync
(3:20) Commit barrier: sync: t17:Sync
(0:) Barrier propagate to thread: t17:Sync to Thread 0
(1:) Barrier propagate to thread: t17:Sync to Thread 1
(2:) Barrier propagate to thread: t17:Sync to Thread 2
Acknowledge sync: Sync t17:Sync

Result:
0:r4=0; 2:r4=1; 3:r3=1; 3:r4=2;

```