

Number 745



UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

A better x86 memory model: x86-TSO (extended version)

Scott Owens, Susmit Sarkar, Peter Sewell

March 2009

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<http://www.cl.cam.ac.uk/>

© 2009 Scott Owens, Susmit Sarkar, Peter Sewell

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

<http://www.cl.cam.ac.uk/techreports/>

ISSN 1476-2986

A better x86 memory model: x86-TSO (extended version)

Scott Owens Susmit Sarkar Peter Sewell

University of Cambridge

<http://www.cl.cam.ac.uk/users/pes20/weakmemory>

March 25, 2009 *Revision* : 1746

Abstract

Real multiprocessors do not provide the sequentially consistent memory that is assumed by most work on semantics and verification. Instead, they have relaxed memory models, typically described in ambiguous prose, which lead to widespread confusion. These are prime targets for mechanized formalization. In previous work we produced a rigorous *x86-CC* model, formalizing the Intel and AMD architecture specifications of the time, but those turned out to be unsound with respect to actual hardware, as well as arguably too weak to program above. We discuss these issues and present a new *x86-TSO* model that suffers from neither problem, formalized in HOL4. We believe it is sound with respect to real processors, reflects better the vendor's intentions, and is also better suited for programming. We give two equivalent definitions of x86-TSO: an intuitive operational model based on local write buffers, and an axiomatic total store ordering model, similar to that of the SPARCv8. Both are adapted to handle x86-specific features. We have implemented the axiomatic model in our `memevents` tool, which calculates the set of all valid executions of test programs, and, for greater confidence, verify the witnesses of such executions directly, with code extracted from a third, more algorithmic, equivalent version of the definition.

Contents

1	Introduction	2
2	Many Memory Models	3
2.1	pre-IWP (before Aug. 2007)	3
2.2	IWP/AMD64-3.14/x86-CC	3
2.3	Intel SDM rev-29 (Nov. 2008)	4
3	The x86-TSO Model	5
3.1	The x86-TSO Abstract Machine Memory Model	6
3.2	The x86-TSO Axiomatic Memory Model	9
3.3	The Machine and Axiomatic x86-TSO Models are Equivalent	12
4	Verified Checker and Results	12
5	Related Work	12
6	Conclusion	13
A	Litmus Tests and Discussion	14
A.1	Load/Store Reordering	14
A.2	Independent Reads of Independent Writes	17
A.3	Locked Instructions: Tests iwp2.7/amd7, iwp2.8.a, iwp2.8.b, n8, and n3	18
A.4	Fence Instructions: Tests amd5 and amd10	19
A.5	The Unsoundness of IWP/AMD3.14/x86-CC: Test n6	20

A.6	The Weakness of Rev-29: Tests n4 and n5	20
A.7	Interpreting the rev-29 “not reordered with”: Test n7	21
A.8	Other Tests	22
A.9	Summary of Test Results	24
B	The IWP and Rev-29 Principles and Litmus Tests	25
B.1	Principles	25
B.2	Litmus Tests	25
C	Well-Formed Event Structures	26
D	Auxiliary Definitions	27
D.1	Axiomatic Memory Model	27
D.2	Abstract Machine Memory Model	28
E	Proof Outlines	31
E.1	Event-annotated machine	31
E.2	Linear valid executions	36
E.3	Abstract machine model/axiomatic model equivalence	38
E.3.1	Abstract machine validity	38
E.3.2	Abstract machine completeness	41
E.4	Executable checker	44
F	Change History	49

1 Introduction

Most previous research on the semantics and verification of concurrent programs assumes sequential consistency: that accesses by multiple threads to a shared memory occur in a global-time linear order. Real multiprocessors, however, incorporate many performance optimisations. These are typically unobservable by single-threaded programs, but some have observable consequences for the behaviour of concurrent code. For example, on standard Intel or AMD x86 processors, given two memory locations x and y (initially holding 0), if two processors proc:0 and proc:1 respectively write 1 to x and y and then read from y and x , as in the program below, it is possible for both to read 0 *in the same execution*.

iwp2.3.a/amd4	proc:0	proc:1
poi:0	MOV $[x] \leftarrow \$1$	MOV $[y] \leftarrow \$1$
poi:1	MOV EAX $\leftarrow [y]$	MOV EBX $\leftarrow [x]$
Allow: $0:\text{EAX}=0 \wedge 1:\text{EBX}=0$		

One can view this as a visible consequence of *write buffering*: each processor effectively has a FIFO buffer of pending memory writes (to avoid the need to block while a write completes), so the reads from y and x can occur before the writes have propagated from the buffers to main memory. Such optimisations destroy the illusion of sequential consistency, making it impossible (at this level of abstraction) to reason in terms of an intuitive notion of global time.

To describe what programmers can rely on, processor vendors document *architectures*. These are loose specifications, claimed to cover a range of past and future actual processors, which should reveal enough for effective programming, but without unduly constraining future processor designs. In practice, however, they are informal prose documents, e.g. the Intel 64 and IA-32 Architectures SDM [2] and AMD64 Architecture Programmer’s Manual [1]. Informal prose is a poor medium for loose specification of subtle properties, and, as we shall see in §2, such documents are often ambiguous, are sometimes incomplete (too weak to program above), and are sometimes unsound (with respect to the actual processors). Moreover, one cannot test programs above such a vague specification (one can only run programs on particular actual processors), and one cannot use them as criteria for testing processor implementations.

Architecture specifications are, therefore, prime targets for rigorous mechanised formalisation. In previous work [20] we introduced a rigorous x86-CC model, formalised in HOL4 [11], based on the

informal prose causal-consistency descriptions of the then-current Intel and AMD documentation. Unfortunately those, and hence also x86-CC, turned out to be unsound, forbidding some behaviour which actual processors exhibit.

In this paper we describe a new model, x86-TSO, also formalised in HOL4. To the best of our knowledge, x86-TSO is sound, is strong enough to program above, and is broadly in line with the vendors’ intentions. We present two equivalent definitions of the model: an abstract machine, in §3.1, and an axiomatic version, in §3.2. We compensate for the main disadvantage of formalisation, that it can make specifications less widely accessible, by extensively annotating the mathematical definitions. To explore the consequences of the model, we have a hand-coded implementation in our `memevents` tool, which can explore all possible executions of litmus-test examples such as that above, and for greater confidence we have a verified execution checker extracted from the HOL4 axiomatic definition, in §4. We discuss related work in §5 and conclude in §6.

Further details can be found in the appendices. In Appendix A we discuss a number of litmus tests, comparing their behaviour in the model (as calculated by `memevents`) and on actual processors (as observed with our `litmus` tool). In Appendix B we reproduce the key “principles” from the informal-prose vendor documentation, for reference. Appendix C makes the notion of well-formed event structure precise, Appendix D collects a number of routine auxiliary definitions used in the definition of our models, and Appendix E gives an overview of the proofs of our results.

2 Many Memory Models

We begin by reviewing the informal-prose specifications of recent Intel and AMD documentation. There have been several versions, some differing radically; we contrast them with each other, and with what we know of the behaviour of actual processors.

2.1 pre-IWP (before Aug. 2007)

Early revisions of the Intel SDM (e.g. rev-22, Nov. 2006) gave an informal-prose model called ‘processor ordering’, unsupported by any litmus-test examples. It is hard to give a precise interpretation of this description.

2.2 IWP/AMD64-3.14/x86-CC

In August 2007, an Intel White Paper [12] (IWP) gave a somewhat more precise model, with 8 informal-prose principles supported by 10 litmus tests. This was incorporated, essentially unchanged, into later revisions of the Intel SDM (including rev.26–28), and AMD gave similar, though not identical, prose and tests [1]. These are essentially causal-consistency models. They allow independent readers to see independent writes (by different processors to different addresses) in different orders, as below (IRIW, see also [5]),

amd6	proc:0	proc:1	proc:2	proc:3
poi:0	MOV [x]←\$1	MOV [y]←\$1	MOV EAX←[x]	MOV ECX←[y]
poi:1			MOV EBX←[y]	MOV EDX←[x]
Final: 2:EAX=1 ∧ 2:EBX=0 ∧ 3:ECX=1 ∧ 3:EDX=0				
cc : Allow; tso : Forbid				

but require that, in some sense, causality is respected: “P5. In a multiprocessor system, memory ordering obeys causality (memory ordering respects transitive visibility)”. These were the basis for our x86-CC model, for which a key issue was giving a reasonable interpretation to this “causality”. Apart from that, the informal specifications were reasonably unambiguous — but they turned out to have two serious flaws.

First, they are arguably rather weak for programmers. In particular, they admit the IRIW behaviour above but, under reasonable assumptions on the strongest x86 memory barrier, MFENCE, adding MFENCES would not suffice to recover sequential consistency [20, §2.12]. Here the specifications seem to be much looser than the behaviour of implemented processors: to the best of our knowledge, and following some testing, IRIW is not observable in practice. It appears that some JVM implementations depend on this fact, and would not be correct if one assumed only the IWP/AMD64-3.14/x86-CC architecture [9].

