

# A Formalization of Java's Concurrent Access Modes

ANONYMOUS AUTHOR(S)

Java's memory model was recently updated and expanded with new access modes. The accompanying documentation for these access modes is intended to make strong guarantees about program behavior that the Java compiler must enforce, yet the documentation is frequently unclear. This makes the intended program behavior ambiguous, impedes discussion of key design decisions, and makes it impossible to prove general properties about the semantics of the access modes.

In this paper we present the first formalization of Java's access modes. We have constructed an axiomatic model for all of the modes using the Herd modeling tool. This allows us to give precise answers to questions about the behavior of example programs, called litmus tests. We have validated our model using a large suite of litmus tests from existing research which helps to shed light on the relationship with other memory models. We have also modeled the semantics in Coq and proven several general theorems including a DRF guarantee, which says that if a program is properly synchronized then it will exhibit sequentially consistent behavior. Finally, we use our model to prove that the unusual design choice of a partial order among writes to the same location is unobservable in any program.

## 1 INTRODUCTION

The original Java memory model [Manson et al. 2005] included an early attempt to define the semantics of lock-free shared memory programs running on the Java platform, but the definitions were hard to understand and there was no easy way to check the behavior of example programs. It was also later discovered that it ruled out existing compiler optimizations which it claimed to support [Ševčík and Aspinall 2008]. Since then, researchers have made great advances in memory model design while studying other weak memory models like those for ARM [Alglave et al. 2008; Pulte et al. 2017], C11 [Batty et al. 2011; Kang et al. 2017; Lahav et al. 2017; Vafeiadis et al. 2015], Power [Alglave et al. 2014], and x86 [Owens et al. 2009].

Recently, the ninth version of the Java Development Kit updated and expanded Java's memory model using new "access modes". Though the design of the access modes is inspired by C11's memory orders [Committee et al. 2010], it differs in a few key ways. First, it sheds complicated legacy features like release sequences and release-consume accesses. Second, it includes a broad but simple mechanism to forbid so called "out of thin-air" behavior [Batty and Sewell 2014]. Finally, it makes no provision for a total order on writes to the same location. Taken together this suggests new opportunities to use a simpler model, develop metatheory, and verify lock-free algorithms for the Java platform.

However, the documentation [JDK9 2017; Lea 2017, 2018] is frequently ambiguous. This makes it extremely difficult to provide definitive answers about program behavior and there is little hope of proving important properties about the semantics. Further, it impedes the discussion of key features of the model's design.

To address these issues, we present the first formalization of Java's access modes. Critically, our model is *precise* and *complete*, formalizing all of the main features of Java's access modes. Additionally, we have endeavoured to make the model as *readable* as possible, defining all modes and fences as small extensions to an intuitive notion of visibility. Specifically, we make the following contributions:

- an axiomatic model for all of Java's access modes, fences, and atomic read-writes;

- an instantiation of our model for the Herd tool, which allowed us to examine the outcomes for more than 80 test programs; and
- an instantiation of our model in Coq, which allowed us to prove three key theorems about the semantics that are novel for Java’s access modes.

We have constructed our model using the cat language. This allowed us to leverage the Herd tool [Alglave et al. 2014] to give definitive answers to questions about example programs, called litmus tests. Litmus tests are designed to highlight specific behaviors in memory models. Herd enumerates all possible executions for a litmus test and determines which are allowed according to the model. Then, if at least one execution is allowed, the behavior illustrated by the litmus test is allowed by the model.

We used Herd with more than 80 litmus tests drawn from prior research to help validate our model by comparing it to ARMv8, C11 and x86. Our goal for these comparisons is to show that there are no unexpected differences in behavior. For example, the conventional wisdom is that language memory models will exhibit more behaviors than architecture memory models due to aggressive compiler optimizations. Thus, Java’s access modes should permit more behaviors than ARMv8. If that is not the case then it should be due to a deliberate design decision and not a bug in the definitions. For all 80 litmus tests our model behaves according to our expectations and the documented design of the access modes.

We have also formalized our model in Coq and we prove three key theorems: absence of causal cycles when all reads are “release” reads, sequentially consistent semantics under proper synchronization (DRF), and a guarantee that each stronger mode admits fewer executions [Vafeiadis et al. 2015]. These theorems further validate the definitions of our model, give more evidence that the model is complete with respect to the documentation, and clearly demonstrate that the semantics is suitable for formal reasoning.

Finally, the partial order on same-location writes in Java’s access modes represents a significant departure from the conventions of existing memory models. We show that the impact of switching to a partial order is “unobservable” in any example program executed using our model and thereby show that Java’s access modes can adopt a total order on writes to the same location. The simplicity of our model shines through in our proof, which makes it clear that our reasoning is not applicable to the more complex axiomatic models of RC11 and ARMv8.

The rest of this paper is laid out as follows. In Section 2 we use a simple example to illustrate the ambiguity of the documentation. In Section 3 we detail the features which distinguish Java’s access modes from other memory models. In Section 4 we give a complete account of our axiomatic model. In Section 5 we detail the results of our comparisons with ARMv8, C11 and x86. In Section 6 we give a formalization of the model in Coq and use it to prove the theorems listed above. In Section 7 we demonstrate that the choice of a partial coherence order is not observable in our model. In Section 8 we discuss related work.

## 2 JAVA’S ACCESS MODES

Here we will introduce the basics of the Java access mode API and motivate our formalization by way of an example program. We will show how different interpretations of the documentation can cause one of three unwanted outcomes when using the access mode API.

Developers can make use of Java’s Access Modes, hereafter “the JAM”, through the `VarHandle` API included in the JDK version 9. There are four access modes: plain, opaque, release-acquire, and volatile. Regular reads and writes of shared variables are considered plain mode, and the `VarHandle` API allows reads and writes to be annotated with one of the other three modes. The specified intent of the JAM is that each mode provides progressively more guarantees about the behavior of its

99 accesses at the expense of some performance.

100 plain  $\sqsubseteq$  opaque  $\sqsubseteq$  release-acquire  $\sqsubseteq$  volatile

101  
102 Plain mode gives virtually no guarantees and the compiler is allowed free reign in optimizing such  
103 memory accesses. On the other hand volatile mode provides sequentially consistent (SC) semantics  
104 when it is used for all accesses, but requires a memory barrier for each access which can make  
105 execution slow <sup>1</sup>.

106 The VarHandle documentation [JDK9 2017; Lea 2017, 2018] is intended to answer questions  
107 about which access mode should be used for a read or write to maximize performance while  
108 providing enough guarantees to ensure program correctness. It cites three bad outcomes arising  
109 from the confusion around the old Java memory model as motivation for the new access modes:  
110 undersynchronized and broken code, oversynchronized and slow code, and platform specific code.  
111 However, the documentation is written in loose prose which can result in any one of the same  
112 three outcomes which the access modes were created to address.

113 To see this, consider the example execution of a message  
114 passing program in Figure 1. An execution graph like this  
115 catalogues the relationships and effects of each memory access  
116 from an example program.  $M(x, n)$  represents a memory  
117 access of  $x$  with the value  $n$  where  $M$  can be write,  $W$ , a read,  
118  $R$  and later an atomic read-write  $RW$ . Relationships between  
119 memory access are represented as directed edges. The dashed  
120 edge labeled with  $rf$  represent reads-from and the edges labeled with  $po$  represent the sequence of  
121 instructions as defined in the program. Note that  $a$  and  $b$  are concurrent with  $c$  and  $d$ .

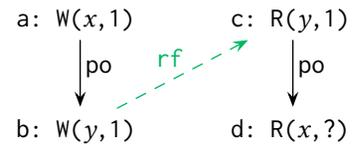


Fig. 1. Event Graph

122 Intuitively, if the read of  $y$  produces the value 1, it acts as a signal to the rest of the second thread  
123 that the write to  $x$  has completed. However, in the presence of weak memory accesses the sequence  
124 of execution may not follow program order. Thus, an important question is: which access modes  
125 should one use so that  $d$  reads the value 1 from  $a$  in keeping with program order, and not 0 from  
126 the initialization of  $x$ ? Depending on a reasonable interpretation of the documentation, our answer  
127 to this question may result in one of the aforementioned bad outcomes.

128 **Undersynchronized, broken code.** Consulting the documentation [Lea 2018] we find that  
129 opaque mode makes the following guarantee:

130 If Opaque (or any stronger) mode is used for all accesses to a variable, updates do not  
131 appear out of order.  
132

133 Whether opaque mode is the correct mode for each access to ensure that  $d$  reads 1, depends  
134 heavily on the definition of “out of order”. Intuitively, we assume that  $rf$  implies that the read has  
135 observed the effects of the write. Then, if the “order” from the documentation for opaque mode  
136 is program order, we might reasonably conclude that updates (writes) are observed in program  
137 order. In which case,  $d$  observes the effects of  $c$  and  $b$  observes the effects of  $a$ . Then, through the  
138  $rf$  edge we can conclude that  $d$  observes the effects of  $a$  and  $d$  will read 1.

139 Unfortunately this approach actually leaves the code undersynchronized and it will not work as  
140 described. It turns out, the documentation also says that opaque mode makes no guarantees about  
141 information that travels through reads:

142 ... reading a value in Opaque mode need not tell you anything about values of any  
143 other variables.  
144

145 <sup>1</sup>This property, of progressively greater guarantees, is called monotonicity and we prove it as a theorem for our model in  
146 Section 6

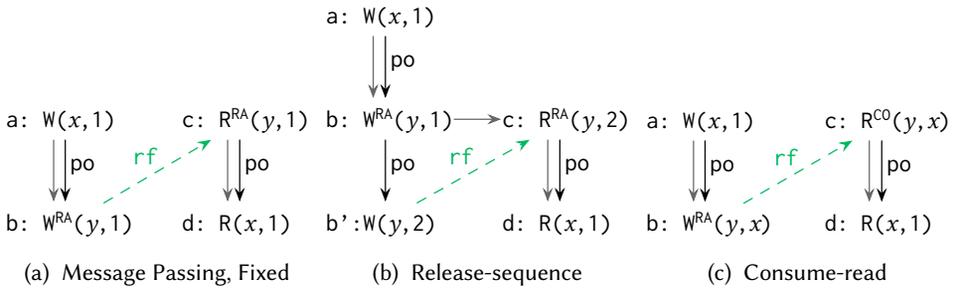


Fig. 2. Release sequences and Consume-reads

Thus, whether or not writes are observed in program order, the read  $c$  may not tell us anything about what happened with  $a$  opening the way for  $d$  to read 0 from the initial write to  $x$ . These two pieces of documentation can lead to an inconsistent understanding of opaque mode. A reasonable reader may react defensively and look to the stronger modes to guarantee the correct program behavior.

**Oversynchronized, slow code.** A defensive approach to getting the right program behavior is to employ volatile mode. This would give all of the reads and writes in the program SC semantics, guaranteeing that everything is observed in program order. Thus, the  $po$  and  $rf$  edges would indeed result in  $d$  observing  $a$ , and  $d$  will read 1. However, if the compiler takes the approach outlined in the documentation it will insert a full fence after every SC access which will slow down execution.

**Platform specific implementation.** Alternately, we may note that x86 processors do not reorder writes with writes or reads with reads. Assuming opaque mode accesses will translate directly to hardware instructions, we can tag each access as opaque and use same reasoning that we used for volatile mode. Unfortunately this reasoning is unsound for execution on ARM processors, which allow such reorderings, and the code would not be portable.

As it happens, the best solution is to use release-write for  $b$  and acquire-read for  $c$ . These modifications appear as the RA annotations for the write and read of  $y$  in Figure 2a. This will ensure the orderings, depicted as gray edges, for any access before  $b$  and any access after  $c$  with a minimum of synchronization. Thus,  $a$  is observed by  $d$  through those orderings and  $d$  will read the value 1. Critically, we can only demonstrate this after more precisely defining the guarantees provided by the opaque and release-acquire modes of the JAM. Our formalization brings this much needed clarity to the current specification.

### 3 DISTINGUISHING FEATURES OF THE JAM

While the JAM was inspired by the access modes of C11, it makes several departures from C11 and other memory models that are worthy of consideration and a challenge for formalization. First, it sheds legacy features like the release sequences and consume-reads of C11. Second, it contains a broad and simple definition of causal cycles for the purposes of ruling out thin-air reads. Finally it includes a non-total ordering on writes to a particular memory address, which is called the coherence order. Here we will discuss how each of these differences will impact any formalization effort for the JAM.

#### 3.1 Letting Go of Release Sequences

Release sequences and consume-reads can be seen as specialized variants of the release-acquire memory order in C11 and release-acquire mode in the JAM. The idea is that, in certain cases, it's possible to get the same guarantees as a release-write and acquire-read pair, like the one Figure 2a, with less synchronization. The result is faster execution in some contexts.