Second, more seriously, they are unsound with respect to current processors. The following n6 example, due to Paul Loewenstein [14], shows a behaviour that is observable (e.g. on an Intel Core 2 duo), but that is disallowed by x86-CC, and by any interpretation we can make of IWP and AMD64-3.14.

n6	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [y]←\$2
poi:1	MOV EAX←[x]	MOV [x]←\$2
poi:2	MOV EBX←[y]	
Final: 0:EAX=1 ∧ 0:EBX=0 ∧ [x]=1		
cc : Forbid; tso : Allow		

To see why this may be allowed by multiprocessors with FIFO write buffers, suppose that first the proc:1 write of [y]=2 is buffered, then proc:0 buffers its write of [x]=1, reads [x]=1 from its own write buffer, and reads [y]=0 from main memory, then proc:1 buffers its [x]=2 write and flushes its buffered [y]=2 and [x]=2 writes to memory, then finally proc:0 flushes its [x]=1 write to memory.

2.3 Intel SDM rev-29 (Nov. 2008)

The most recent x86 vendor specification, at the time of writing, is revision 29 of the Intel SDM (we are told that there will be a future revision of the AMD specification on similar lines). Key extracts are summarised in Appendix B. This is in a similar informal-prose style to previous versions, again supported by litmus tests, but is significantly different to IWP/AMD64-3.14/x86-CC. First, the IRIW final state above is forbidden [Example 7-7, rev-29], and the previous coherence condition: “P6. In a multiprocessor system, stores to the same location have a total order” has been replaced by: “P9. Any two stores are seen in a consistent order by processors other than those performing the stores”.

Second, the memory barrier instructions are now included, with “P11. Reads cannot pass LFENCE and MFENCE instructions” and “P12. Writes cannot pass SFENCE and MFENCE instructions”.

Third, same-processor writes are now explicitly ordered (we regarded this as implicit in the IWP “P2. Stores are not reordered with other stores”): “P10. Writes by a single processor are observed in the same order by all processors”.

This specification appears to deal with the unsoundness, admitting the n6 behaviour above, but, unfortunately, it is still problematic. The first issue is, again, how to interpret “causality” as used in P5. The second issue is one of weakness: the new P9 says nothing about observations of two stores by those two processors themselves (or by one of those processors and one other). This is illustrated by the following n5 and n4 examples. These final states were not allowed in x86-CC, and we would be surprised if they were allowed by any reasonable implementation (they are not allowed in a pure write-buffer implementation). We have not observed them on actual processors, and programming above a model that permitted them would be problematic. However, rev-29 appears to allow them.

n5	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [x]←\$2
poi:1	MOV EAX←[x]	MOV EBX←[x]
Forbid: 0:EAX=2 ∧ 1:EBX=1		

n4	proc:0	proc:1
poi:0	MOV EAX←[x]	MOV ECX←[x]
poi:1	MOV [x]←\$1	MOV [x]←\$2
poi:2	MOV EBX←[x]	MOV EDX←[x]
Forbid: 0:EAX=2 ∧ 0:EBX=1 ∧ 1:ECX=1 ∧ 1:EDX=2		

Summarising the key litmus-test differences, we have:

	IWP/AMD64-3.14/x86-CC	rev-29	actual processors
IRIW	allowed	forbidden	not observed
n6	forbidden	allowed	observed
n4/n5	forbidden	allowed	not observed

There are also many non-differences: tests for which the behaviours coincide in all three cases. The test details can be found in Appendix A. They include the 9 other IWP tests, illustrating that various load and store reorderings other than that shown in iwpt2.3.a/amd4 (§1) are not possible; the AMD MFENCE tests amd5 and amd10, and several other tests.

3 The x86-TSO Model

Given these problems with the informal specifications, we cannot produce a useful rigorous model by formalising the “principles” they contain (as we attempted with x86-CC [20]). Instead, we have to build a reasonable model that is consistent with the given litmus tests, observable processor behaviour, and with what we know of the needs of programmers and of the vendors intentions.

The fact that write buffering is observable (iwp2.3.a/amd4 and n6) but IRIW is not, together with the other tests that prohibit many other reorderings, strongly suggests that, apart from write buffering, all processors share the same view of memory (in contrast to x86-CC, where each processor had a separate view order). This is broadly similar to the SPARC Total Store Ordering (TSO) memory model [21, 23], which is essentially an axiomatic description of the behaviour of write-buffer multiprocessors. Moreover, while the term “TSO” is not used, informal discussions suggest this matches the intention behind the rev.29 informal specification. Accordingly, we define here a rigorous x86-TSO model.

After some preliminaries, we give two equivalent definitions of x86-TSO. The first, in §3.1, is an abstract machine, with explicit write buffers. The second, in §3.2, is an axiomatic model, defining valid executions in terms of memory orders and reads-from maps. In both, we deal with the x86 CISC instructions with multiple memory accesses, with x86 LOCK’d instructions (CMPXCHG, LOCK;INC, etc.), with potentially non-terminating computations, and with dependencies through registers. Together with our earlier instruction semantics, x86-TSO thus defines a complete semantics of programs.

The intended scope of x86-TSO, as for the x86-CC model, covers typical user code and most kernel code: programs using coherent write-back memory, without exceptions, misaligned or mixed-size accesses, ‘non-temporal’ operations (e.g. MOVNTI), self-modifying code, or page-table changes.

Basic Types: Actions, Events, and Event Structures As in our earlier work, the action of (any particular execution of) a program is abstracted into a set of *events* (with additional data) called an *event structure*. An event represents a read or write of a particular value to a memory address, or to a register, or the execution of a fence. Our earlier work includes a definition of the set of event structures generated by an assembly language program. For any such event structure, the memory model (there x86-CC, here x86-TSO) defines what a *valid execution* is.

In more detail, each machine-code instruction may have multiple events associated with it: events are indexed by an instruction ID *iiid* that identifies which processor the event occurred on and the position in the instruction stream of the instruction it comes from (the *program order index*, or *poi*). Events also have an event ID *eiid* to identify them within an instruction (to permit multiple, otherwise identical, events). An event structure indicates when one of an instruction’s events has a dependency on another event of the same instruction with an *intra-causality* relation, a partial order over the events of each instruction. An event structure also records which events occur together in a locked instruction with *atomicity* data, a set of (disjoint, non-empty) sets of events which must occur atomically together.

Expressing this in HOL, we index processors by a type `proc = num`, take types `address` and `value` to both be the 32-bit words, and take a location to be either a memory address or a register of a particular processor:

```
location = LOCATION_REG of proc 'reg
          | LOCATION_MEM of address
```

The model is parameterised by a type `'reg` of x86 registers, which one should think of as an enumeration of the names of ordinary registers EAX, EBX, etc., the instruction pointer EIP, and the status flags. To identify an instance of an instruction in an execution, we specify its processor and its program order index.

```
iiid = ⟨ proc : proc; poi : num ⟩
```

This introduces a type of records with two fields, a *proc* of type `proc` and a program order index *poi* of type `num`. An action is either a read or write of a value at some location, or a barrier:

```
dirn = R | W
barrier = LFENCE | SFENCE | MFENCE
action = ACCESS of dirn ('reg location) value | BARRIER of barrier
```

Finally, an event has an instruction instance id, an event id (of type `eiid = num`, unique per `iiid`), and an action:

$\text{event} = \langle \langle \text{eiid} : \text{eiid}; \text{iiid} : \text{iiid}; \text{action} : \text{action} \rangle \rangle$


and an event structure E comprises a set of processors, a set of events, an intra-instruction causality relation, and a partial equivalence relation (PER) capturing sets of events which must occur atomically, all subject to some well-formedness conditions which we omit here.

$\text{event_structure} = \langle \langle \text{procs} : \text{proc set};$
 $\text{events} : ('reg \text{ event})\text{set};$
 $\text{intra_causality} : ('reg \text{ event})\text{reln};$
 $\text{atomicity} : ('reg \text{ event})\text{set set} \rangle \rangle$

Example We show a very simple event structure below, for the program:

tso1	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [x]←\$2
poi:1	MOV EAX←[x]	

There are four events — the inner (blue) boxes. The event ids are pretty-printed alphabetically, as a,b,c,d, etc. We also show the assembly instruction that gave rise to each event, e.g. MOV [x]←\$1, though that is not formally part of the event structure.



Note that events contain concrete values: in this particular event structure, there are two writes of x , with values 1 and 2, a read of $[x]$ with value 2, and a write of proc:0 's EAX register with value 2. Later we show two valid executions for this program, one for this event structure and one for another (note also that some event structures may not have any valid executions).

In the diagram, the instructions of each processor are clustered together, into the outermost (magenta) boxes, with program order (po) edges between them, and the events of each instruction are clustered together into the intermediate (green) boxes, with intra-causality edges as appropriate — here, in the MOV EAX←[x], the write of EAX is dependent on the read of x .

3.1 The x86-TSO Abstract Machine Memory Model

To understand our x86-TSO machine model, consider an idealised x86 multiprocessor system partitioned into two components: its memory and register state (of all its processors combined), and the rest of the system (the other parts of all the processor cores). Our abstract machine is a labelled transition system: a set of states, ranged over by s , and a transition relation $s \xrightarrow{l} s'$. An abstract machine state s models the state of the first component, the memory and register state of a multiprocessor system, and the machine interacts with the rest of the system by synchronising on labels l (the interface of the abstract machine), which include register and memory reads and writes. One should think of the machine as operating in parallel with the processor cores (absent their register/memory subsystems), executing their instruction streams in program order; the latter data is provided by an event structure. This partitioning does not correspond directly to the microarchitecture of any realistic x86 implementation, in which memory and registers would be implemented by separate and intricate mechanisms, but it is useful and sufficient for describing the programming model, which is the proper business of an architecture description. It also supports a precise correspondence with our axiomatic memory model. In more detail, the labels l are the values of the HOL type:

$\text{label} = \text{TAU} \mid \text{EVT of proc } ('reg \text{ action}) \mid \text{LOCK of proc} \mid \text{UNLOCK of proc}$

- TAU, for an internal action by the machine;

Read from memory

$$\frac{\text{not_blocked } s \ p \wedge (s.M \ a = \text{SOME } v) \wedge \text{no_pending } (s.B \ p) \ a}{s \xrightarrow{\text{EVT } p \ (\text{ACCESS } R \ (\text{LOCATION_MEM } a) \ v)} s}$$

Read from write buffer

$$\frac{\text{not_blocked } s \ p \wedge (\exists b_1 \ b_2. (s.B \ p = b_1 \ ++ [(a, v)] \ ++ b_2) \wedge \text{no_pending } b_1 \ a)}{s \xrightarrow{\text{EVT } p \ (\text{ACCESS } R \ (\text{LOCATION_MEM } a) \ v)} s}$$

Read from register

$$\frac{(s.R \ p \ r = \text{SOME } v)}{s \xrightarrow{\text{EVT } p \ (\text{ACCESS } R \ (\text{LOCATION_REG } p \ r) \ v)} s}$$

Write to write buffer

$$\frac{\mathbf{T}}{s \xrightarrow{\text{EVT } p \ (\text{ACCESS } W \ (\text{LOCATION_MEM } a) \ v)} s \oplus \langle B := s.B \oplus (p \mapsto [(a, v)] \ ++ (s.B \ p)) \rangle}$$

Write from write buffer to memory

$$\frac{\text{not_blocked } s \ p \wedge (s.B \ p = b \ ++ [(a, v)])}{s \xrightarrow{\text{TAU}} s \oplus \langle M := s.M \oplus (a \mapsto \text{SOME } v); B := s.B \oplus (p \mapsto b) \rangle}$$

Write to register

$$\frac{\mathbf{T}}{s \xrightarrow{\text{EVT } p \ (\text{ACCESS } W \ (\text{LOCATION_REG } p \ r) \ v)} s \oplus \langle R := s.R \oplus (p \mapsto ((s.R \ p) \oplus (r \mapsto \text{SOME } v))) \rangle}$$

Barrier

$$\frac{(b = \text{MFENCE}) \implies (s.B \ p = [])}{s \xrightarrow{\text{EVT } p \ (\text{BARRIER } b)} s}$$

Lock

$$\frac{(s.L = \text{NONE}) \wedge (s.B \ p = [])}{s \xrightarrow{\text{LOCK } p} s \oplus \langle L := \text{SOME } p \rangle}$$

Unlock

$$\frac{(s.L = \text{SOME } p) \wedge (s.B \ p = [])}{s \xrightarrow{\text{UNLOCK } p} s \oplus \langle L := \text{NONE} \rangle}$$

Figure 1: The x86-TSO Machine Behaviour

- EVT p a , where a is an action, as defined above (a memory or register read or write, with its value, or a barrier), by processor p ;
- LOCK p , indicating the start of a LOCK'd instruction by processor p ; or
- UNLOCK p , for the end of a LOCK'd instruction by p .

(Note that there is nothing specific to any particular memory model in this interface.) The states of the x86-TSO machine are records, with fields R , giving a value for each register on each processor; M , giving a value for each shared memory location; B , modelling a write buffer for each processor, as a list of address/value pairs; and L , which is a global lock, either SOME p , if p holds the lock, or NONE. The HOL type is below.

```
machine_state = ⟨⟨  $R : \text{proc} \rightarrow 'reg \rightarrow \text{value option};$  (* per-processor registers *)
                  $M : \text{address} \rightarrow \text{value option};$  (* main memory *)
                  $B : \text{proc} \rightarrow (\text{address}\#\text{value})\text{list};$  (* per-processor write buffers *)
                  $L : \text{proc option}(* \text{ which processor holds the lock }*) \rangle \rangle$ 
```

The behaviour of the x86-TSO machine, the transition relation $s \xrightarrow{L} s'$, is defined by the rules in Fig. 1. The rules use two auxiliary definitions: processor p is *not blocked* in machine state s if either it holds the lock or no processor does; and there are *no pending* writes in a buffer b for address a if there are no (a, v) pairs in b . Restating the rules informally:

1. p can read v from memory at address a if p is not blocked, has no buffered writes to a , and the memory does contain v at a ;
2. p can read v from its write buffer for address a if p is not blocked and has v as the newest write to a in its buffer;
3. p can read the stored value v from its register r at any time;
4. p can write v to its write buffer for address a at any time;
5. if p is not blocked, it can silently dequeue the oldest write from its write buffer to memory;
6. p can write value v to one of its registers r at any time;
7. if p 's write buffer is empty, it can execute an MFENCE (so an MFENCE cannot proceed until all writes have been dequeued, modelling buffer flushing); LFENCE and SFENCE can occur at any time, making them no-ops;
8. if the lock is not held, and p 's write buffer is empty, it can begin a LOCK'd instruction; and
9. if p holds the lock, and its write buffer is empty, it can end a LOCK'd instruction.

We emphasise that this is an *abstract* machine: we are concerned with its extensional behaviour, the (completed, finite or infinite) traces of labelled transitions it can perform (which should include the behaviour of real implementations), not with its internal states and the transition rules. The machine should provide a good model for programmers, but may bear little resemblance to the internal structure of implementations. Indeed, a realistic design would certainly not implement LOCK'd instructions with a global lock, and would have many other optimisations — the force of the x86-TSO model is that none of those have *programmer-visible* effects, except perhaps via performance observations.

One can imagine several variants of the machine with different degrees of locking. We conjecture that one could additionally make the read-register, write-register, write-buffer, and barrier transitions dependent on a `not_blocked` premise to arrive at a stricter, but observationally equivalent, machine, that might be simpler for programmers to reason with. However, the additions would make the proof of equivalence to the axiomatic memory model more difficult. We also conjecture that removing the `not_blocked` premise from the read-memory transition, and removing the requirement that the buffer is empty on a Lock label, would give a more liberal, but equivalent, machine. Again the equivalence proof would be more difficult. Should either of these variant machines become useful, we anticipate being able to prove them equivalent by working solely on the machines.

We relate the machine to event structures in two steps, which we summarise here (the HOL details can be found on-line [22] and are summarised in Appendix E). First, we define a more intensional event-machine: we annotate each memory and register location with an event option, recording the most recent write event (if any) to that location, refine write buffers to record lists of events rather than of plain location/value pairs, and annotate labels with the relevant events.

Theorem 1 *The annotation-erasure of the event-machine is exactly the machine presented above.*

[HOL proof]

Let *path* range over finite or infinite sequences of states and labels $s_0 \xrightarrow{l_1} s_1 \xrightarrow{l_2} s_2 \dots$. We define (okMpath *path*) to hold if *path* is a path through the event-machine, the first state is an initial machine state, with empty write buffers, etc., and a progress condition holds: for each memory write in *path*, the corresponding TAU is eventually performed. Finally, given an event structure *E*, we say such a *path* is a machine execution corresponding to that event structure if (okEpath *E path*), i.e., if there is a 1:1 correspondence between non-TAU/LOCK/UNLOCK labels of *path* and the events of *E*, consistent with program order and intra-causality, and atomic sets are properly bracketed by LOCK/UNLOCK pairs.

3.2 The x86-TSO Axiomatic Memory Model

Our x86-TSO axiomatic memory model is based on the SPARCV8 memory model specification [21, 23], but adapted to x86 and in the same terms as our earlier x86-CC model. Compared with the SPARCV8 TSO specification, we omit instruction fetches (*IF*), instruction loads (*IL*), flushes (*F*), and stbars (*S*). The first three deal exclusively with instruction memory, which we do not model, and the last is useful only under the SPARC PSO memory model. To adapt it to x86 programs, we add register and fence events, generalize to support instructions that give rise to many events (partially ordered by an intra-instruction causality relation), and generalize atomic load/store pairs to locked instructions.