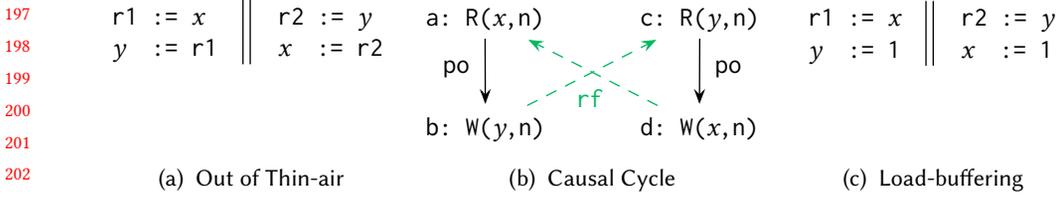


Fig. 3. Causal Cycles

Release sequences can be used when a write that is after a release-write is read by an acquire-read. Consider the message passing variant in 2b where b is followed by another distinct write to  $y$ ,  $b'$ . If c reads from  $b'$  then the release sequences of C11 would guarantee that d would see a through the orderings depicted with gray edges. Then d will read the value 1, just as it would if c read from b. Without the guarantees of release sequences,  $b'$  must also be a release-write to get the desired outcome.

Consume-reads are used with a release-write when the memory accesses which must be ordered after the consume-read are data dependent on the value of the read. Consider the message passing variant in 2c where c is annotated with CO and reads a pointer that determines the memory address that d reads from. Some architectures enforce the ordering of c and d in the presence of such data dependencies without the extra synchronization that would result from making c an acquire-read. This can speed up execution in those settings.

The JAM does not include the guarantees of release sequences or any way to annotate a read as a consume-read. This is a design choice in favor of simplicity in the model and it gives us the opportunity to build a more readable and more easily testable model.

### 3.2 Acyclic Causality

In Figure 3a we have a classic example program which can exhibit a so called “thin-air” read under sufficiently weak memory models. The question is, at the end of execution can  $x = y = 42$ ? Intuitively, assuming both  $x$  and  $y$  are initialized to 0, this program can’t generate 42 “out of thin-air”, but many axiomatic models do not exclude such executions from the set of all candidate executions.

Observe that, in any execution that allows  $x = y$ , there must be a cycle in the program order and reads-from relations, as illustrated in Figure 3b. The JAM explicitly forbids such cycles, but at the cost of forbidding some behaviors which may be beneficial for performance.

For example, consider the classic load-buffering (LB) litmus test in Figure 3c, where the question is, can  $r1 = r2 = 1$ ? If performance was the sole concern in the design of the JAM this behavior would be allowed because the reads can be reordered with the writes with the aim of improving performance. Unfortunately, this example also exhibits the same cycle in the program order and reads-from, so it is forbidden by the JAM.

The problem of differentiating these kinds of examples has been studied at great length by memory model researchers. Another approach is to forbid cycles in rf and a subset of program order based on a notion of dependency. Sadly, this too has very subtle issues, as outlined by Batty and Sewell [2014]. The original, formal Java Memory Model [Manson et al. 2005] attempted to address the issue of causal cycles in its full generality. More recently the Promising semantics of [Kang et al. 2017] introduced a novel “promise” mechanism to model compiler optimizations for this purpose.

In all of these cases the complexity of the resulting models makes them hard to understand and hard to test. Instead, the JAM specification adopts a simple solution by forbidding cycles in the program order and reads-from. While this does forbid the behavior of the second example at the

246 cost of some performance [Ou and Demsky 2018], it gives us yet another opportunity to build a  
 247 simpler model.

### 248 3.3 Partial Coherence Order

249 The JAM specification makes no provision for a total ordering of writes to a given memory location  
 250 which is a standard feature in other memory models. The consequences of this design choice  
 251 manifest in subtle ways.

252 For example, the standard definition of atomicity for read-writes relies on a total coherence order.  
 253 Borrowing the definition from Vafeiadis et al. [2015], for a read-write, RW, the write it reads-from,  
 254  $W_1$ , and the coherence order,  $\xrightarrow{\text{co}}$ , we have:

$$255 W_1 \xrightarrow{\text{rf}} RW \implies \neg \exists W_2, W_1 \xrightarrow{\text{co}} W_2 \xrightarrow{\text{co}} RW$$

256 Taken together with a total coherence order this means that the atomic pairs of writes and read-  
 257 writes are always totally ordered. Since the JAM makes no such guarantees regarding normal writes  
 258 there are ambiguities like the one in Figure 4. There, the write-read-write pairs are not ordered  
 259 across threads since there is no global relationship between writes.

260 As a consequence, our model must include extra constraints  
 261 for read-writes while being careful not to over-constrain  
 262 normal writes which would otherwise be concurrent. We  
 263 detail our approach in Section 4.5.

264 Separately, our drive for simplicity in the definitions of  
 265 the model has yielded a key insight where the coherence  
 266 order is concerned. While we have remained faithful to the  
 267 documentation and personal communications with the au-  
 268 thors in modeling a partial coherence order, we will show  
 269 that the effects of that choice are not observable under our  
 270 model. Specifically, we show that it is impossible to construct  
 271 a litmus test that behaves differently in the presence or absence of a total coherence order.

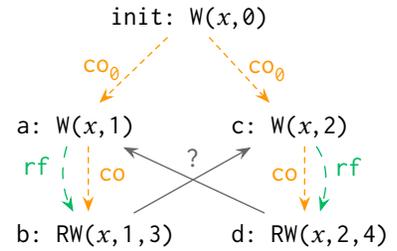


Fig. 4. Concurrent Read-writes

## 274 4 AXIOMATIC MODEL

275 The JAM has six components: plain mode, opaque mode, release-acquire mode, volatile mode,  
 276 fences, and atomic read-writes. Each of the modes, from plain to volatile, provides strictly more  
 277 guarantees than the previous mode.

278 To model the JAM we have constructed an axiomatic semantics. The definitions of our model act  
 279 as a predicate over candidate executions which include memory events, e.g. reads and writes, and  
 280 relations over those events, e.g. reads-from and program order, in the style of Alglave et al. [2014].

281 Our definitions focus on two key concepts. The first is an acyclicity requirement for the coherence  
 282 order. Aside from the restriction of causal cycles, this is the only mechanism by which executions  
 283 are forbidden. Thus, every unwanted execution must exhibit a cycle in the coherence order. The  
 284 second is an intuitive notion of *visibility*, which represents when one memory access has "seen"  
 285 the effects of another memory access. The behavior of the three modes and fences are modeled as  
 286 small extensions to this relationship.

287 In each of the following subsections we detail the extensions to our model for each component.  
 288 The full model can be viewed in Appendix A.

### 290 4.1 Plain & Opaque Mode

291 In the JAM documentation, plain mode accesses are given virtually no guarantees when they occur  
 292 in different threads without correct synchronization. Opaque mode, on the other hand, does provide  
 293 some cross thread guarantees which form the basis for the rest of the memory model. There are  
 294

some subtleties involved in the documented relationship between plain and opaque mode accesses so we will address them together. First, the main properties that opaque mode accesses guarantee are:

**Bitwise Atomicity** Reads will see the value of only one write. Opaque mode guarantees that reads will not see mixed bits from different writes.

**Write Availability** Writes can be read by later reads. The intent is to avoid a situation (e.g. in a spin-loop) where repeated reads never see a write in another thread because they are optimized by the compiler to execute only one time. When an optimization like this happens, the availability of the write for the read depends on when the read is executed [Corbet 2012].

**Acyclic Causality** A read should not influence its own value. As described in Section 3.2, this forbids counter-intuitive behavior like thin-air reads.

**Coherence** The order of writes should respect visibility and should agree with the way that reads observe their order. For example, one guarantee (of four we will define later) is that a read should be paired with the last write that it knows about and not an earlier one.

Our model of opaque mode focuses on acyclic causality and coherence. In keeping with other axiomatic models built for Herd [Alglave et al. 2008; Lahav et al. 2017; Pulte et al. 2017], our model pairs each read with a single write and there is no accounting for optimizing reads out of loops as Herd does not support them.

**Plain Coherence.** The JAM documentation does not include a coherence guarantee for plain mode accesses. This follows the approach in the C11 documentation but it departs from formal models for C11.

To see why plain mode accesses should be included in the coherence ordering guarantees, note that plain mode accesses are safe to use within a critical section guarded by a lock according to the documentation. The idea is that code which is properly synchronized with locks will have single threaded semantics for the duration of the critical section. Thus, the model shouldn't require accesses to be annotated with anything stronger than plain mode.

Then, consider the message passing variant in Figure 5. Here, the purpose of this pattern is to signal the end of the critical section through atomic read-writes that unlock, b, and lock, c, the variable  $y$ . For any lock to work correctly, the accesses which are program order before the unlock,  $a \xrightarrow{po} b$ , must be visible to accesses after the lock,  $c \xrightarrow{po} d$  (for now we leave the mechanism that enforces these orderings unspecified). In particular the effects of the write a should be visible to the write d,  $a \xrightarrow{co?} d$ .

However, even if the unlock and lock enforce program order of the plain writes, and thereby show that a is visible to d, we would not be able to derive  $a \xrightarrow{co} d$  because the coherence rules do not apply to plain accesses. This stands in contrast to formal models of C11 [Lahav et al. 2017] where the coherence order and happens-before relations apply to plain accesses. As a result, our model extends the coherence order guarantees to plain mode accesses and only the extra guarantee of acyclic causality is left to opaque mode.

**Herd Model.** Figure 6 details our axiomatic model of the JAM's plain and opaque modes as defined in the cat modeling language. cat includes the following built in operations:  $|$ ,  $\&$ ,  $;$ ,  $+$ ,  $\sim$  are relational union, intersection, composition, transitive closure, and complement. New constructs are defined with let. Filters, like `wco`, are defined on relations using `let F(R) = ...` and applied with `F(R)`. Finally, models can include checks like `acyclic` for relations.

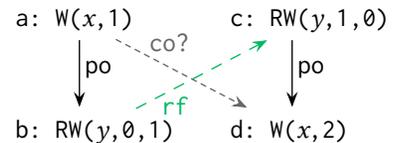


Fig. 5. Message Passing Coherence

344 Herd provides several built-in, ambient sets and relations  
 345 for models defined in `cat`. `IW`, `W`, `FW`, `R`, `O`, `RA`, `V` are the sets of  
 346 initial writes for each location, plain mode writes, the final  
 347 plain mode writes to each location, plain mode reads, opaque  
 348 accesses, release-acquire accesses, and volatile accesses. `po`,  
 349 `po-loc`, `rf`, and `invf` represent the program order, program  
 350 order per-location, reads-from, and inverted reads-from rela-  
 351 tions. `loc` relates memory accesses to the same location and  
 352 `id` is the identity relation. Note, that we do not use Herd's  
 353 built-in coherence relation but rather define our own without  
 354 a total ordering.

355 We define visibility order, `vo`, using two properties, `po-loc`  
 356 and `rf+`. In the first case, given two memory accesses to the  
 357 same location in the same thread the second should see the  
 358 effects of the first. In the second case, the sequence of one  
 359 or more reads guarantees that the final read is executing  
 360 in a context where the effects of the write are visible. This  
 361 definition for `vo` guarantees only the most basic forms of  
 362 ordering which is important given the inclusion of plain  
 363 mode accesses.

364 Then, to satisfy the coherence requirement we first ensure  
 365 that the coherence order includes only distinct writes to the same location, `wwco`, and adopt the  
 366 four behaviors originally identified for C11 by Batty et al. [2011]:

- 367 **coww** Two such writes ordered by visibility are similarly coherence ordered.
- 368 **cowr** A read chooses the last write it has seen.
- 369 **corw** A write that follows a read of the same location in program order is coherence ordered  
 370 after the write paired with the read.
- 371 **corr** Program ordered reads to the same location order their writes in the same way.

372 We also order the initialization write before all writes of the same location, `coinit`, and we  
 373 order all writes before the final write of the same location, `cofw`. In the rest of the paper these  
 374 relationships will be distinguished as `co0` for clarity but they are treated the same as any other `co`  
 375 edge by the model.

376 These six rules make up the definition of the coherence order, `co`, until we discuss read-writes.  
 377 Then, the primary mechanism by which the model forbids executions is through requiring the  
 378 coherence order to be acyclic, `acyclic co`. In Figure 7, we give examples of how each coherence  
 379 rule forbids an execution by showing a cycle in the coherence order. In every example, the final  
 380 write of 1 is assumed to be coherence order after all other writes, `co0`.

381 In Figure 7a, if we the write, `a`, is visible write, `b`, to the same location then `a` is not the last write.  
 382 In Figure 7b, if the read `c` has seen the last write `b` then it should not be able to read an older write  
 383 `a`. In Figure 7c, if we can follow the read of `a` to `b` to a later write, `c`, then `a` can't be the last write.  
 384 In Figure 7d, given two reads in program order, `b` and `c`, if `b` has seen the last write then `c` should  
 385 not be able to read an older write.

386 Finally, we must forbid causal cycles for opaque mode access. We define `opq` to be any opaque or  
 387 stronger access since any guarantee provided by opaque should apply to stronger access modes.  
 388 Then, we forbid cycles in the union of program order and reads-from, qualified for opaque mode  
 389 accesses: `acyclic (po | rf) & opq`.