An execution is permitted by our memory model if there exists an *execution witness* *X* for its event structure *E* that is a *valid execution*. An execution witness contains a *memory_order*, an *rfmap*, and an *initial_state*; the rest of this section defines when these are valid.

execution_witness =

```
⟦ memory_order : ('reg event)reln;
  rfmap : ('reg event)reln;
  initial_state : ('reg location → value option) ⟧
```

The memory order is a partial order that records the global ordering of memory events. It must be a total order on memory writes, and corresponds to the $\leq$ relation in SPARCV8, as constrained by the SPARCV8 **Order** condition (in figures, we use the label mo_non-po_write_write for the otherwise-unforced part of this order).

partial_order ($<_{X.memory_order}$)(mem_accesses *E*)

linear_order ($(<_{X.memory_order})|_{(mem_writes\ E)}$)(mem_writes *E*)

The initial state is a partial function from locations to values. Each read event's value must come either from the initial state or from a write event: the rfmap ('reads-from map') records which, containing (*ew*, *er*) pairs where the read *er* reads from the write *ew*. The *reads_from_map_candidates* predicate below ensures that the rfmap only relates such pairs with the same address and value. (The SPARCV8 model does not have an explicit representation of reads-from maps, and does not deal with initial states.)

reads_from_map_candidates *E* rfmap =

$$\forall (ew, er) \in rfmap. (er \in reads\ E) \wedge (ew \in writes\ E) \wedge$$

$$(loc\ ew = loc\ er) \wedge (value_of\ ew = value_of\ er)$$

We lift program order from instructions to a relation po_iico *E* over events, taking the union of program order of instructions and intra-instruction causality. This corresponds roughly to the ; in SPARCV8. However, *intra_causality* might not relate some pairs of events in an instruction, so our po_iico *E* will not generally be a total order for the events of a processor.

po_strict *E* =

$$\{(e_1, e_2) \mid (e_1.iid.proc = e_2.iid.proc) \wedge e_1.iid.poi < e_2.iid.poi \wedge e_1 \in E.events \wedge e_2 \in E.events\}$$

$$<_{(po_iico\ E)} = po_strict\ E \cup E.intra_causality$$

The *check_rfmap_written* below ensures that the rfmap relates a read to the most recent preceding write. For a register read, this is the most recent write in program order. For a memory read, this is the most recent write in memory order among those that precede the read in either memory order or program order (intuitively, the first case is a read of a committed write and the second is a read from the local write buffer). The *check_rfmap_written* and *reads_from_map_candidates* predicates implement the SPARCv8 **Value** axiom above the rfmap witness data. The *check_rfmap_initial* predicate extends this to handle initial state, ensuring that any read not in the rfmap takes its value from the initial state, and that that read is not preceded by a write in memory order or program order.

$$previous_writes\ E\ er\ <_{order} =$$

$$\{ew' \mid ew' \in writes\ E \wedge ew' <_{order}\ er \wedge (loc\ ew' = loc\ er)\}$$

$$check_rfmap_written\ E\ X =$$

$$\forall(ew, er) \in (X.rfmap).$$

$$\text{if } ew \in mem_accesses\ E \text{ then}$$

$$ew \in maximal_elements\ (previous_writes\ E\ er\ (<_{X.memory_order}) \cup previous_writes\ E\ er\ (<_{(po_iico\ E)}))$$

$$(<_{X.memory_order})$$

$$\text{else } (*\ ew\ IN\ reg_accesses\ E\ *)$$

$$ew \in maximal_elements\ (previous_writes\ E\ er\ (<_{(po_iico\ E)}))(<_{(po_iico\ E)})$$

$$check_rfmap_initial\ E\ X =$$

$$\forall er \in (reads\ E \setminus range\ X.rfmap).$$

$$(\exists l.(loc\ er = SOME\ l) \wedge (value_of\ er = X.initial_state\ l)) \wedge$$

$$(previous_writes\ E\ er\ (<_{X.memory_order}) \cup$$

$$previous_writes\ E\ er\ (<_{(po_iico\ E)}) = \{\})$$

We now further constrain the memory order, to ensure that it respects the relevant parts of program order, and that the memory accesses of a LOCK'd instruction do occur atomically.

- Program order is included in memory order, for a memory read before a memory access (labelled *mo_po_read_access* in figures) (SPARCv8's **LoadOp**):

$$\forall er \in (mem_reads\ E). \forall e \in (mem_accesses\ E).$$

$$er <_{(po_iico\ E)} e \implies er <_{X.memory_order} e$$

- Program order is included in memory order, for a memory write before a memory write (*mo_po_write_write*) (the SPARCv8 **StoreStore**):

$$\forall ew_1\ ew_2 \in (mem_writes\ E).$$

$$ew_1 <_{(po_iico\ E)} ew_2 \implies ew_1 <_{X.memory_order} ew_2$$

- Program order is included in memory order, for a memory write before a memory read, *if* there is an MFENCE between (*mo_po_mfence*). (There is no need to include fence events themselves in the memory ordering.)

$$\forall ew \in (mem_writes\ E). \forall er \in (mem_reads\ E). \forall ef \in (mfences\ E).$$

$$(ew <_{(po_iico\ E)} ef \wedge ef <_{(po_iico\ E)} er) \implies ew <_{X.memory_order} er$$

- Program order is included in memory order, for any two memory accesses where at least one is from a LOCK'd instruction (*mo_po_access/lock*):

$$\forall e_1\ e_2 \in (mem_accesses\ E). \forall es \in (E.atomicity).$$

$$((e_1 \in es \vee e_2 \in es) \wedge e_1 <_{(po_iico\ E)} e_2) \implies e_1 <_{X.memory_order} e_2$$

- The memory accesses of a LOCK'd instruction occur atomically in memory order (mo_atomicity), i.e., there must be no intervening memory events. Further, all program order relationships between the locked memory accesses and other memory accesses are included in the memory order (this is a generalization of the SPARCV8 **Atomicity** axiom):

$$\begin{aligned} & \forall es \in (E.\text{atomicity}). \forall e \in (\text{mem_accesses } E \setminus es). \\ & (\forall e' \in (es \cap \text{mem_accesses } E). e <_{X.\text{memory_order}} e') \vee \\ & (\forall e' \in (es \cap \text{mem_accesses } E). e' <_{X.\text{memory_order}} e) \end{aligned}$$

To deal properly with infinite executions, we also require that the prefixes of the memory order are all finite, ensuring that there are no limit points, and, to ensure that each write eventually takes effect globally, there must not be an infinite set of reads unrelated to any particular write, all on the same memory location (this formalizes the SPARCV8 **Termination** axiom).

$$\text{finite_prefixes } (<_{X.\text{memory_order}})(\text{mem_accesses } E)$$

$$\begin{aligned} & \forall ew \in (\text{mem_writes } E). \\ & \text{finite}\{er \mid er \in E.\text{events} \wedge (\text{loc } er = \text{loc } ew) \wedge \\ & \quad er \not<_{X.\text{memory_order}} ew \wedge ew \not<_{X.\text{memory_order}} er\} \end{aligned}$$

A final state of a valid execution takes the last write in memory order for each memory location, together with a maximal write in program order for each register (or the initial state, if there is no such write). This is uniquely defined assuming that no instruction has multiple unrelated writes to the same register — a reasonable property for x86 instructions.

The definition of valid_execution $E \ X$ comprising the above conditions is equivalent to one in which $<_{X.\text{memory_order}}$ is required to be a linear order, not just a partial order (again, the full details are on-line):

Theorem 2

1. If linear_valid_execution $E \ X$ then valid_execution $E \ X$.
2. If valid_execution $E \ X$ then there exists an $\hat{X}$ with a linearisation of X 's memory order such that linear_valid_execution $E \ \hat{X}$. [HOL proof]

Interpreting “not reordered with” Perhaps surprisingly, the above definition does not require that program order is included in memory order for a memory write followed by a read from the same address. The definition does imply that any such read cannot be speculated before the write (by check_rfnmap_written, as that takes both $<_{(\text{po_iico } E)}$ and $<_{X.\text{memory_order}}$ into account). However, if one included a memory order edge, perhaps following a naive interpretation of the rev-29 “P4. Reads may be reordered with older writes to different locations but not with older writes to the same location”, then the model would be strictly stronger: the n7 example below would become forbidden, whereas it is allowed on x86-TSO. We conjecture that this would correspond to the (rather strange) machine with the Fig. 1 rules but without the read-from-write-buffer rule, in which any processor would have to flush its write buffer up to (and including) a local write before it can read from it.

n7	proc:0	proc:1	proc:2
poi:0	MOV [x]←\$1	MOV [y]←\$1	MOV ECX←[y]
poi:1	MOV EAX←[x]		MOV EDX←[x]
poi:2	MOV EBX←[y]		
Allow: 0:EAX=1 ∧ 0:EBX=0 ∧ 2:ECX=1 ∧ 2:EDX=0			

Examples We show two valid executions of the previous example program in Fig. 2. In both executions, the proc:0 W x=1 event is before the proc:1 W x=2 event in memory order (the bold mo_non-po_write_write edge). In the first execution, on the left, the proc:0 read of x reads from the most recent write in memory order (the combination of the bold mo_non-po_write_write edge and the mo_rf edge), which is the proc:1 W x=2. In the second execution, on the right, the proc:0 read of x reads from the most recent write in program order, which is the proc:0 W x=1. This example also illustrates some register events: the MOV EAX←[x] instruction gives rise to a memory read of x, followed by (in the intra-instruction causality relation) a register write of EAX.