392

```

let vo = rf+ | po-loc
let wwco(rel) = rel & ~id
                & loc & (W * W)

let coww = wwco(vo)
let cowr = wwco(vo; invrf)
let corw = wwco(vo; po-loc)
let corr = wwco(rf; po-loc; invrf)

let coinit = wwco((IW * W))
let cofw = wwco((W * FW))

let co = coww | cowr | corw
        | corr | cofw | coinit

acyclic co

let opq = O | RA | V
acyclic (po | rf) & opq
  
```

Fig. 6. Opaque Mode

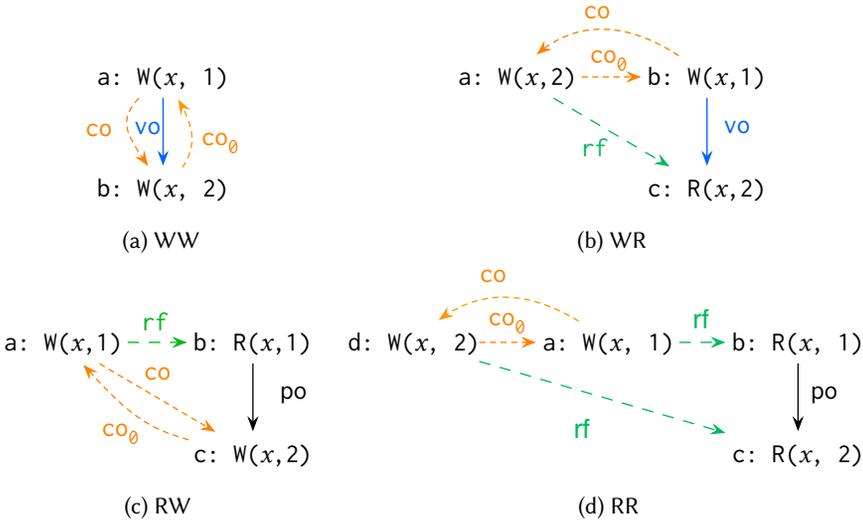


Fig. 7. Coherence

As we will see in the following sections, these base definitions allow us to model nearly every other component by extending `vo`. The lone exception is atomic read-writes, for which we extend the coherence order directly.

## 4.2 Fences

The JAM supports five types of fences: release, acquire, load-load, store-store, and full. Programs often include fences to enforce the order of memory access effects before and after the fence. We update `vo` by extending visibility to `rfso+` and abstracting over these fence types with *specified orders* [Bender et al. 2015; Cray and Sullivan 2015]. We therefore require that programs relate the ordering of pairs of memory accesses using either `svo` or `spush` in place of fences.

```

421 with to from linearisations( $\mathbb{M} \sqcup \text{IW}$ , cofw | rf | into)
422 let spushto = to+ & (domain(spsh) * domain(spsh))
423 let rfso    = rf | svo | spush | spushto; spush
424 let vo      = rfso+ | po-loc

```

The effects of intra-thread ordering fences (release, acquire, load-load, and store-store) are modeled with specified visibility orders, `svo`. Full fences, are modeled as specified push orders, `spush` and `spushto`, in the style of Cray and Sullivan [2015].

This approach allows us to extend visibility to account for synchronization in a way that is uniform for fences and their related access modes. Moreover, the abstraction allows the compiler to make smart choices to ignore specified visibility based on existing intra-thread ordering of instructions guaranteed by a target architecture as discussed in the work of Bender et al. [2015]; Sullivan [2015]. We will see an example of this when discussing atomic read-writes.

Note that trace order, `to`, is a total ordering of all memory accesses (except for initial writes) as constructed by Herd's built-in `linearisations` function. Trace order respects `cofw`, `rf`, and all intra-thread orderings, `into`, induced by specified visibility orders, push orders and later release-acquire memory accesses and volatile memory accesses. Intuitively, each of these synchronization mechanisms guarantees that the related memory accesses are executed in program order.

**Specified Visibility Orders.** To see how specified visibility orders create intra-thread visibility between accesses, consider the event graph in Figure 8 for the message passing example from Section 2.

442 By specifying that  $a \xrightarrow{\text{svo}} b$  and  $c \xrightarrow{\text{svo}} d$  we can show that  $d$  could not have read 0 from the  
 443 initial write to  $x$ . First, we assume that the initial write to  $x$  is coherence order before  $a$ ,  $\text{init} \xrightarrow{\text{co}_0} a$ .  
 444 Then, assume that  $d$  has read from the initialization,  $R(x, 0)$ . We will show a contradiction. In the  
 445 graph we have that,  $a \xrightarrow{\text{svo}} b \xrightarrow{\text{rf}} c \xrightarrow{\text{svo}} d$ . By the definition of  $\text{rfso}$  we have the three edges  
 446 that combine to show  $a \xrightarrow{\text{vo}} d$ . Then, by  $\text{cowr}$  (Fig. 7b) we have that  $a \xrightarrow{\text{co}} \text{init}$  which is a cycle  
 447 in  $\text{co}$  and a contradiction.

448 **Specified Push Orders.** Specified push orders create visibility  
 449 relationships in two ways. The first is an intra-thread  
 450 visibility ordering between the two push ordered instructions  
 451 like  $\text{svo}$ . This appears as the  $\text{spush}$  in the definition of  $\text{rfso}$ .

452 The second, is a cross-thread ordering that emulates the  
 453 standard total ordering of two full fences. Given two push  
 454 orders, the head of the first will be visible to tail of the second,  
 455 or vice-versa based on the current ordering of the heads in  
 456 to. The ordering of the heads is recorded as  $\text{spushto}$  which  
 457 we connect with the additional visibility ordering of one tail  
 458 using,  $\text{spushto}; \text{spush}$ .

459 To see how push orders emulate full fences, consider the  
 460 store-buffering litmus test in Figure 9. The question is, can  
 461 both reads take their values from the initialization?

462 First, we assume the gray edge between the two fence  
 463 instructions,  $\text{fence} \rightarrow \text{fence}$ , which represents one side of  
 464 the total ordering provided by full fences. Then, we assume  
 465 that  $d$  reads from the initial write to  $x$  and show a contra-  
 466 diction. Since fences also provide the intra-thread orderings  
 467  $a \rightarrow \text{fence}$  and  $\text{fence} \rightarrow d$ , we can see that  $a$  is ordered  
 468 before  $d$ . Then by  $\text{cowr}$  it must be that  $a \xrightarrow{\text{co}} \text{init}$ , which  
 469 creates a cycle with  $\text{init} \xrightarrow{\text{co}_0} a$  and a contradiction. On the  
 470 other side of the ordering between fences a similar argument  
 471 applies if  $b$  reads from initial write to  $y$ . Thus,  $b$  and  $c$  could  
 472 not have read from the initial writes to  $x$  and  $y$  in the same  
 473 execution.

474 Push orders are more direct but capture the same ordering.  
 475 In Figure 10 both write-read pairs are push ordered.

476 First we assume one side of the total trace order between  
 477 the two writes,  $a \xrightarrow{\text{to}} c$ . Again, we assume that  $d$  reads  
 478 from the initial write to  $x$  and show a contradiction. By  
 479  $\text{spushto}; \text{spush}$ , we have  $a \xrightarrow{\text{vo}} d$  which means that  $d$  has  
 480 seen  $a$  and the same reasoning with  $\text{cowr}$  that we used with  
 481 the fences applies.

482 Note that including  $\text{spushto}$  in  $\text{vo}$  in place of  $\text{spushto}; \text{spush}$  would give the desired visibility  
 483 between the heads and tails of two orders by transitivity, but it would also give an ordering to the  
 484 heads of the push orders which does not otherwise exist.

485 Specified orders have a natural implementation as fences which has been studied by Bender et al.  
 486 [2015] and Sullivan [2015]. We give such a mapping in Appendix B.

487  
 488  
 489  
 490

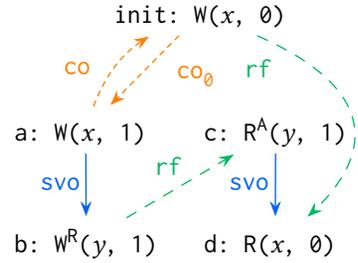


Fig. 8. Specified Visibility




Fig. 9. Fences, One side

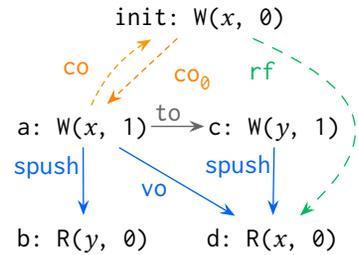


Fig. 10. Push order, One side

### 4.3 Release-Acquire Mode

We can now make small extensions to incorporate other access modes into our model. First, we define release-acquire mode following the standard set by other models like C11 [Committee et al. 2010] and ARMv8 [Pulte et al. 2017].

```

495 let rel = W & (RA | V)
496 let acq = R & (RA | V)
497 let ra = po;[rel] | [acq];po | rfso
498 let vo = ra+ | po-loc

```

We define release writes,  $rel$ , to be any write marked as release or volatile. We define acquire reads,  $acq$  to be any read marked as acquire or volatile. Again, any guarantee provided by release-acquire mode should hold for volatile mode.

We update the  $vo$  definition from opaque mode with fences by extending what was previously  $rfso+$  to be  $ra$ . We add edges from memory accesses for any location to a release write that is later in program order,  $po;[rel]$ . We also add edges from an acquire read to program order later opaque memory accesses for any location,  $[acq];po$ . Note that the documentation makes clear that plain accesses should be ordered by release writes and acquire reads so long as the types are bitwise atomic.

To see, how the release-acquire extension to opaque mode works, consider the message passing example from the Section 2. Note that, we use superscript  $RA$  for release writes,  $W^{RA}$ , and acquire reads,  $R^{RA}$ . Later we will use  $V$  for volatile accesses. Then, if we adopt a release write for  $b$  and an acquire read for  $c$  we can show that,  $a \xrightarrow{vo} d$ , and the reasoning is the same as for specified visibility orders in Section 4.2.

We note that our definitions suggest that, if all reads were release reads, we could derive a visibility relationship between a read and itself in executions exhibiting causal cycles. In section 6, we prove that visibility cycles are a contradiction in our model and, as such, causal cycles are forbidden when all reads are release-reads.

### 4.4 Volatile Mode

Volatile mode is a further extension of the visibility order in release-acquire mode. We update  $vo$  by extending  $ra$  to  $vvo$ . Here  $volint$  creates edges from any access to a volatile read that is later in program order,  $po;[V \& R]$ , and from an volatile write to any program order later opaque memory access,  $[V \& W];po$ .

```

528 let vol = V
529 let volint = po;[vol & R] | [vol & W];po
530 let push = spush | volint
531 let vvo = ra | pushto;push | push
532 let vo = vvo+ | po-loc

```

These new edges preserve program order for any access before or after the volatile access when combined with the visibility definition of release-acquire. We also extend  $spush$  with  $volint$  edges so that we can leverage the same visibility relationships induced by push orders for volatile accesses.

A simpler approach would be to translate the total trace ordering into visibility edges. That is, we could replace the definition of  $vol$  above with the following:

```

538 let vol = ra | spush | volint | to & (V * V)

```

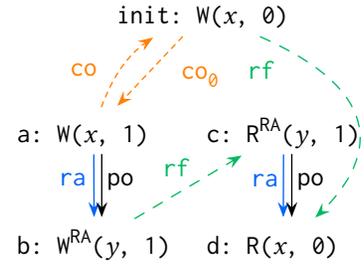


Fig. 11. Release-Acquire Visibility

540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588

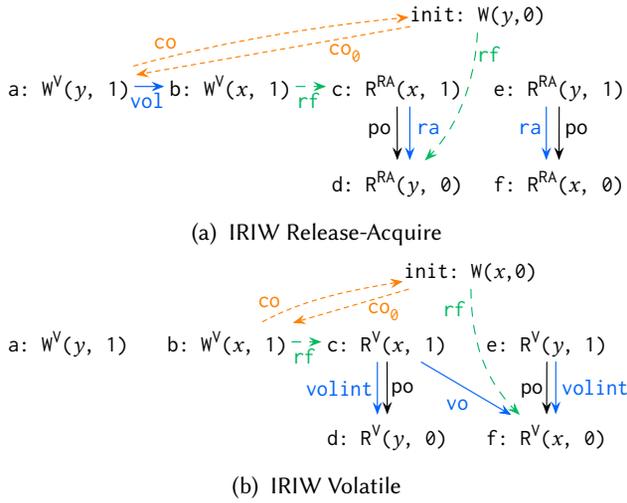


Fig. 12. IRIW Variants

This approach more directly encodes the cross thread ordering guarantee of sequential consistency, but, unfortunately, this is too strong.

Consider a modified version of the classic IRIW litmus test in Figure 12a. The question is, can both  $d$  and  $f$  read from the initial write to  $x$  and  $y$ ? As outlined in the JAM documentation the acquire reads in this example are allowed to see the writes in different orders, so both can read from the initial writes for  $x$  and  $y$ . However, if volatile mode accesses are totally ordered as in the proposed definition, then we have either  $a \xrightarrow{vo} b$  or  $b \xrightarrow{vo} a$ . In Figure 12a we have the first case. Then we have  $a \xrightarrow{vo} d$  because  $a \xrightarrow{vo1} b \xrightarrow{rf} c \xrightarrow{ra} d$ . Then by cowr, we have that  $a \xrightarrow{co} \text{init}$ , which is a contradiction. In the other case,  $b \xrightarrow{vo} a$ , we have a contradiction when  $f$  reads from the initial write to  $y$ .

Instead, we extend our notion of push orders and define push to include both `push` and `volint` edges. This has the same effect as push ordering accesses related by `volint`. As we will demonstrate later, using a matching litmus tests, this guarantees the correct behavior for the release-acquire variant of IRIW because it enforces no direct ordering between the two writes.

Importantly, it also gives the correct behavior when the reads are volatile, which should be SC semantics. Consider Figure 12b which has one of the two possible orderings given by `pushto`; push when the reads are volatile. This correctly establishes  $c \xrightarrow{vo} f$  creating a contradiction for the opposite thread's final read. As before the other direction of the total order forbids the other read.

#### 4.5 Atomic Read-Writes

The behavior of atomic read-writes is the only part of the JAM that is not modeled by extending `vo`. Recall Figure 4 from Section 3. We must take care to ensure that we do not allow concurrent read-writes in the presence of a non-total coherence order. To achieve this we update the coherence order `co` with two additional rules:

```

let corwexcl = wwco((rf; [RW])^-1; co')
let corwtotal = wwco(((RW * W) | (W * RW)) & to)
let rec co = ... | crwexcl | corwtotal

```

The first, `corwexcl` ensures exclusivity in the relationship between the read-write and its paired write. Note that, for clarity, we separate out `corwexcl` even though it recursively refers to `co` in its definition. As illustrated in Figure 13, if there is a write `co`-after the one paired with the read-write,  $a \xrightarrow{co} b$  then it is `co`-after the read-write,  $c \xrightarrow{co} b$ .

589 Recalling the example of Figure 4, this exclusivity is not  
 590 enough to prevent concurrent read-read-write pairs to the  
 591 same location. Since exclusivity uses the ordering of other  
 592 writes with the paired write we could consider a total ordering  
 593 just for paired writes.

```
594 let pairedw = domain(rf; [RW])
595 let corwtotal = wwco((pairedw * pairedw) & to)
```

596 This suffices to forbid concurrent read-write chains to the same location. In Figure 4, by  
 597 corwtotal we have  $a \xrightarrow{\text{co}} c$  or  $c \xrightarrow{\text{co}} a$ . Then, in the first case, we have  $b \xrightarrow{\text{co}} c$  by corwexcl. The  
 598 other side is similar. However, this approach inadvertently orders writes that could be concurrent  
 599 under a partial coherence order. Instead, as in the first definition, we choose a total ordering with  
 600 the read-write itself.

601 Now any chain of read-writes will be exclusive without unnecessarily ordering regular writes.  
 602 We will show that in Figure 4 either,  $b \xrightarrow{\text{co}} c$  or  $d \xrightarrow{\text{co}} a$ . By corwtotal we have  $d \xrightarrow{\text{co}} a$  or  
 603  $a \xrightarrow{\text{co}} d$ . In the first case we are done. In the second case, we consult corwtotal again and we  
 604 have  $b \xrightarrow{\text{co}} c$  or  $c \xrightarrow{\text{co}} b$ . In the first case we are done. In the second case, by assumption we have  
 605  $a \xrightarrow{\text{co}} d$  and  $c \xrightarrow{\text{co}} b$ . Then we can apply corwexcl to both to derive  $c \xrightarrow{\text{co}} d$  and  $d \xrightarrow{\text{co}} c$ , which  
 606 is cycle and a contradiction. Importantly this reasoning can be applied repeatedly to any chain of  
 607 read writes to achieve exclusivity.

608 We note here, that JAM makes no intra-thread ordering guarantees for atomic read-writes even  
 609 though they exist on some architectures like x86 [Owens et al. 2009]. We discuss this in more detail  
 610 in Appendix B.

## 612 4.6 Summary

613 Our axiomatic model of the JAM is complete, covering all four modes, atomic read-writes, and  
 614 all five fence types in less than 40 lines of definitions. Moreover each component is implemented  
 615 as a modest extension to just two relations which makes it readable. Now we must demonstrate  
 616 that the model is consistent with expectations about the behavior of the JAM as outlined in the  
 617 documentation. We discuss the results of our comparison with ARMv8, RC11 and x86 in the next  
 618 section.

## 620 5 VALIDATION

621 In this section we validate our formalization of the JAM by comparing litmus test outcomes for the  
 622 JAM with the outcomes for ARMv8 [Pulte et al. 2017], RC11 [Lahav et al. 2017] and the model for  
 623 x86 included with Herd. We use these comparisons to show that there are no unexpected differences  
 624 in behavior between each pair of models. For example, we expect the JAM to permit more behaviors  
 625 than ARMv8 with the exception of litmus tests like load-buffering (Fig. 3c) because the behavior is  
 626 forbidden by the JAM's definition of causal cycles.

627 We run the tests using the Herd tool [Alglave et al. 2014]. Herd exhaustively enumerates all  
 628 possible executions of a litmus test and checks if each execution is allowed by the model being  
 629 tested. We can then compare the three possible results for each model to see how often executions  
 630 are allowed: Always, Sometimes, and Never. In terms of the behavior being tested Always and  
 631 Sometimes mean that the behavior is allowed, while Never means that it is not allowed. Weaker  
 632 memory models should allow more behaviors and see more Always and Sometimes results, while  
 633 stronger memory models should allow fewer behaviors and see more Never results.

634 Our test suite is built with litmus tests taken from existing research for the three models we  
 635 compare against. Importantly we did not modify any of these tests. However, Herd uses different  
 636 built-in relations to refer to the access types of each model. For example, Herd provides the L built-in

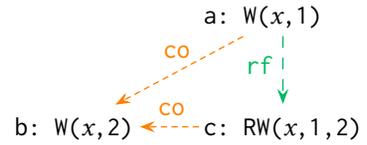


Fig. 13. Read-write Exclusivity

	Pulte et al. [2017]	
	ARMv8	JAM
638	name	
639	WRC+adrrs	Never Never
640	LB+data+data-wsi	Never Never
641	W+RR	Never Never
642	totalco	Never Never
643	PPOCA	Sometimes Sometimes
644	IRIW	Sometimes Sometimes
645	IRIW+adrrs	Never Sometimes
646	IRIW+poaas+LL	Never Sometimes
647	IRIW+poaps+LL	Never Sometimes
648	MP+dmb.sy+addr-ws-rf-addr	Never Sometimes
649	WW+RR+WW+RR+wsilp+poaa+wsilp+poaa	Never Sometimes
650	LB	<b>Sometimes Never</b>

Fig. 14. ARMv8 Litmus Test Comparison

653 to refer to release-writes in ARMv8 litmus tests and the REL built-in to refer to release-writes in  
 654 C11 litmus tests. Thus, for each comparison we include a mapping between the access types of the  
 655 other model and the access modes of the JAM.

656 For each comparison we define the mapping between the built-in relations of the two models,  
 657 we detail our expectations for the results, and then discuss the results of the comparison. As we  
 658 will see our model behaves as expected in all cases.

## 660 5.1 Comparison with ARMv8

661 Herd provides several built-in relations for the ARMv8 litmus  
 662 tests: M for normally memory operations, L for release writes, A for  
 663 acquire reads, and DMB.SY for full fences. There are no SC/Volatile  
 664 accesses.

665 **Mapping.** We map every memory access to opaque mode with  
 666 `opq = M`. This is conservative with respect to our expectation that  
 667 the JAM permits more behaviors than ARMv8. If we were to map some accesses to plain mode they  
 668 would be allowed to exhibit causal cycles. Thus, mapping all memory accesses to opaque mode  
 669 means the JAM model will permit fewer behaviors.

670 We map release-writes to release-acquire mode writes, `rel = L` and acquire-reads to release-  
 671 acquire mode reads, `acq = A`. Note, that M includes L and A so release writes and acquire reads will  
 672 also be treated as opaque mode accesses. Finally we treat all accesses with a full fence between  
 673 them in program order having a specified push order.

674 **Expectations.** As stated, we expect the JAM to be weaker than ARMv8 since the JAM is a  
 675 language memory model which is subject to aggressive compiler optimizations. That is, we expect  
 676 that any time ARMv8 exhibits a behavior the JAM should too and there should be instances where  
 677 the JAM exhibits behaviors that ARMv8 doesn't. The lone exception is cases where the broad  
 678 definition of causal cycles adopted by the JAM will rules out behavior like the load-buffering  
 679 example of Figure 3c.

680 **Results.** Aside from `totalco` and `LB`, all the tests come from the supplementary material accom-  
 681 panying the ARMv8 model of Pulte et al. [2017] which we use for comparison. Figure 14 shows  
 682 the results of our comparison and they agree with our expectations: the JAM is at least as weak as  
 683 ARMv8 except in the case of load-buffering (`LB`).

684 Many of the results owe to the fact that the JAM is not multi-copy atomic. That is, unlike ARMv8,  
 685 different threads can see writes to different locations in different orders. So, `IRIW+*` and `WRC+*` are  
 686

```

661 let opq = M
662 let rel = L
663 let acq = A
664 let spush = po; [DMB.SY]; po

```

Fig. 15. ARMv8 Mapping

687 allowed for the JAM but not for ARMv8. The WW+RR+WW+RR+wsilp+poaa+wsilp+poaa litmus test  
 688 is a variant of IRIW where all writes are release writes and all reads are acquire reads. The writes in  
 689 this test use a weaker form of synchronization than the example in Figure 12a in Section 4.4 where  
 690 the behavior is allowed by the JAM. As a result this behavior is allowed by the JAM. Finally, the  
 691 LB test is identical to the example in Figure 3c from Section 3. This execution is allowed by the  
 692 ARMv8 model but it is a cycle in (po | rf) & opq which is explicitly disallowed by the JAM.

## 694 5.2 Comparison with RC11

695 C11's atomic memory accesses can be annotated with memory orders. Herd provides the following  
 696 built-in relations for the memory orders and accesses in the C11 litmus tests: M for plain memory  
 697 access, RLX for relaxed memory order accesses, REL for release memory order accesses, ACQ for  
 698 acquire memory order accesses, REL\_ACQ for release-acquire memory order read-write accesses, SC  
 699 for sequentially consistent memory order accesses, F & REL for release fences, F & AQR for acquire  
 700 fences, F & SC for sequentially consistent (full) fences, and F \ SC for all other fences.

```
701
702 let opq = RLX | ACQ | REL | ACQ_REL | SC
703 let rel = REL | ACQ_REL | SC
704 let acq = ACQ | ACQ_REL | SC
705 let vol = SC
706 let svo = po;[F & REL];po;[W] | [R];po;[F & ACQ];po
707 let spush = po;[F & SC];po
```

708 **Mapping.** Our mapping for the relations provided by Herd for C11 follows the informal relation-  
 709 ship outlined in the documentation for the JAM [Lea 2018]. All plain accesses in C11 are treated  
 710 as plain accesses in our mapping to the JAM. We map relaxed memory order accesses or stronger  
 711 to opaque mode, opq = RLX | . . . , release memory order accesses or stronger to release mode,  
 712 rel = REL | . . . , acquire memory order accesses or stronger to acquire mode, acq = ACQ |  
 713 . . . . sequentially consistent memory order accesses to volatile mode, vol = SC. We also map  
 714 release fences before writes, po;[F & REL];po;[W], and acquire fences after reads, [R];po;[F &  
 715 ACQ];po, to specified visibility orders, svo. Finally, we map sequentially consistent fences, po;[F  
 716 & SC];po, to push orders between program order earlier and program order later accesses.

717 **Expectations.** The access modes of the JAM are inspired by the memory orders of C11 so this  
 718 comparison is of particular importance. We expect the JAM to match C11 except in cases where  
 719 release-sequences, consume-reads, or causal cycles are involved.

720 **Results.** In Figure 16 we have the results of the comparison with the RC11 model of [Lahav et al.  
 721 2017]. We include the tests from research by Wickerson and Batty [2015], [Vafeiadis et al. 2015],  
 722 [Lahav et al. 2017]. In the case of [Wickerson and Batty 2015] and [Vafeiadis et al. 2015] we used  
 723 the tests directly. In the case of [Lahav et al. 2017] we translated the tests from the paper to Herd  
 724 litmus tests ourselves. The results largely agree with our expectations. The exceptions are places  
 725 where the RC11 model breaks with the C11 specification. We will discuss each in turn.

726 The cyc\_na test is the same as the load-buffering example from Section 3 but all the accesses are  
 727 plain. The JAM allows this cycle in plain mode po | rf because its acyclicity requirement only  
 728 applies to opaque mode or stronger accesses. The RC11 model breaks with C11 by including an  
 729 acyclicity requirement for po | rf for all memory accesses.

730 The lb test is the same as cyc\_na except that all of the accesses are relaxed memory order. In  
 731 the mapping to the JAM this translates into opaque mode accesses which are not allowed to exhibit  
 732 a cycle in po | rf. Thus both models forbid the load-buffering behavior in this case.

733 The mp\_relacq\_rs test leverages the release sequences of C11. Since the JAM does not include  
 734 release sequences it does not forbid the behavior of this test.