Figure 2: Example valid execution witnesses (for two different event structures)

3.3 The Machine and Axiomatic x86-TSO Models are Equivalent

To show these two definitions equivalent, given an event-machine *path*, we first build an execution witness *path_to_X path*, putting the memory reads and write-buffer flushes (TAUS) in their (linear) order from *path*. Now:

Theorem 3 *For any well-formed event structure E and event-machine path $path$, if $\text{okEpath } E \text{ path}$ and $\text{okMpath } path$, then $path_to_X \text{ path}$ is a valid execution for E .* [HOL proof]

Theorem 4 *For any well-formed event structure E , and valid execution X for E , there exists some event-machine path, such that $\text{okEpath } E \text{ path}$ and $\text{okMpath } path$, in which the memory reads and write-buffer flushes both respect $<_{X.\text{memory_order}}$.* [hand proof, with some parts in HOL]

4 Verified Checker and Results

To explore the consequences of x86-TSO, we implemented the axiomatic model in our **memevents** tool, which exhaustively explores candidate execution witnesses. For greater confidence, we added to this a verified witness checker: we defined variants of event structures and execution witnesses, using lists instead of sets, wrote algorithmic versions of `well_formed_event_structure` and `valid_execution`, proved these equivalent (in the finite case) to our other definitions, extracted OCaml code from the HOL, and integrated that into **memevents**. (Obviously, this only provides assurance for positive tests, those with allowed final states.)

The **memevents** results coincide with our observations on real processors and the vendor specifications, for the 10 IWP tests, the (negated) IRIW test, the two MFENCE tests `amd5` and `amd10`, our `n2-n6` and `n8`, and `rcw-fenced`. The remaining tests (`amd3`, `n1`, `n7`, and `rcw-unfenced`) are “allow” tests for which we have not observed the specified final state in practice.

5 Related Work

There is an extensive literature on relaxed memory models, but most of it does not address x86, and we are not aware of any previous model that addresses the concerns of §2. We touch here on some of the most closely related work.

There are several surveys of weak memory models, including those by Adve and Gharachorloo [3], by Luchango [16], and by Higham, Kawash, and Verwaal [13]. The latter, in particular, formalises a range of models, including a TSO model, in both operational and axiomatic styles, and proves equivalence results. Their axiomatic TSO model is in a different style to ours, rather closer to the machine behaviour, and idealised rather than x86-specific. Burckhardt and Musuvathi [7, Appendix A] also give operational and axiomatic definitions of a TSO model and prove equivalence, but only for finite executions. Their models treat memory read, memory write and barrier events, but lack register events and locked instructions that force multiple events to happen atomically. Hangel et al. [10] describe the Sun TSOtool, checking

the observed behaviour of pseudo-randomly generated programs against a TSO model. Roy et al. [17] describe an efficient algorithm for checking whether an execution lies within an approximation to a TSO model, used in Intel’s Random Instruction Test (RIT) generator. Boudol and Petri [6] give an operational model with hierarchical write buffers (thereby permitting IRIW behaviours), and prove sequential consistency for data-race-free (DRF) programs. Burckhardt et al. [8] define an x86 memory model based on IWP [12] (together with two MS CLR models). The mathematical form of their definitions is rather different to ours, using rewrite rules to re-order or eliminate memory accesses in sets of traces. Their model validates the 10 IWP tests and also some instances of IRIW (depending on how parallel compositions are associated), so it will not coincide with x86-TSO or x86-CC. Loewenstein et al. [15] describe a “golden memory model” for SPARC TSO, somewhat closer to a particular implementation microarchitecture than the abstract machine we give in §3.1, that they use for testing implementations. They argue that the additional intensional detail increases the effectiveness of simulation-based verification. Roychoudhury [18] describes a system for exhaustive search of executions in a logic programming system for versions of SPARC TSO and the Java Memory Model. Saraswat et al. [19] also define memory models in terms of local reordering, and prove a DRF theorem, but focus on high-level languages rather than processors. Several groups have used proof tools to tame the intricacies of these models, including Yang et al. [24], using Prolog and SAT solvers to explore an axiomatic Itanium model, and, turning briefly to high-level languages, Aspinall and Sevcik [4], who formalised and identified problems with the Java Memory Model using Isabelle/HOL.

6 Conclusion

We have described x86-TSO, a memory model for x86 processors that does not suffer from the ambiguities, weaknesses, or unsoundnesses of earlier models. Its abstract-machine definition should be intuitive for programmers, whereas its equivalent axiomatic definition supports the `memevents` exhaustive search and permits an easy comparison with related models; the similarity with SPARCv8 suggests x86-TSO is strong enough to program above. Mechanisation in HOL4 revealed a number of subtle points of detail, including some of the well-formed event structure conditions that we depend on (e.g. that instructions have no *internal* data races). We hope this will clarify the semantics of x86 architectures.

Acknowledgements We thank Luc Maranget for his work on `memevents`, and David Christie, Dave Dice, Doug Lea, Paul Loewenstein, Gil Neiger, and Francesco Zappa Nardelli for their helpful remarks. We acknowledge funding from EPSRC grant EP/F036345.

A Litmus Tests and Discussion

In this appendix we go through a number of litmus tests, including all those from the Intel IWP [12], the Intel SDM (Rev-29) [2], the AMD64 APM [1], and some others.

Tests are named `iwpXXX`, for the test numbered XXX from IWP, `amdNNN`, for the NNN'th test from the AMD64 APM, or `nNNN`, for other tests. We let `x` and `y` range over distinct aligned memory locations. Unless otherwise stated, all tests are considered with respect to an initial state in which those locations and all registers contain 0. We write assembly instructions in Intel syntax (as opposed to the AT&T syntax used by `gcc` and `gas`) except that we write an arrow $\leftarrow$ instead of a comma, to clarify the direction of data flow.


The results of the tests, in x86-CC, x86-TSO, and our observations on actual processors, are summarised in §A.9. In all cases the x86-TSO behaviours are consistent with our observations, though it does permit some behaviours that we have not observed. In other words, this data does not contradict the claim that x86-TSO is sound.

A.1 Load/Store Reordering

Rev-29 Example 7-1. Stores Are Not Reordered with Other Stores.

iwp2.1/amd1	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV EAX←[y]
poi:1	MOV [y]←\$1	MOV EBX←[x]
Forbid: 1:EAX=1 ∧ 1:EBX=0		

This is Test 2.1 in the Intel White Paper (IWP), and also the first test in the AMD64-3.14. There are two stores on one processor to two different locations, and the other processor loads from those locations in the opposite order. Executions in which the first load reads from the second store, but the second load reads from the initial state, not from the first store, are forbidden. We illustrate such executions below, in a diagram showing the reads-from map over the memory events (eliding register events and the *eiid* and *iiid* data of each memory event). This test also shows that loads are not reordered with other loads.




Rev-29 Example 7-2. Stores Are Not Reordered with Older Loads.

iwp2.2/amd2	proc:0	proc:1
poi:0	MOV EAX←[x]	MOV EBX←[y]
poi:1	MOV [y]←\$1	MOV [x]←\$1
Forbid: 0:EAX=1 ∧ 1:EBX=1		

Here the two processors read a value, from two different locations respectively, and then each stores a value to the other location. Executions as shown below, in which each reads from the write of the other


processor, are forbidden.



Rev-29 Example 7-3. Loads May be Reordered with Older Stores.

iwp2.3.a/amd4	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [y]←\$1
poi:1	MOV EAX←[y]	MOV EBX←[x]
Allow: 0:EAX=0 ∧ 1:EBX=0		


This test illustrates non-sequentially-consistent behaviour allowed by store buffering, as discussed in §1. Each processor stores to one location and loads from the other location. Both loads are allowed to read from the initial state. We illustrate one such allowed execution below.



Rev-29 Example 7-4. Loads Are not Reordered with Older Stores to the Same Location.

iwp2.3.b	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [y]←\$1
poi:1	MOV EAX←[x]	MOV EBX←[y]
Require: 0:EAX=1 ∧ 1:EBX=1		


This is variant of the previous test in which each load reads from the same location as that processor stored to. Here they are forced to read from the local store. Operationally, loads must read from the local store buffer if it contains a write to the address in question. All executions must satisfy the ‘Require’ condition; one such is shown below. (The Rev-29 example is identical to the proc:0 part of iwp2.3.b; the proc:1 part adds no more force).



Rev-29 Example 7-5. Intra-Processor Forwarding is Allowed.

iwp2.4/amd9	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [y]←\$1
poi:1	MOV EAX←[x]	MOV ECX←[y]
poi:2	MOV EBX←[y]	MOV EDX←[x]
Allow: 0:EBX=0 ∧ 1:EDX=0		


This test is a variant of Example 7-3 (iwp2.3.a/amd4) in which each processor stores to a location *then* reads from that location before reading from the other location. Again, store buffers allow the specified final state, and one such execution is shown below.