735

736  
737  
738  
739  
740  
741  
742  
743  
744

745  
746  
747  
748  
749  
750  
751  
752  
753

754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770

771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784

#### Wickerson and Batty [2015]

name	RC11	JAM
cppmem_iriw_relacq	Sometimes	Sometimes
cppmem_sc_atomics	Never	Never
iriw_sc	Never	Never
mp_fences	Never	Never
mp_relacq	Never	Never
mp_relacq_rs	<b>Never</b>	<b>Sometimes</b>
mp_relaxed	Sometimes	Sometimes
mp_sc	Never	Never

#### Lahav et al. [2017]

name	RC11	JAM
2+2W	Never	Never
IRIW-acq-sc	<b>Sometimes</b>	<b>Never</b>
RWC+syncs	Never	Never
W+RWC	Never	Never
Z6.U	<b>Sometimes</b>	<b>Never</b>

#### Herd X86 Tests

name	x86	JAM
CoRWR	Never	Never
SB+mfences	Never	Never
4.SB	Sometimes	Sometimes
iriw-internal	Sometimes	Sometimes
podrw000	Sometimes	Sometimes
podrw001	Sometimes	Sometimes
SB	Sometimes	Sometimes
SB+rfi-pos	Sometimes	Sometimes
SB+SC	Sometimes	Sometimes
X000	Sometimes	Sometimes
X001	Sometimes	Sometimes
X002	Sometimes	Sometimes
X003	Sometimes	Sometimes
X004	Sometimes	Sometimes
X005	Sometimes	Sometimes
X006	Sometimes	Sometimes
6.SB	Sometimes	<b>timed out</b>
6.SB+prefetch	Sometimes	<b>timed out</b>
iriw	Never	Sometimes
x86-2+2W	Never	Sometimes

Fig. 16. RC11 & x86 Litmus Test Comparison

Our test runs for `fig6` and `fig6_translated` timed out at 5 minutes. The behavior modeled by these tests highlights a quirk in C11's rules for SC accesses. Together they demonstrate that strengthening the memory order of a particular relaxed store in `fig6` to an SC store in `fig6_translated` creates new behaviors. That is, the tests demonstrate that the memory orders of C11 are not monotonic. RC11 includes a fix proposed by [Vafeiadis et al. 2015] which forbids this behavior. In Section 6 we prove that the JAM's access modes are indeed monotonic in our model which means that stronger access modes exhibit fewer behaviors, thus the behavior in these tests is forbidden.

In the case of `IRIW-acq-sc`, `Z6.U`, and `IRIW-sc-rlx-acq` our model forbids the behavior in keeping with the C11 specification. We will consider each case and discuss why RC11 does not forbid each behavior.

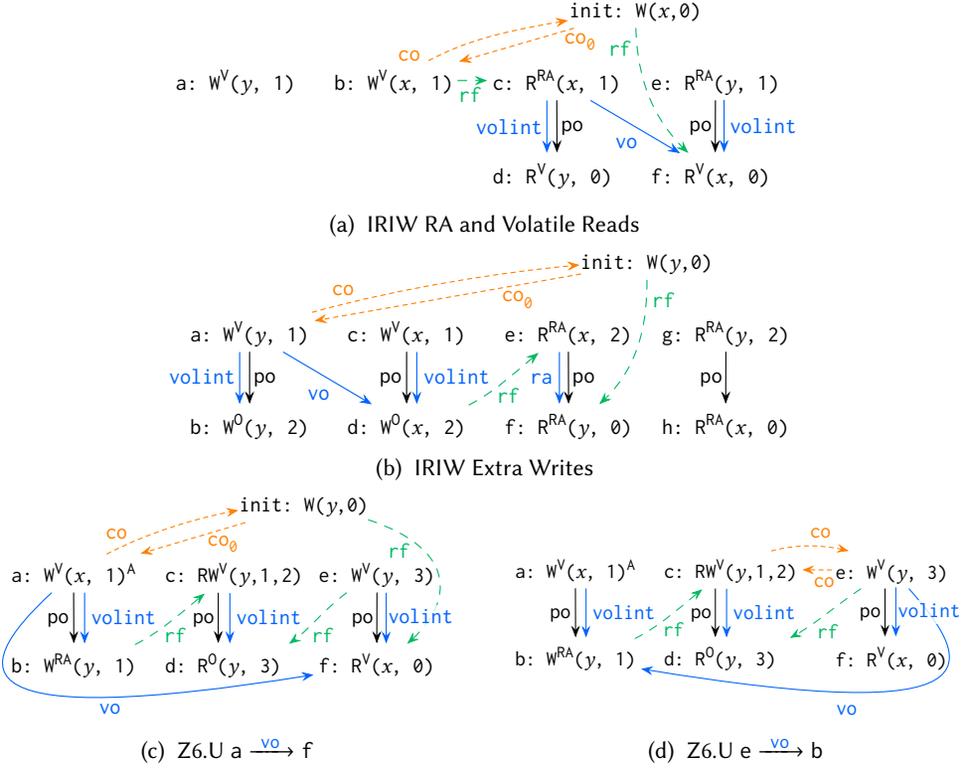


Fig. 17. Z6.U &amp; IRIW RA with Volatile Reads Litmus Tests

The IRIW-acq-sc test appears in Figure 17a. The reason this is forbidden in our model is that all accesses before a SC/Volatile read are ordered by `volint`. Then, because there are two such orders in the reading threads there is either a visibility order,  $c \xrightarrow{vo} f$  or  $d \xrightarrow{vo} e$ . Thus, one of the two reads must see the non-initialization write using the same reasoning from Section 4.4 for Figure 12b. This effectively emulates the "leading" fence compilation scheme described in [Lahav et al. 2017], where a full fence is placed before the SC/Volatile accesses. The authors (personal communication) point out that this scheme should forbid this behavior. By contrast RC11 relaxes the C11 model to accommodate a leading *or* trailing fence compilation scheme. In this case, if the trailing fence scheme is used there's no ordering provided to any of the SC accesses and the execution is allowed.

Two possible executions for Z6.U test appears in Figures 17c and 17d. They represent each case of the visibility orders induced by  $a \xrightarrow{volint} b$  and  $e \xrightarrow{volint} f$ . In 17c we have  $a \xrightarrow{vo} f$ , which means that  $f$  must have seen  $a$  and could not read from the initialization by `cowr`. In 17d we have  $e \xrightarrow{vo} b$ . We also have  $b \xrightarrow{rf} c$ . Then together we have  $e \xrightarrow{vo} b \xrightarrow{rf} c$ . Then because  $c$  and  $e$  are both writes to  $y$  we have  $e \xrightarrow{co} c$  by `coww`. Separately, since  $c \xrightarrow{volint} d$  we have  $c \xrightarrow{vo} d$ . Then since  $e \xrightarrow{rf} d$  we have  $c \xrightarrow{co} e$  by `cowr`. Thus we have a cycle in `co` and the execution is forbidden.

This effectively emulates the "trailing" fence compilation scheme described by Lahav et al. [2017], where a full fence is placed after SC/Volatile accesses. They point out that this scheme should forbid this behavior. Again, RC11 relaxes C11 to model both schemes and under a leading fence compilation scheme the behavior this allowed.

Finally, in Figure 17b we have our own variant of IRIW called IRIW-sc-r1x-acq based on the WW+RR+. . . test in the ARMv8 suite. In this case, if the compiler is inserting trailing fences after  $a$  and  $c$ , then either  $a \xrightarrow{vo} d$  or  $b \xrightarrow{vo} c$ . Taking the first case we can construct  $b \xrightarrow{vo+} f$  from

834 a  $\xrightarrow{vo}$  d  $\xrightarrow{rf}$  a  $\xrightarrow{ra}$  f and a cycle in `co` by `cowr`. The second case is similar but with the read of `x`  
 835 in the 4th thread.

836 Again, the RC11 model does not forbid this execution. In our discussion with the authors, they  
 837 suggested two possible interpretations for this behavior. The first is a leading fence compilation  
 838 scheme which we have discussed above. The second is a merge of each pair of writes into a single  
 839 write. In the second case we suggest that the Java compiler should not merge opaque mode (or  
 840 stronger) accesses.

841 In summary, our model correctly forbids the behavior in each of these cases. The reason is that  
 842 volatile reads would be preceded by a full fence and volatile writes would be followed by a full  
 843 fence in the presence of mixed mode programs with SC accesses. Importantly, the tension between  
 844 optimal compilation and a simple model exists here as it does with C11.

### 846 5.3 Comparison with x86

847 Herd provides the following built in relations for x86 litmus tests: M for memory accesses, SFENCE  
 848 for fences for intra-thread ordering of writes with other accesses, LFENCE for ordering intra-thread  
 849 ordering of reads with other accesses, and MFENCE as a full fence with cross thread ordering  
 850 guarantees.

```
851 let opq = M
852 let svo = [W];po;[SFENCE];po | [R];po;[LFENCE];po
853 let spush = po;[MFENCE];po
854
```

855 **Mapping.** We map from all regular memory accesses to opaque mode `opq = M` in keeping  
 856 with the opaque mode mapping from our comparison with ARMv8. We map SFENCE to specified  
 857 visibility orders from writes to other memory accesses across the fence, `[W];po;[SFENCE];po`. We  
 858 map LFENCE to specified visibility orders from reads to other memory accesses across the fence,  
 859 `[R];po;[LFENCE];po`. Finally we map MFENCE as a specified push orders between any access before  
 860 the fence to any access after the fence, `po;[MFENCE];po`.

861 **Expectations.** We expect the JAM to be weaker than x86 in all cases. The only weak behavior  
 862 that x86 exhibits is reordering writes with reads (store-buffering) which the JAM also allows in  
 863 opaque mode.

864 **Results.** Figure 16 shows the results of our comparison with the x86 model included with  
 865 Herd. The litmus tests are also included with Herd for the model. The two tests which timed out,  
 866 `6.SB+prefetch` and `6.SB`, are store buffering variants. Given that the JAM can reorder writes with  
 867 reads the behavior of these tests is allowed. Otherwise the tests confirm our expectations.

## 869 6 METATHEORY

870 Here, we develop a metatheory for our model of the JAM. First, we detail the semantics and then  
 871 sketch the proofs for the theorems listed in the introduction. We show that each mode, from  
 872 Volatile to Opaque, admits strictly more executions, otherwise referred to as monotonicity. We  
 873 also show that properly synchronized programs, i.e. race-free programs as defined by Boehm and  
 874 Adve [2008], will only exhibit sequentially consistent behaviors as required by Java's data-race-free  
 875 (DRF) guarantee. Finally, we show that acquire mode for all reads obviates the inclusion of the  
 876 JAM's acyclic causality requirement.

877 These theorems further validate the definitions of our model, give more evidence that the model  
 878 is complete with respect to the documentation, and clearly demonstrate that the semantics is  
 879 suitable for formal reasoning. The semantics, lemmas and theorems have been mechanized in Coq.  
 880 The source is available for inspection in the supplementary material and it includes instructions on  
 881 where to find everything included below.

882

## 6.1 Semantics

The axioms of our model apply to memory event graphs, each of which represents one possible execution of a program. Thus far we have relied on Herd to generate executions for a program and then apply the axioms of our model to rule on whether an execution is allowed. Our formalization in Coq replaces the Herd machinery with the history fragment of the semantics of Crary and Sullivan [2015] to model the set of possible executions. Here we will give enough detail about the history semantics to understand the content of our theorems and proofs. We provide a more comprehensive set of definitions and intuition for the semantics in Appendix C. The full semantics is available in the supplementary material.

We use  $n$  and  $l$  to range over natural numbers and memory locations. We use  $i$  to range over unique memory access identifiers where previously we have used  $a, b, c$ , etc. We use  $m$  to represent one of the four access modes, P for plain, O for opaque, RA for release acquire and V for volatile. Memory accesses,  $a$ , can take the form of reads,  $l_m$ , and writes,  $l_m := n$  with their accompanying modes as well as read-writes,  $RW(l, n)$ .

We use  $H$  to range over lists of memory events that represent executions. We call these lists histories and we use  $H(h)$  means that the memory event  $h$  is in  $H$ . Memory events record the program's interactions with a history. For example,  $H(\text{is}(i, l_m := n))$  records that the identifier  $i$  is a write to  $l$  of the value  $n$  and  $H(\text{exec}(i))$  records that  $i$  has executed, subject to the axioms of the memory model. We use  $i \xrightarrow{R}_H i$  to represent our model's relations as they apply to  $H$ .

We also require that histories satisfy some basic well-formedness conditions. For example, a read must take its value from a write to the same location which has been executed. Importantly, we require that the trace order respects any intra-thread ordering created by visibility generated from specified orders, release-acquire accesses, or volatile accesses. This is in keeping with our restriction on trace order as detailed in Section 4.2. We call this well-formedness condition executable. In what follows, the well-formedness conditions are unified as *trace coherence*. Otherwise, the only restrictions on histories come from the acyclicity conditions outlined in Section 4.1.

## 6.2 Theorems

To begin, we demonstrate the monotonicity of our access mode definitions. We define the reflexive ordering of the access modes as  $P \sqsubseteq O \sqsubseteq RA \sqsubseteq V$  and extend it to accesses  $l_{m_1} \sqsubseteq l_{m_2}, l_{m_1} := n_1 \sqsubseteq l_{m_2} := n_2, RW(l, n_1) \sqsubseteq RW(l, n_2)$  whenever  $m_1 \sqsubseteq m_2$ . As a technical matter we treat read-writes as always having the same order. We extend the order to histories by matching identifiers and ordering the accesses.

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

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

**THEOREM 1 (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  $\text{acyclic}(\xrightarrow{\text{co}}_{H_2})$*

We make two notes. First the absence of specified orders in  $H_2$  is a technical convenience since specified order edges are not related to the strength of the access modes for reads and writes. Second, we focus on the acyclic coherence requirement because the match assumption means that it would be trivial to satisfy the acyclic causality requirement for  $H_2$  supposing it is true of  $H_1$  because po and rf have fewer edges in  $H_2$ .