Rev-29 Example 7-6. Stores Are Transitively Visible.

iwp2.5/amd8	proc:0	proc:1	proc:2
poi:0	MOV [x]←\$1	MOV EAX←[x]	MOV EBX←[y]
poi:1		MOV [y]←\$1	MOV ECX←[x]
Forbid: 1:EAX=1 ∧ 2:EBX=1 ∧ 2:ECX=0			


This test shows that a particular transitive chain of causality must be respected. Here proc:1 reads proc:0's write to x, then (in program order) proc:1 writing to y; proc:2 reads from that write of y, then (in program order) proc:1 reads x. That final read is not allowed to be from the initial state. The forbidden reads-from relationship is shown below. The Write-to-Read Causality (WRC) test of Boehm and Adve [5, Fig. 5] is similar to this test, but with fences between the proc:1 and proc:2 pairs of instructions. In x86-TSO the final state is forbidden without requiring such fences.



AMD3.14 Test 3

amd3	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [y]←\$1
poi:1	MOV [x]←\$2	MOV [y]←\$2
poi:2	MOV EAX←[y]	MOV EBX←[x]
Allow: 0:EAX=1 ∧ 1:EBX=1		


This test is another variant of Example 7-3 (iwp2.3.a/amd4) in which each processor stores to a location *then stores again to that location* before reading from the other location. Again, store buffers allow each processor to read from the first write of the other processor, and one such execution is shown below.



IWP Test 2.6. Total Order on Stores to the Same Location

iwp2.6	proc:0	proc:1	proc:2	proc:3
poi:0	MOV [x]←\$1	MOV [x]←\$2	MOV EAX←[x]	MOV ECX←[x]
poi:1			MOV EBX←[x]	MOV EDX←[x]
Forbid: 2:EAX=1 ∧ 2:EBX=2 ∧ 3:ECX=2 ∧ 3:EDX=1				


This test requires stores to the same location by two processors to be observed in the same order by two other processors. The forbidden reads-from relation is shown below. The test was in IWP, but is no longer present in rev-29.



A.2 Independent Reads of Independent Writes

The following IRIW example, discussed here in §2.2 and by Boehm and Adve [5], has a final state permitted in AMD64-3.14 (explicitly), in IWP (implicitly), and in x86-CC. Rev-29, however, contains the same example but with the given final state forbidden (rev-29 Example 7-7. Stores Are Seen in a Consistent Order by Other Processors). We have not observed the final state in practice, and it is forbidden in x86-TSO. The reads-from relation required for the given final state is below; in such executions, proc:0 and proc:1 write to two different locations, and proc:2 and proc:3 read from those locations, seeing the writes in opposite orders (proc:2 sees the write to x but not the write to y, reading y from the initial state, whereas proc:3 sees the write to x but reads x from the initial state).

amd6	proc:0	proc:1	proc:2	proc:3
poi:0	MOV [x]←\$1	MOV [y]←\$1	MOV EAX←[x]	MOV ECX←[y]
poi:1			MOV EBX←[y]	MOV EDX←[x]
Final: 2:EAX=1 ∧ 2:EBX=0 ∧ 3:ECX=1 ∧ 3:EDX=0				
cc : Allow; tso : Forbid				



A.3 Locked Instructions: Tests iwp2.7/amd7, iwp2.8.a, iwp2.8.b, n8, and n3

Only the following three IWP litmus tests involve locked instructions, and none of the three forbidden results are allowed in x86-TSO. However, only iwp2.8.a is forbidden by the parts of x86-TSO dealing with locking; the other two are forbidden by the core TSO properties directly (i.e., even if hypothetical non-locked versions of the XCHG instructions were used).

Rev-29 Example 7-8. Locked Instructions Have a Total Order.

iwp2.7/amd7	proc:0	proc:1	proc:2	proc:3
poi:0	XCHG [x]←EAX	XCHG [y]←EBX	MOV ECX←[x]	MOV ESI←[y]
poi:1			MOV EDX←[y]	MOV EDI←[x]
Initial state: 0:EAX = 1; 1:EBX = 1 (elsewhere 0)				
Forbid: 2:ECX=1 ∧ 2:EDX=0 ∧ 3:ESI=1 ∧ 3:EDI=0				

This is a variant of the IRIW example amd6 but using XCHG instructions (which are implicitly LOCK'd) instead of MOV instructions on proc:0 and proc:1.

Rev-29 Example 7-9. Loads Are not Reordered with Locks.

iwp2.8.a	proc:0	proc:1
poi:0	XCHG [x]←EAX	XCHG [y]←ECX
poi:1	MOV EBX←[y]	MOV EDX←[x]
Initial state: 0:EAX = 1; 1:ECX = 1 (elsewhere 0)		
Forbid: 0:EBX=0 ∧ 1:EDX=0		

Rev-29 Example 7-10. Stores Are not Reordered with Locks.

iwp2.8.b	proc:0	proc:1
poi:0	XCHG [x]←EAX	MOV EBX←[y]
poi:1	MOV [y]←\$1	MOV ECX←[x]
Initial state: 0:EAX = 1 (elsewhere 0)		
Forbid: 1:EBX=1 ∧ 1:ECX=0		

Test n8. Loads Are not Reordered with Locks — Single-XCHG Variant.

n8	proc:0	proc:1
poi:0	XCHG [x]←EAX	MOV [y]←\$1
poi:1	MOV EBX←[y]	MOV EDX←[x]
Initial state: 0:EAX = 1; [y] = 1 (elsewhere 0)		
Allow: 0:EAX=0 ∧ 1:EDX=0		

This test is a variant of iwp2.8.a, but with an unlocked write in place of one of the XCHG instructions. The x86-TSO model allows the result where EBX and EDX are both 0.

Test n3: Independent Locked Instructions and Stores

n3	proc:0	proc:1	proc:2	proc:3
poi:0	XCHG EAX←[x]	MOV [y]←\$1	MOV EBX←[y]	MOV ESI←[x]
poi:1			MOV ECX←[x]	MOV EDI←[y]
poi:2			MOV EDX←[x]	MOV EBP←[y]
Initial state: 0:EAX = 1 (elsewhere 0)				
Final: 2:EBX=1 ∧ 2:ECX=0 ∧ 2:EDX=1 ∧ 3:ESI=1 ∧ 3:EDI=0 ∧ 3:EBP=1				
cc : Allow; tso : Forbid				




IWP left open the question of whether a locked instruction and a store to a different location could be seen in different orders by two other processors. We introduced the n3 test above to cover this case [20]. The given final state was allowed by x86-CC, but is not allowed by x86-TSO (again without the LOCK having anything to do with the reasoning). We have not observed the final state to occur in practice.

A.4 Fence Instructions: Tests amd5 and amd10

We handle fences by forcing write-to-read program order dependencies into the memory order. The following amd10 example is very similar to iw2.8.a, and the forbidden behavior is prohibited by x86-TSO by essentially the same mechanism. Test amd5 below is similar. Our interpretation seems consistent with rev-29 principles P11 and P12.

amd10	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [y]←\$1
poi:1	MFENCE	MFENCE
poi:2	MOV EAX←[x]	MOV ECX←[y]
poi:3	MOV EBX←[y]	MOV EDX←[x]
Final: 0:EBX=0 ∧ 1:EDX=0		
cc : Allow; tso : Forbid		

amd5	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [y]←\$1
poi:1	MFENCE	MFENCE
poi:2	MOV EAX←[y]	MOV EBX←[x]
Final: 0:EAX=0 ∧ 1:EBX=0		
cc : Allow; tso : Forbid		




A.5 The Unsoundness of IWP/AMD3.14/x86-CC: Test n6

The n6 test below, by Loewenstein, was discussed in §2.2, where we explained how the final state is allowed by write-buffer implementations.

n6	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [y]←\$2
poi:1	MOV EAX←[x]	MOV [x]←\$2
poi:2	MOV EBX←[y]	
Final: 0:EAX=1 ∧ 0:EBX=0 ∧ [x]=1		
cc : Forbid; tso : Allow		

The final state is also allowed in x86-TSO, as shown in the execution below.



The final state is allowed (as far as we can interpret the principles) in rev-29. The final state is observable on Intel processors: we find witnesses on an Intel Core 2, with our `litmus` tool. However, the given final state is forbidden in our x86-CC model, and by any interpretation we can make of the IWP or AMD3.14 principles. We have:


- (a, b) in preserved_program_order (by IWP P4)
- (b, d) in preserved_program_order (by IWP P1)
- (d, f) in the view order of proc:0 (otherwise d could not read value 0)
- (f, g) in preserved_program_order (by IWP P2)
- So the view order of proc:0 must be a, b, d, f, g .
- So (a, g) in the write serialization for location x (by IWP P6)
- So the final state must have $[x] = 2$, not $[x] = 1$.

A.6 The Weakness of Rev-29: Tests n4 and n5


As discussed in §2.3, the following examples were not allowed in x86-CC, nor are they allowed by x86-TSO. We would be surprised if they were allowed by any reasonable implementation, and have not observed them on Intel processors. Programming above a model that permitted them would be problematic. However, the rev-29 principles seem to allow them. In particular, P9 has no force because

it only involves processors other than those performing the two stores.

n5	proc:0	proc:1
poi:0	MOV [x]←\$1	MOV [x]←\$2
poi:1	MOV EAX←[x]	MOV EBX←[x]
Forbid: 0:EAX=2 ∧ 1:EBX=1		



n4	proc:0	proc:1
poi:0	MOV EAX←[x]	MOV ECX←[x]
poi:1	MOV [x]←\$1	MOV [x]←\$2
poi:2	MOV EBX←[x]	MOV EDX←[x]
Forbid: 0:EAX=2 ∧ 0:EBX=1 ∧ 1:ECX=1 ∧ 1:EDX=2		




A.7 Interpreting the rev-29 “not reordered with”: Test n7

Key phrases in the vendor documentation (Intel IWP, AMD64-3.14, and Intel SDM rev-29) are not given definite meanings, making it hard use them to justify x86-TSO. Principles P1, P2, P3, P4, and P8 of rev-29 (reproduced in Appendix B) all refer to some events being “not reordered with” others. These principles were also present in IWP/AMD64-3.14, and in x86-CC we interpreted them as giving rise to a preserved-program-order relation, which was included in a transitive happens-before relation, with which each processor’s view order had to be consistent. In a TSO-based model, however, one might expect to interpret two events being “not reordered” by requiring that any program order relation between them must also appear in the transitive memory order (as suggested by P5). This seems reasonable, and is what we do, for P1, P2, P3, and P8. However, P4 states READS MAY BE REORDERED WITH OLDER WRITES TO DIFFERENT LOCATIONS BUT NOT WITH OLDER WRITES TO THE SAME LOCATION. There is no problem with the first part, but for the second part, that interpretation would forbid the final state of litmus test n7 below (discussed in §3.2), whereas it is allowed on x86-TSO.

n7	proc:0	proc:1	proc:2
poi:0	MOV [x]←\$1	MOV [y]←\$1	MOV ECX←[y]
poi:1	MOV EAX←[x]		MOV EDX←[x]
poi:2	MOV EBX←[y]		
Allow: 0:EAX=1 ∧ 0:EBX=0 ∧ 2:ECX=1 ∧ 2:EDX=0			

We illustrate this for the example x86-TSO valid execution below. If the second part of P4 were interpreted to give rise to preserved program order edges appearing in memory order, then there would be a

memory order edge from a to b. Read d would then have to read from f, not from the initial state.




There are other valid executions, but not with the given final state.

A.8 Other Tests

Test n1: Reordering of Loads with Older Stores to Different Locations

n1	proc:0	proc:1	proc:2
poi:0	MOV [x]←\$2	MOV [y]←\$1	MOV EBX←[x]
poi:1	MOV EAX←[y]	MOV [x]←\$1	MOV ECX←[x]
Allow: 0:EAX=0 ∧ 2:EBX=1 ∧ 2:ECX=2			

In x86-CC, the allowed final result of iwp2.3.a/amd4 did not require the reordering of loads with older stores to different locations that the test description spoke of. We introduced Test n1 [20] as an example where such reordering was (in x86-CC) essential. The final behaviour is still allowed in x86-TSO, as shown below.




Test n2: Transitive Causality via Preserved Program Order and the Write Serialisation

n2	proc:0	proc:1	proc:2	proc:3
poi:0	MOV [y]←\$1	MOV [x]←\$2	MOV EAX←[x]	MOV ECX←[z]
poi:1	MOV [x]←\$1	MOV [z]←\$1	MOV EBX←[x]	MOV EDX←[y]
Forbid: 2:EAX=1 ∧ 2:EBX=2 ∧ 3:ECX=1 ∧ 3:EDX=0				

We introduced n2 [20] to show transitivity through preserved program order (of the proc:0 and proc:1 events) and the write serialisation for x (which has W [x]=1 before W [x]=2 by the proc:2 observations).


There are no x86-TSO executions with the reads-from map below.



Test RWC (unfenced): Read-to-Write Causality

rw-unfenced	proc:0	proc:1	proc:2
poi:0	MOV [x]←\$1	MOV EAX←[x]	MOV [y]←\$1
poi:1		MOV EBX←[y]	MOV ECX←[x]
Allow: 1:EAX=1 ∧ 1:EBX=0 ∧ 2:ECX=0			


This test is based on the Read-to-Write Causality (RWC) test of Boehm and Adve [5, Fig. 6]. With no fences, the non-SC behaviour below is allowed in x86-TSO.



Test RWC (singly fenced): Read-to-Write Causality

rw-fenced	proc:0	proc:1	proc:2
poi:0	MOV [x]←\$1	MOV EAX←[x]	MOV [y]←\$1
poi:1		MOV EBX←[y]	MFENCE
poi:2			MOV ECX←[x]
Final: 1:EAX=1 ∧ 1:EBX=0 ∧ 2:ECX=0			
cc : Allow; tso : Forbid			

Adding an MFENCE to proc:2 rules out that behaviour:



A.9 Summary of Test Results

The table below summarises the results for all the previous tests in the x86-CC model, the x86-TSO model, and our observations on actual processors. The `memevents-cc.txt` and `memevents-tso.txt` columns give the behaviour in our x86-CC and x86-TSO models respectively¹. These results are from our `memevents` tool, which includes hand-written executable implementations of the formal models. For x86-TSO, “Allow” results are validated with the verified checker described in §4. The “actual processors” data is from observations with our `litmus` tool on an Intel Core 2 Duo, and indicates the number of positive (for “Allow” or “Require” tests) or negative (for “Forbid” tests) observations out of the number of runs.

	memevents-cc	memevents-tso	Actual hardware	
amd10	Allow	Forbid	Forbid validated	0/200000
amd3	Allow	Allow	Allow not validated	0/200000
amd5	Allow	Forbid	Forbid validated	0/200000
amd6	Allow	Forbid	Forbid validated	0/200000
iwp2.1/amd1	Forbid	Forbid	Forbid validated	0/200000
iwp2.2/amd2	Forbid	Forbid	Forbid validated	0/200000
iwp2.3.a/amd4	Allow	Allow	Allow validated	46/200000
iwp2.3.b	Require	Require	Require validated	200000/200000
iwp2.4/amd9	Allow	Allow	Allow validated	193/200000
iwp2.5/amd8	Forbid	Forbid	Forbid validated	0/200000
iwp2.6	Forbid	Forbid	Forbid validated	0/200000
iwp2.7/amd7	Forbid	Forbid	Forbid validated	0/200000
iwp2.8.a	Forbid	Forbid	Forbid validated	0/200000
iwp2.8.b	Forbid	Forbid	Forbid validated	0/200000
n1	Allow	Allow	Allow not validated	0/200000
n2	Forbid	Forbid	Forbid validated	0/200000
n3	Allow	Forbid	Forbid validated	0/200000
n4	Forbid	Forbid	Forbid validated	0/200000
n5	Forbid	Forbid	Forbid validated	0/200000
n6	Forbid	Allow	Allow validated	9/200000
n7	Allow	Allow	Allow not validated	0/200000
n8	Allow	Allow	Allow validated	130832/200000
rwc-fenced	Allow	Forbid	Forbid validated	0/200000
rwc-unfenced	Allow	Allow	Allow not validated	0/200000

Discrepancies between x86-TSO and the observed behaviour could be of two kinds: For a “Forbid” or “Require” test, an observation of the forbidden result, which we would indicate with “Forbid NOT validated” or “Require NOT validated” respectively, would indicate that the formal x86-TSO model is unsound with respect to actual hardware. We see no such observations.

For an “Allow” test, the absence of an observation of the allowed result, which we indicate with “Allow not validated” (in blue), may suggest that x86-TSO is weaker (allowing more behaviours) than actual processors, or may merely mean that the test has not been run sufficiently often or with the processor in a state to exhibit the behaviour.

Note also that this experimental data is from just a few tests run on just one machine of a specific implementation, whereas many x86 implementations, with radically different microarchitectures, are in widespread use. Further testing would be desirable to increase confidence in the soundness of x86-TSO.

It would be very surprising if x86-TSO were not more liberal in some ways than actual processors. For example, actual processors will have size-bounded write buffers, whereas x86-TSO does not, and probably should not (otherwise reasoning above the model would be needlessly complex and implementation-specific).

¹For the fence tests (amd10, amd5, and rwc-fenced) the x86-CC column indicates the behaviour in the extension of our formal x86-CC model (which did not cover fences) with MFENCE enforcing local write-read ordering in happens-before.

B The IWP and Rev-29 Principles and Litmus Tests

B.1 Principles

The Intel documentation, both IWP and the SDM Vol.3A rev-29 §7.2.2, states a number of “principles”. We reproduce them here for reference.

Eight of the rev-29 principles correspond to principles from IWP, in the sense that they have the same force for the fragment we are considering, but may have additional conditions, here greyed out, relevant to other instructions. For consistency with our previous work we tag these P1–P8, following the IWP numbering. In addition there are five new clauses, which we tag P9–P14. We list them in the same order as Rev-29.