PROOF SKETCH. We assume  $i \xrightarrow{\text{co}}_{H_2} i$  for some  $i$  and show that this must mean  $i \xrightarrow{\text{co}}_{H_1} i$  which is a contradiction. This is straight forward by induction on  $i \xrightarrow{\text{co}}_{H_2} i$ , noting that each case of visibility in  $H_2$  will exist in  $H_1$  because of stronger access modes in  $H_1$ .  $\square$

For the last two theorems will first establish that the main component of visibility in our model, vvo, is irreflexive.

LEMMA 1 (IRREFLEXIVE VISIBILITY). *If  $H$  is trace coherent then, for all  $i$ ,  $\neg i \xrightarrow{\text{vvo}^+}_H i$ .*

PROOF SKETCH. This follows from two facts. First, vvo derives from orderings that are always either program ordered from intra-thread synchronization (svo, spush, ra, volint) or trace ordered (pushto, rf). Second, both relations are total and to is consistent with the po edges for intra-thread visibility by the executable well-formedness condition.  $\square$

Next we show that if a program is properly synchronized then it will only exhibit sequentially consistent behavior. We require the following standard definitions including the traditional notion of sequential consistency [Shasha and Snir 1988]:

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

We also require a definition of proper synchronization in keeping with our focus on visibility.

$$i_1 \xrightarrow{\text{sync}}_H i_2 \triangleq \exists i_3, i_1 \xrightarrow{\text{vvo}^+}_H i_3 \xrightarrow{\text{po}^*}_H i_2 \wedge \forall i_4, i_3 \xrightarrow{\text{po}^*}_H i_4 \Rightarrow i_3 \xrightarrow{\text{vvo}}_H i_4$$

The idea is that any conflicting access is visibility ordered by some mechanism, be it a specified order (fence) or a strong mode for  $i_3$ . Then, following the definition of “type 2” data-races from Boehm and Adve [2008], we say that  $H$  is *race-free* when, for all conflicting accesses  $i_1$  and  $i_2$ , we have  $i_1 \xrightarrow{\text{sync}}_H i_2$  or  $i_2 \xrightarrow{\text{sync}}_H i_1$ .

THEOREM 2 (DRF-SC). *If  $H$  is trace coherent, race free and, acyclic( $\xrightarrow{\text{co}}_H$ ), then acyclic( $\xrightarrow{\text{sc}}_H$ ).*

PROOF SKETCH. We assume  $i \xrightarrow{\text{sc}^+}_H i$  for some  $i$  and show a contradiction. By Lemma 1 it is enough to demonstrate a cycle in vvo+.

We can show that for any  $i_1$  and  $i_2$ , if we have  $i_1 \xrightarrow{\text{com}}_H i_2$  then we have  $i_1 \xrightarrow{\text{sync}}_H i_2$ . Note that any accesses related by com are conflicting. Then we have either  $i_1 \xrightarrow{\text{sync}}_H i_2$  or  $i_2 \xrightarrow{\text{sync}}_H i_1$ . In each case for com we can show that  $i_2 \xrightarrow{\text{sync}}_H i_1$  creates a cycle in the coherence order so it must be  $i_1 \xrightarrow{\text{sync}}_H i_2$ .

Then, since po is irreflexive, we have that any sequence  $i \xrightarrow{\text{sc}^+}_H i$  must include at least one com edge. Then since com edges are also sync edges, when we have  $i_1 \xrightarrow{\text{sc}^+}_H i_2$  we also have  $i_1 \xrightarrow{\text{po}|\text{com}^+}_H i_2$  and therefore  $i_1 \xrightarrow{\text{po}|\text{sync}^+}_H i_2$ .

But then we can rearrange a cycle in  $i \xrightarrow{\text{po}|\text{sync}^+}_H i$  to be  $i \xrightarrow{\text{sync}^+}_H i$  by appending a leading po edge to the end. Observe that for any sequence  $i_1 \xrightarrow{\text{sync}^+}_H i_2$  we have  $i_1 \xrightarrow{\text{vvo}^+}_H i_2$ , so we have  $i \xrightarrow{\text{vvo}^+}_H i$  as required.  $\square$

Finally, we show that if all reads are acquire-reads then the JAM’s acyclic causality requirement is unnecessary. This theorem demonstrates the soundness of proposed compiler implementations for satisfying the acyclic causality requirement of the JAM [Ou and Demsky 2018].

THEOREM 3 (CAUSAL ACQUIRE-READS). *If  $H$  is trace coherent and all reads in  $H$  are acquire-reads, then acyclic( $\xrightarrow{\text{po}|\text{rf}}_H$ ).*

981 PROOF SKETCH. We assume  $i \xrightarrow{po|rf+}_H i$  for some  $i$  and show a contradiction. By Lemma 1, it is  
 982 enough to demonstrate a cycle in  $vvo+$ .

983 First, note that any sequence  $i_1 \xrightarrow{rf}_H i_2 \xrightarrow{po}_H i_3$  implies  $i_1 \xrightarrow{vvo+}_H i_3$  because  $rf$  implies  $vvo$  and  
 984 because, by assumption, the read  $i_2$  is an acquire read and  $i_3$  is program order later, so  $i_2 \xrightarrow{vvo}_H i_3$ .

985 Let  $\xrightarrow{rfpoq}_H$  be a reads-from edge followed by an optional program order edge. Note that, by  
 986 induction and the fact that  $rf$  and  $rfpo$  imply  $vvo$  we can show that  $i_1 \xrightarrow{rfpoq}_H i_2$  implies  $i_1 \xrightarrow{vvo+}_H i_2$

987 Then, since  $po$  is irreflexive, we have that any sequence  $i \xrightarrow{po|rf+}_H i$  must include at least one  
 988  $rf$  edge. Thus we can rearrange to get  $i' \xrightarrow{rfpoq+}_H i'$  and we have  $i' \xrightarrow{vvo+}_H i'$  by the above, as  
 989 required.  $\square$

## 990 7 UNOBSERVABLE TOTAL COHERENCE ORDER

991 Here we will demonstrate that the effects of a total coherence order are unobservable in our model.  
 992 This is a surprising result since a total coherence order is intuitively associated with maintaining  
 993 per-location state and we would expect concurrent writes to have a material impact on the behavior  
 994 of programs.

995 To start, recall that a feature of a memory model as exhibited by an example program is defined  
 996 in binary terms. The feature is present when any executions are allowed and the feature is absent  
 997 when no executions are allowed. Then we say that a feature is *observable* when we can construct  
 998 any example program such that the feature's presence in the model changes whether any executions  
 999 are allowed. So, to observe a total coherence order, we must construct a program that changes  
 1000 whether executions are allowed based on the presence of the total order. We will show that this is  
 1001 impossible under our model.

1002 First, note that in our model there is no way to allow previously forbidden executions when  
 1003 adding coherence order edges. This means we can't observe the feature by going from some  
 1004 executions *with* a total order to no executions *without* a total order. Thus, any program that allows  
 1005 us to observe the total coherence order must forbid all executions in the presence of a total order.  
 1006 We will show that such a program will also forbid all executions when the total order is removed.  
 1007

1008 **THEOREM 4.** *It is impossible to construct an example program for the JAM, such that, for any two*  
 1009 *writes, if they are totally ordered there are no valid executions, and if they are not totally ordered there*  
 1010 *is at least one valid execution.*

1011 **PROOF.** From a total order we have that all candidate executions can be divided between the two  
 1012 directions. For some writes to the same location  $a$  and  $b$  we have:

$$1013 \quad a \xrightarrow{co} b \vee \quad (1)$$

$$1014 \quad b \xrightarrow{co} a \quad (2)$$

1015 We begin by eliminating all executions for the first direction, (1). This means we must construct  
 1016 a cycle in  $co$ , such that, for all executions including (1) we have :

$$1017 \quad a \xrightarrow{co} b \quad (1)$$

$$1018 \quad b \xrightarrow{co} \dots \xrightarrow{co} a \quad (3)$$

1019 With this cycle, executions including (1) are forbidden by the acyclicity requirement for  $co$ .  
 1020 Observe that (3) must be present in all executions of the program. If it is not then we can divide  
 1021 the executions that do not have (3) between executions with (1) and executions with (2). For  
 1022 the executions without (3) and with (1) the execution is allowed and we have failed to forbid all  
 1023 executions. Now we must eliminate all executions for (2) with another cycle:

$$1024 \quad b \xrightarrow{co} a \quad (2)$$

$$1025 \quad a \xrightarrow{co} \dots \xrightarrow{co} b \quad (4)$$

1030 Moreover (4) must persist for all executions for the same reason that (3) persisted. Then it must  
 1031 be that, for all executions we have:

$$1032 \quad b \xrightarrow{\text{co}} \dots \xrightarrow{\text{co}} a \quad (3)$$

$$1033 \quad a \xrightarrow{\text{co}} \dots \xrightarrow{\text{co}} b \quad (4)$$

1034 Then, even if we remove the total order, (1) and (2), from the model, all executions of the program  
 1035 will still have a forbidden cycle by (3) and (4) and we can not demonstrate a single execution that is  
 1036 allowed.  $\square$

1037 Critically, this line of reasoning depends on the fact that we have a single acyclicity requirement  
 1038 in which `co` participates and both (3) and (4) form a cycle that violates that requirement. If that  
 1039 were not the case then (3) and (4) would not necessarily create a cycle or forbid any executions. As a  
 1040 result this reasoning does not apply to the ARMv8 and RC11 models as they have multiple acyclicity  
 1041 requirements involving the coherence order. In Appendix D we have included a litmus test for  
 1042 ARMv8 that shows it is possible to construct a program for that model that behaves differently  
 1043 with and without the total order.

1044 Demonstrating that a total coherence order is not observable in our model means that the JAM  
 1045 may adopt a total coherence order without affecting the outcomes of the model. This would allow  
 1046 it to emulate other mainstream memory models in this respect and recover the intuitive notion of  
 1047 per-location state.

## 1048 8 RELATED WORK

1049 The JAM is intended as an update and expansion of Java's memory model. As a result, we have  
 1050 taken a fresh look at how to construct our model, but our work takes inspiration from a large body  
 1051 of research on such formalization efforts.

### 1052 8.1 The Original Java Memory Model

1053 The original Java Memory Model [Manson et al. 2005] included an early attempt to model standard  
 1054 compiler optimizations while ruling out thin-air reads. The specification in that work was part  
 1055 prose and part formal definitions and the "causality" mechanisms at the heart of the model made  
 1056 the definitions complex. Taken together these issues made the model unsuitable for the tasks which  
 1057 one normally formalizes a programming language, namely accurate discourse, metatheory, and  
 1058 algorithm verification. Later work by Ševčík and Aspinnall [2008] manually examined a large suite  
 1059 of litmus tests to show that the model disallowed some standard compiler optimizations. Eventually  
 1060 the model was fully formalized in Coq by Huisman and Petri [2007] and Aspinnall and Ševčík [2007],  
 1061 but, to the best of our knowledge, there is no way to easily test the behavior of example programs.  
 1062 By contrast we have constructed a readable and testable model for Java's new access modes.

### 1063 8.2 C11, ARM, and x86

1064 The C11 memory model, which served as the inspiration for Java's access modes, has seen extensive  
 1065 study. It was originally formalized by Batty et al. [2011] with later revisions to include read-modify-  
 1066 writes and fences [Sarkar et al. 2012]. The work of Vafeiadis et al. [2015], from which we draw  
 1067 the largest set of our C11 litmus tests, studied the soundness of common compiler optimizations  
 1068 under the model of C11 by Morriset et al. [2013]. Our monotonicity theorem is modeled after the  
 1069 same theorem from Vafeiadis et al. [2015]. Most recently, the axiomatic model of Lahav et al. [2017]  
 1070 incorporated the proposed fixes of Vafeiadis et al. [2015] and addressed the unsound compilation  
 1071 strategies which we discussed in Section 5. Also, further progress has been made on new models  
 1072 for C11 that more successfully support standard compiler optimizations without the problem of  
 1073 thin-air reads [Kang et al. 2017; Pichon-Pharabod and Sewell 2016]. However, like the original Java  
 1074 memory model, these models rely on complex formal constructs (promises and event structures  
 1075  
 1076  
 1077  
 1078

1079 respectively). Again, our semantics remains relatively simple thanks to the JAM's broad definition  
1080 of causal cycles.

1081 As outlined previously there are several important differences between C11 and Java. First, the  
1082 partial coherence order required careful consideration. The consequences of this design choice  
1083 manifests most clearly where atomic read-writes are concerned. Second, the lack of legacy features,  
1084 like C11's release sequences, and the simple mechanism that forbids causal cycles allowed us to  
1085 build a simple model. In turn, the simplicity of the model makes it more readable than existing C11  
1086 models and it allowed us to argue forcefully that the model should adopt a total coherence order.

1087 These differences appear most clearly in our litmus test comparison with RC11 model of [Lahav  
1088 et al. \[2017\]](#). Of particular note are the `mp_relacq_rs` and `lb` tests. In the first case the JAM is  
1089 weaker than C11 because it does not support release sequences, in the second case it is stronger  
1090 because the C11 specification gives no concrete definition for how to rule out thin-air reads. Notably,  
1091 the RC11 herd model also forbids causal cycles in `po | rf` so the load buffering behavior is forbidden.  
1092 This method of preventing thin-air reads in RC11 is included to enable their proofs of soundness  
1093 for compilation to POWER and not as a representation of the C11 specification.

1094 Programs that mix SC/volatile access modes are the primary focus of [Lahav et al. \[2017\]](#). In  
1095 particular the leading and trailing fence insertion approach to compilation was shown to be  
1096 unsound for Power under models prior to RC11. We compared our model with RC11 in Section 5.  
1097 Our semantics correctly forbids the behavior described in `IRIW-sc-rlx-acq`, `IRIW-aqc-sc` and  
1098 `Z6.U` in keeping with the JAM documentation.

1099 Hardware memory models have also seen extensive study. x86 was studied by [Owens et al. \[2009\]](#)  
1100 and [Alglave et al. \[2014\]](#). We use the model included with Herd in our litmus test comparisons and  
1101 we saw that x86 was stronger than the JAM, as expected. ARM processors have traditionally had  
1102 a much weaker memory model when compared with x86. Recently the model for ARMv8 [[Pulte  
1103 et al. 2017](#)] expanded the guarantees made by the architecture to include multi-copy-atomicity. We  
1104 saw the effects of this in the behavior of the `IRIW-*` and `WRC-*` litmus tests from our comparison  
1105 between the JAM and ARMv8. As expected the JAM is weaker than ARMv8 in every case except  
1106 where cycles in `po | rf` are concerned.

### 1107 8.3 The Relaxed Memory Calculus

1108 Our mechanized semantics is based on the history fragment of the Relaxed Memory Calculus of  
1109 [Crary and Sullivan \[2015\]](#). We use their concept of specified push orders and we draw inspiration  
1110 for our definitions from their notion of visibility. Also, the proof of our theorems benefited greatly  
1111 from the library of lemmas included with the RMC mechanization. However, their purpose was to  
1112 model a weaker version of C11 in the interest of generality while, our goal is to model the JAM.  
1113 Importantly, we do not employ the execution orders of RMC, our coherence definition is far more  
1114 compact and we have added the `corr` rule to follow a more standard notion of coherence.  
1115

## 1116 9 CONCLUSION

1117 We have presented the first formal model for Java's access modes. Our model is precise, complete,  
1118 and testable. We have validated our model against the expected behavior of example programs  
1119 relative to three other mainstream memory models. We have further validated the model by using  
1120 it to prove general theorems about the semantics of the Java access modes. Finally, we used our  
1121 model to demonstrate that Java's access modes can adopt a total coherence order.

1122 We believe our model is a strong foundation for a larger formalization effort around weak-  
1123 memory concurrency on the Java platform. In particular we would like expand the model and use  
1124 it to verify the lock-free algorithms of the standard library.  
1125  
1126  
1127

## REFERENCES

- 1128  
1129 Jade Alglave, Anthony Fox, Samin Ishtiaq, Magnus O. Myreen, Susmit Sarkar, Peter Sewell, and Francesco Zappa Nardelli.  
1130 2008. The Semantics of Power and ARM Multiprocessor Machine Code. In *Proceedings of the 4th Workshop on Declarative  
1131 Aspects of Multicore Programming (DAMP '09)*. ACM, New York, NY, USA, 13–24. <https://doi.org/10.1145/1481839.1481842>
- 1132 Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding Cats: Modelling, Simulation, Testing, and Data Mining  
1133 for Weak Memory. *ACM Trans. Program. Lang. Syst.* 36, 2, Article 7 (July 2014), 74 pages. <https://doi.org/10.1145/2627752>
- 1134 David Aspinall and Jaroslav Ševčík. 2007. Formalising Java's Data Race Free Guarantee. In *Proceedings of the 20th International  
1135 Conference on Theorem Proving in Higher Order Logics (TPHOLS'07)*. Springer-Verlag, Berlin, Heidelberg, 22–37. <http://dl.acm.org/citation.cfm?id=1792233.1792237>
- 1136 Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ Concurrency. *SIGPLAN  
1137 Not.* 46, 1 (Jan. 2011), 55–66. <https://doi.org/10.1145/1925844.1926394>
- 1138 Mark Batty and Peter Sewell. 2014. The Thin-air Problem. <http://www.cl.cam.ac.uk/~pes20/cpp/notes42.html>. [Online,  
1139 accessed Dec 2018].
- 1140 John Bender, Mohsen Lesani, and Jens Palsberg. 2015. Declarative Fence Insertion. In *Proceedings of the 2015 ACM SIGPLAN  
1141 International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 2015)*. ACM,  
1142 New York, NY, USA, 367–385. <https://doi.org/10.1145/2814270.2814318>
- 1143 Hans-J. Boehm and Sarita V. Adve. 2008. Foundations of the C++ Concurrency Memory Model. In *Proceedings of the 29th  
1144 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '08)*. ACM, New York, NY, USA,  
1145 68–78. <https://doi.org/10.1145/1375581.1375591>
- 1146 C++ Standards Committee, Pete Becker, et al. 2010. Programming Languages-C++(final committee draft). C++ standards  
1147 committee paper WG21/N3092= J16/10-0082.
- 1148 Jon Corbet. 2012. ACCESS\_ONCE(). <https://lwn.net/Articles/508991/>. [Online, accessed Dec 2018].
- 1149 Karl Crary and Michael Sullivan. 2015. A Calculus for Relaxed Memory. In *Proceedings of POPL'15, ACM Symposium on  
1150 Principles of Programming Languages*.
- 1151 Marieke Huisman and Gustavo Petri. 2007. The Java memory model: a formal explanation. *VAMP 7* (2007), 81–96.
- 1152 JDK9. 2017. VarHandle (Java SE 9 & JDK 9). <https://docs.oracle.com/javase/9/docs/api/java/lang/invoke/VarHandle.html>.  
1153 [Online, accessed March 2019].
- 1154 Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A Promising Semantics for Relaxed-  
1155 memory Concurrency. In *Proceedings of POPL'17, SIGPLAN–SIGACT Symposium on Principles of Programming Languages*.
- 1156 Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing Sequential Consistency in  
1157 C/C++11. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2017)*. ACM, New York, NY, USA, 618–632. <https://doi.org/10.1145/3062341.3062352>
- 1158 Doug Lea. 2017. JEP 193. <http://openjdk.java.net/jeps/193>. [Online, accessed March 2019].
- 1159 Doug Lea. 2018. Using JDK 9 Memory Order Modes. <http://gee.cs.oswego.edu/dl/html/j9mm.html>. [Online, accessed March  
1160 2019].
- 1161 Jeremy Manson, William Pugh, and Sarita V. Adve. 2005. The Java Memory Model. In *Proceedings of the 32nd ACM  
1162 SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '05)*. ACM, New York, NY, USA, 378–391.  
1163 <https://doi.org/10.1145/1040305.1040336>
- 1164 Robin Morisset, Pankaj Pawan, and Francesco Zappa Nardelli. 2013. Compiler testing via a theory of sound optimisations in  
1165 the C11/C++ 11 memory model. In *ACM SIGPLAN Notices*, Vol. 48. ACM, 187–196.
- 1166 Peizhao Ou and Brian Demsky. 2018. Towards Understanding the Costs of Avoiding Out-of-thin-air Results. *Proc. ACM  
1167 Program. Lang.* 2, OOPSLA, Article 136 (Oct. 2018), 29 pages. <https://doi.org/10.1145/3276506>
- 1168 Scott Owens, Susmit Sarkar, and Peter Sewell. 2009. A Better x86 Memory Model: X86-TSO. In *Proceedings of the 22nd  
1169 International Conference on Theorem Proving in Higher Order Logics (TPHOLS '09)*. Springer-Verlag, Berlin, Heidelberg,  
1170 391–407. [https://doi.org/10.1007/978-3-642-03359-9\\_27](https://doi.org/10.1007/978-3-642-03359-9_27)
- 1171 Jean Pichon-Pharabod and Peter Sewell. 2016. A Concurrency Semantics for Relaxed Atomics That Permits Optimisation  
1172 and Avoids Thin-air Executions. *SIGPLAN Not.* 51, 1 (Jan. 2016), 622–633. <https://doi.org/10.1145/2914770.2837616>
- 1173 Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. 2017. Simplifying ARM Concurrency:  
1174 Multicopy-atomic Axiomatic and Operational Models for ARMv8. *Proc. ACM Program. Lang.* 2, POPL, Article 19  
1175 (Dec. 2017), 29 pages. <https://doi.org/10.1145/3158107>
- 1176 Susmit Sarkar, Kayvan Memarian, Scott Owens, Mark Batty, Peter Sewell, Luc Maranget, Jade Alglave, and Derek Williams.  
2012. Synchronising c/c++ and power. *Acm Sigplan Notices* 47, 6 (2012), 311–322.
- Jaroslav Ševčík and David Aspinall. 2008. On Validity of Program Transformations in the Java Memory Model. In *ECOOP  
2008 – Object-Oriented Programming: 22nd European Conference Paphos, Cyprus, July 7-11, 2008 Proceedings*, Jan Vitek  
(Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 27–51. [https://doi.org/10.1007/978-3-540-70592-5\\_3](https://doi.org/10.1007/978-3-540-70592-5_3)
- Dennis Shasha and Marc Snir. 1988. Efficient and Correct Execution of Parallel Programs That Share Memory. *ACM Trans.  
Program. Lang. Syst.* 10, 2 (April 1988), 282–312. <https://doi.org/10.1145/42190.42277>

- 1177 Michael J. Sullivan. 2015. Low-level Concurrent Programming Using the Relaxed Memory Calculus. <https://www.msully.net/stuff/thesprop.pdf>. PhD. Dissertation [Online, accessed July 2017].
- 1178
- 1179 Viktor Vafeiadis, Thibaut Balabonski, Soham Chakraborty, Robin Morisset, and Francesco Zappa Nardelli. 2015. Common
- 1180 Compiler Optimisations Are Invalid in the C11 Memory Model and What We Can Do About It. In *Proceedings of the*
- 1181 *42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '15)*. ACM, New York,
- 1182 NY, USA, 209–220. <https://doi.org/10.1145/2676726.2676995>
- 1183 John Wickerson and Mark Batty. 2015. Taming the complexities of the C11 and OpenCL memory models. *arXiv preprint*
- 1184 *arXiv:1503.07073* (2015).
- 1185
- 1186
- 1187
- 1188
- 1189
- 1190
- 1191
- 1192
- 1193
- 1194
- 1195
- 1196
- 1197
- 1198
- 1199
- 1200
- 1201
- 1202
- 1203
- 1204
- 1205
- 1206
- 1207
- 1208
- 1209
- 1210
- 1211
- 1212
- 1213
- 1214
- 1215
- 1216
- 1217
- 1218
- 1219
- 1220
- 1221
- 1222
- 1223
- 1224
- 1225

1226 **A FULL HERD MODEL**

```

1227 let opq = O | RA | V
1228 let rel = W & RA
1229 let acq = R & RA
1230 let vol = V
1231
1232 (* release acquire ordering *)
1233 let ra = po;[rel] | [acq];po
1234
1235 (* intra-thread volatile ordering *)
1236 let volint = po;[vol & R] | [vol & W];po
1237
1238 (* intra-thread ordering constraints *)
1239 let into = svo | spush | ra | volint
1240
1241 (* define trace order, ensure it respects rf and intra-thread specified orders *)
1242 (* Note that ((W * FW) & loc & ~id) = cofw *)
1243 with to from linearisations( $\mathbb{M} \sqcup \text{IW}$ , ((W * FW) & loc & ~id) | rf | into)
1244
1245 (* cross thread push ordering extended with volatile memory accesses *)
1246 let push = spush | volint
1247 let pushto = to+ & (domain(push) * domain(push))
1248
1249 (* extend ra visibility *)
1250 let vvo = rf | svo | ra | push | pushto;push
1251 let vo = vvo+ | po-loc
1252
1253 include "filters.cat"
1254 let WWco(rel) = WW(rel) & loc & ~id
1255 let cofw = WWco((W * FW))
1256
1257 (* coherence rules *)
1258 let coint = loc &  $\text{IW}^*(\mathbb{W} \sqcup \text{IW})$ 
1259 let cown = WWco(vo)
1260 let cown = WWco(vo;invrf)
1261 let corw = WWco(vo;po)
1262 let corr = WWco(rf;po;invrf)
1263
1264 (* general definition from RC11, works for atomic rws and split instruction rws *)
1265 let rmw-jom = [RMW] | rmw
1266
1267 (* read-write rules *)
1268 let cormwtotal = WWco(((range(rmw-jom) * _) | (_ * range(rmw-jom))) & to)
1269 let cormwexcl = WWco((rf;rmw-jom)^-1;co-jom)
1270
1271 let rec co-jom = cown | cown | corw | corr
1272 | cofw | coint | cormwtotal
1273 | WWco((rf;rmw-jom)^-1;co-jom)
1274
1275 acyclic (po | rf) & opq
1276 acyclic co-jom
1277
1278
1279

```

## B SPECIFIED ORDERS: A MAPPING FOR ARMV8 AND X86 READ-WRITES

Specified orders can be compiled to existing synchronization using a fairly direct mapping. Here we give an example mapping for ARMv8. We also discuss how specified visibility orders can be elided when the target platform for compilation already enforces them.

$$\begin{aligned} [W]; \text{svo}; [M] &\rightsquigarrow [W]; \text{po}; [\text{DMB ST}]; \text{po}; [M] \\ [R]; \text{svo}; [M] &\rightsquigarrow [R]; \text{po}; [\text{DMB LD}]; \text{po}; [M] \\ [M_1]; \text{push}; [M_2] &\rightsquigarrow [M_1]; \text{po}; [\text{DSB}]; \text{po}; [M_2] \end{aligned}$$

Note that this mapping is informal. The fence instruction for each mapping must be inserted on *all* program order paths between the ordered instruction as demonstrated originally by [Bender et al. \[2015\]](#) and [Sullivan \[2015\]](#).

The value of these orders can be seen in the case of atomic read-writes on x86. The JAM makes no intra-thread ordering guarantees for atomic read-writes even though they exist on some architectures like x86 [\[Owens et al. 2009\]](#). This might result in unnecessary synchronization to achieve a desired outcome that requires such intra-thread ordering. However, using specified orders means that a compiler can make intelligent decisions based on the target platform. For example if we have a specified visibility order between a read-write and a later read target architecture is x86, the compiler can recognize that the head of the order is a compare and exchange instruction and omit any extra synchronization:

$$[\text{RW}]; \text{svo}; [R] \rightsquigarrow [\text{RW}]; \text{po}; [R]$$

## C HISTORY SEMANTICS

The syntax of our formalization appears in Figure 18. We use  $n, l, i,$  and  $p$  to range over natural numbers, memory locations, unique memory access identifiers, and unique thread identifiers. We use  $m$  to represent one of the four access modes, P for plain, O for opaque, RA for release acquire and V for volatile. Note that RA writes are release writes and RA reads are acquire reads. Memory accesses,  $a,$  can take the form of reads,  $l_m,$  and writes,  $l_m := n$  with their accompanying modes as well as read-writes,  $\text{RW}(l, n).$

Memory events,  $h$  record the program's interactions with memory. Notably, we assume that the program and history can record specified visibility and specified push orders,  $\text{vo}(i, i)$  and  $\text{push}(i, i)$  before the execution of the related identifiers. For example, the program may use the labeling and ordering mechanism of Cray and Sullivan [cite]. We discuss the other events in detail below.

We use  $d$  to range over program generated events which are translated by the memory semantics into memory events and we use  $P$  to abstract over an expression language that can produce such events. We use  $H$  to represent a list of memory events, where  $H(h)$  means that  $h \in H.$

Finally we use  $i \xrightarrow{R}_H i$  to represent memory model relations for  $H.$  In what follows, the restriction on the trace order,  $\text{to},$  and the acyclicity requirements form the interface with the relations of our axiomatic model.

Program and history states transition together via the step relation defined in figure 19. The history transition semantics certifies program events  $d@p$  of the form  $i = a@p, i@p,$  or  $i$  to  $n@p$  and turns them into memory events.

$i = a@p$  represents the "initialization" of a memory access  $a$  using the unique identifier  $i$  in thread  $p.$  The history semantics appends  $\text{init}(i, p)$  and  $\text{is}(i, a)$  to  $H$  to record the initialization and the form of the memory access identified by  $i.$  The only certification required of an initialization is that the memory access identified by  $i$  is not already initialized. Importantly, we assume that the program initializes memory accesses in program order, so the subsequence of initialization events records program order in  $H.$

1324	Nat	$n := 0 \mid 1 \mid \dots$	Accesses	$a := l_m \mid l_m := n \mid \text{RW}(l, n)$
1325	Locations	$l := \dots$	Mem. Events	$h := \text{init}(i, p) \mid \text{is}(i, a) \mid \text{exec}(i)$
1326	Modes	$m := P \mid 0 \mid \text{RA} \mid V$		$\mid \text{rf}(i, i) \mid \text{vo}(i, i) \mid \text{push}(i, i)$
1327	Access Ids	$i := \dots$	Prog. Events	$d := i = a \mid i \mid i \text{ to } n$
1328	Thread Ids	$p := \dots$	Program	$P := \dots$
1329			History	$H := \epsilon \mid H, h$
1330				
1331				
1332				
1333				
1334				
1335				

Fig. 18. Syntax

$$\begin{array}{c}
\frac{P \xrightarrow{d@p} P' \quad H \xrightarrow{d@p} H'}{(P, H) \rightarrow (P', H')} \text{ step} \qquad \frac{\text{wf}(H, \text{exec}(i), p) \quad \text{acyclic}(\xrightarrow{\text{co}}_{H, \text{exec}(i)})}{H \xrightarrow{i@p} H, \text{exec}(i)} \text{ write} \\
\frac{\neg H(\text{init}(i, \_))}{H \xrightarrow{i=a@p} H, \text{init}(i, p), \text{is}(i, a)} \text{ init} \qquad \frac{\text{acyclic}(\xrightarrow{\text{co}}_{H, \text{rf}(i_w, i)}) \quad \text{wf}(H, \text{rf}(i_w, i), p, n) \quad \text{acyclic}(\xrightarrow{\text{po} \mid \text{rf}}_{H, \text{rf}(i_w, i)})}{H \xrightarrow{i \text{ to } n@p} H, \text{rf}(i_w, i)} \text{ read} \\
\text{reads}(H, i, l) \triangleq \exists m, H(\text{is}(i, l_m)) \vee \exists n, H(\text{is}(i, \text{RW}(l, n))) \\
\text{writes}(H, i, l, n) \triangleq \exists m, H(\text{is}(i, l_m := n)) \vee H(\text{is}(i, \text{RW}(l, n))) \\
\text{executed}(H, i) \triangleq H(\text{exec}(i)) \vee \exists i_w, H(\text{rf}(i_w, i)) \\
\text{executable}(H, i) \triangleq \neg \text{executed}(H, i) \wedge \forall i', i' \xrightarrow{\text{into}}_H i \implies \text{executed}(H, i') \\
\text{wf}(H, \text{exec}(i), p) \triangleq \begin{cases} H(\text{init}(i, p)) \\ H(\text{is}(i, l_m := \_)) \\ \text{executable}(H, i) \end{cases} \quad \text{wf}(H, \text{rf}(i_w, i), p, n) \triangleq \begin{cases} H(\text{init}(i, p)) \\ \text{reads}(H, i, l) \\ \text{writes}(H, i_w, l, n) \\ \text{executed}(H, i_w) \\ \text{executable}(H, i) \end{cases}
\end{array}$$

Fig. 19. Semantics

$i@p$  represents the execution of a write in thread  $p$  which is recorded in history with  $\text{exec}(i)$ . Writes must be certified by the acyclicity requirement for  $\text{co}$  and a basic well-formedness condition  $\text{wf}(H, \text{exec}(i), p)$ . The well-formedness requires that  $i$  be initialized  $H(\text{init}(i, p))$ , that the memory access associated with  $i$  be a write,  $H(\text{is}(i, l_m := n))$ , and that the write be executable  $\text{executable}(H, i)$ .

The  $\text{executable}(H, i)$  constraint requires that  $i$  has not already executed and that the sequence of execution, to, respects any intra-thread ordering,  $\text{into}$ , as in the description of Section 4.2. The program is otherwise free to execute memory accesses out-of-order.

$i \text{ to } n@p$  represents the execution of a read  $i$ , reading the value  $n$  in thread  $p$  which is recorded in history with  $\text{rf}(i_w, i)$ . Reads must be certified by the same acyclicity requirement for  $\text{co}$  and the additional requirement on  $\text{po} \mid \text{rf}$ . The well-formedness condition for reads requires that  $i$  be initialized, that it be a read or a read-write,  $\text{reads}(H, i, l)$ , and that it be executable. It further requires

1373 that the paired write  $i_w$  is executed,  $\text{executed}(H, i_w)$ , and that  $i_w$  writes the value  $n$  to the same  
 1374 location  $l$ ,  $\text{writes}(H, i, l, n)$ .

1375

## 1376 D OBSERVABLE TOTAL COHERENCE FOR ARMV8

1377 Here we demonstrate that it is possible to construct a program that is only forbidden due to the  
 1378 total coherence order of the ARMv8 memory model from Pulte et al. [2017]. We start by noticing  
 1379 that Herd model for ARMv8 has two acyclicity requirements that involve **co**:

1380 (\* Internal visibility requirement \*)

1381 acyclic po-loc | ca | rf **as** internal

1382

1383 (\* External visibility requirement \*)

1384 irreflexive ob **as** external

1385

1386 In the first requirement  $ca = fr \mid co$ . In the second  $ob = obs \mid \dots$  where  $obs = \dots \mid coe$   
 1387 and  $coe$  is coherence restricted to inter-thread relationships. Critically, as illustrated in the proof  
 1388 for our model, they do not together work to form cycles. So we can use one with each side of the  
 1389 total order to demonstrate its observability in the model.

1390 We have constructed the following program (included in the supplementary material) which will  
 1391 exhibit a cycle in the first requirement for one side of the total order and a different, incompatible  
 1392 cycle, for the second requirement:

1393 Arch64 totalco

```

1394 {
1395   0:X1=x; 0:X3=y;
1396   1:X1=x; 1:X3=y;
1397   2:X1=x; 2:X3=y;
1398 }
1399   P0           | P1           | P2;
1400   LDR X2,[X1]  | LDAR X5, [X3]| LDAR X5, [X1];
1401   MOV X0,#1   | MOV X2,#2   | MOV X0, #1;
1402   STR X0,[X1]  | STR X2,[X1] | STR X0, [X3];

```

1403 exists ( $0:X2=2 \wedge 1:X5=1 \wedge 2:X5=1$ )

1404 We show this by running the program with Herd and allowing executions that violate these two  
 1405 requirements (we mark them with “flag” in the Herd parlance). The result in Figure 20 is exactly  
 1406 two flagged executions. Each figure corresponds to one direction of the total ordering between  
 1407 b and d. We will inspect both to demonstrate that they are only forbidden as consequence of the  
 1408 total order, and thus if the total order was taken away they would both be allowed.

1409 In Figure 20a note the following. All of the cycles for any kind of edge involve a  $\xrightarrow{po-loc}$  b  
 1410 which is not in ob. This means we can avoid a cycle in ob. Recall that if there were a cycle in ob  
 1411 then when we use ob for the other side of the total coherence order it would be a cycle regardless  
 1412 of the total coherence order. Also, note that if we take away  $b \xrightarrow{co}$  d (blue) it will also remove  
 1413  $e \xrightarrow{fr}$  d. Thus, if we take away  $b \xrightarrow{co}$  d by removing the total coherence order of the model, there  
 1414 is no cycle left in the graph and the execution would be allowed.

1415 In Figure 20b note the following. All of the cycles for any kind of edge involve  $c \xrightarrow{bob}$  d which  
 1416 is not in po-loc | ca | rf. This means we can't establish a cycle in po-loc | ca | rf. Thus, if  
 1417 we take away  $d \xrightarrow{co}$  b by removing the total coherence order of the model, there is no cycle in ob  
 1418 left in the graph and the execution would be allowed.

1420

1421

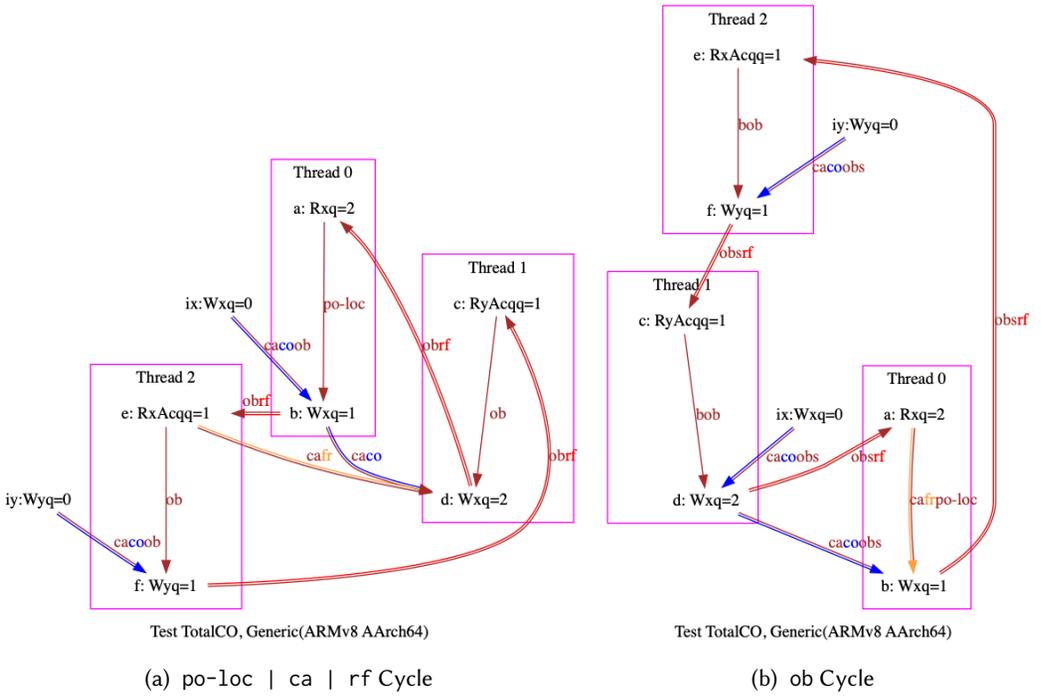


Fig. 20. Two Flagged Executions