P1 Reads are not reordered with other reads.

P3 Writes are not reordered with older reads.

P2 Writes to memory are not reordered with other writes[, with the exception of writes executed with the CLFLUSH instruction, streaming stores (writes) executed with the non-temporal move instructions (MOVNTI, MOVNTQ, MOVNTDQ, MOVNTPS, and MOVNTPD), string operations (see Section 7.2.4.1)].

P4 Reads may be reordered with older writes to different locations but not with older writes to the same location.

P8 Reads or writes cannot be reordered with [I/O instructions,] locked instructions[, or serializing instructions].

P11 Reads cannot pass LFENCE and MFENCE instructions.

P12 Writes cannot pass SFENCE and MFENCE instructions.

In a multiple-processor system, the following ordering principles apply:

P13 Individual processors use the same ordering principles as in a single-processor system.

P10 Writes by a single processor are observed in the same order by all processors.

P14 Writes from an individual processor are NOT ordered with respect to the writes from other processors.

P5 Memory ordering obeys causality (memory ordering respects transitive visibility).

P9 Any two stores are seen in a consistent order by processors other than those performing the stores.

P7 Locked instructions have a total order.

Summarising the differences of rev-29 with respect to IWP, the IWP P6

P6 In a multiprocessor system, stores to the same location have a total order.

has been replaced by P9; P11 and P12, concerning the fence instructions, have been added, and P10, P13 and P14 have been added (though they might have been considered implicit in IWP).

B.2 Litmus Tests

The rev-29 litmus tests, Examples 7-1 to 7-10, are essentially identical to the IWP litmus tests, except that iw2.6 is removed and the IRIW test, with forbidden final state, is added. Rev-29 also adds 6 Litmus tests for string operations, which are not in the fragment of the ISA we consider in this paper.

The vendor litmus tests appear to be a more reliable source than the prose principles: the litmus tests are almost completely unambiguous (the only exception is that in some cases it is not explicitly stated that certain locations are not aliased), and, as far as we can observe, are consistent with the behaviour of actual processors, whereas the principles are hard to interpret, as we have discussed. Of course, the litmus tests do not by themselves *determine* a useful memory model, as they specify behaviour only for a small number of very specific programs.

C Well-Formed Event Structures

In this appendix we define our well-formedness condition for event structures, comprising a number of sanity conditions (that there are only a finite number of processors, that the intra-instruction causality relation is indeed a partial order, etc.). This follows our x86-CC work, in which we proved that all event structures arising from the semantics for assembly programs are indeed well-formed.

`well_formed_event_structure` $E =$

(* The set of events is at most countable *)
`countable` $E.events \wedge$

(* there are only a finite number of processors *)
`(finite` $E.procs) \wedge$

(* all events are from one of those processors *)
 $(\forall e \in (E.events). \text{proc } e \in E.procs) \wedge$

(* the eiid and iiid of an event (together) identify it uniquely *)
 $(\forall e_1 e_2 \in (E.events). (e_1.eiid = e_2.eiid) \wedge (e_1.iiid = e_2.iiid) \implies (e_1 = e_2)) \wedge$

(* intra-instruction causality is a partial order over the events *)
`partial_order` $(E.intra_causality)E.events \wedge$

(* ...and moreover, is *intra*-instruction *)
 $(\forall (e_1, e_2) \in (E.intra_causality). (e_1.iiid = e_2.iiid)) \wedge$

(* the atomicity data is a partial equivalence relation: the atomic sets of events are disjoint *)
`per` $E.events E.atomicity \wedge$

(* atomic sets are *intra* instruction *)
 $(\forall es \in (E.atomicity). \forall e_1 e_2 \in es. (e_1.iiid = e_2.iiid)) \wedge$

(* accesses to a register on a processor can only be by that processor *)
 $(\forall e \in (E.events). \forall p r. (\text{loc } e = \text{SOME } (\text{LOCATION_REG } p \ r)) \implies (p = \text{proc } e)) \wedge$

(* An event never comes after an infinite number of other events in program order *)
`finite_prefixes` $(\text{po_iico } E)E.events \wedge$

(* The additional properties below hold, for the ISA fragment dealt with in [SSFN+09], and were useful for the metatheory there, but seem less essential than those above. *)

(* there is no intra-causality edge *from* a memory write *)
 $(\forall (e_1, e_2) \in (E.intra_causality). e_1 \neq e_2 \implies e_1 \notin \text{mem_writes } E) \wedge$

(* if an instruction two events on a location and one is a write, then there must be an intra-causality edge between them. In other words, there cannot be a local race within an instruction *)
 $(\forall (e_1 \in \text{writes } E) e_2.$

$(e_1 \neq e_2) \wedge$
 $(e_2 \in \text{writes } E \vee e_2 \in \text{reads } E) \wedge$
 $(e_1.iiid = e_2.iiid) \wedge$
 $(\text{loc } e_1 = \text{loc } e_2)$
 $\implies$
 $(e_1, e_2) \in E.intra_causality \vee$
 $(e_2, e_1) \in E.intra_causality) \wedge$

(* each atomic set includes all the events of its instruction *)
 $(\forall es \in (E.atomicity). \forall e_1 \in es. \forall e_2 \in (E.events). (e_1.iiid = e_2.iiid) \implies e_2 \in es) \wedge$

(* all locked instructions include at least one memory read *)
 $(\forall es \in (E.atomicity). \exists e \in es. e \in \text{mem_reads } E)$

D Auxiliary Definitions

This appendix collects the remaining auxiliary definitions used in the x86-TSO axiomatic and abstract machine memory model definitions. These are mostly routine, but presented here for reference.

D.1 Axiomatic Memory Model

type_abbrev `proc` : num

type_abbrev `Ximm` : word32

type_abbrev `address` : Ximm

type_abbrev `value` : Ximm

type_abbrev `eiid` : num

type_abbrev `reln` : 'a#'*a* $\rightarrow$ bool

`is_mem_access` *e* = $\exists d \ a \ v. e.action = \text{ACCESS } d \ (\text{LOCATION_MEM } a) v$

`writes` *E* = $\{e \mid e \in E.events \wedge \exists l \ v. e.action = \text{ACCESS } W \ l \ v\}$

`reads` *E* = $\{e \mid e \in E.events \wedge \exists l \ v. e.action = \text{ACCESS } R \ l \ v\}$

`fences` *E* = $\{e \mid e \in E.events \wedge (\exists f. e.action = \text{BARRIER } f)\}$

`mfences` *E* = $\{e \mid e \in E.events \wedge (e.action = \text{BARRIER } \text{MFENCE})\}$

`mem_writes` *E* = $\{e \mid e \in E.events \wedge \exists a \ v. e.action = \text{ACCESS } W \ (\text{LOCATION_MEM } a) v\}$

`mem_reads` *E* = $\{e \mid e \in E.events \wedge \exists a \ v. e.action = \text{ACCESS } R \ (\text{LOCATION_MEM } a) v\}$

`reg_writes` *E* = $\{e \mid e \in E.events \wedge \exists p \ r \ v. e.action = \text{ACCESS } W \ (\text{LOCATION_REG } p \ r) v\}$

`reg_reads` *E* = $\{e \mid e \in E.events \wedge \exists p \ r \ v. e.action = \text{ACCESS } R \ (\text{LOCATION_REG } p \ r) v\}$

`mem_accesses` *E* = $\{e \mid e \in E.events \wedge (\exists d \ a \ v. e.action = \text{ACCESS } d \ (\text{LOCATION_MEM } a) v)\}$

`reg_accesses` *E* = $\{e \mid e \in E.events \wedge \exists d \ p \ r \ v. e.action = \text{ACCESS } d \ (\text{LOCATION_REG } p \ r) v\}$

`loc` *e* =

case *e.action* **of**
 $\text{ACCESS } d \ l \ v \rightarrow \text{SOME } l$
 $\parallel \text{BARRIER } f \rightarrow \text{NONE}$

value_of $e =$
case $e.action$ **of**
 ACCESS $d\ l\ v \rightarrow \text{SOME } v$
 || BARRIER $f \rightarrow \text{NONE}$

proc $e = e.iid.proc$

The following definition of valid executions collects together the conditions in the text of §3.2.

valid_execution $E\ X =$
 partial_order $X.memory_order$ (mem_accesses E) $\wedge$
 linear_order $(X.memory_order|_{(\text{mem_writes } E)})(\text{mem_writes } E) \wedge$
 finite_prefixes $X.memory_order$ (mem_accesses E) $\wedge$
 $(\forall ew \in (\text{mem_writes } E)).$
 finite $\{er \mid er \in E.events \wedge (\text{loc } er = \text{loc } ew) \wedge$
 $(er, ew) \notin X.memory_order \wedge (ew, er) \notin X.memory_order\} \wedge$
 $(\forall er \in (\text{mem_reads } E)). \forall e \in (\text{mem_accesses } E). (er, e) \in \text{po_iico } E \implies (er, e) \in X.memory_order) \wedge$
 $(\forall ew_1\ ew_2 \in (\text{mem_writes } E). (ew_1, ew_2) \in \text{po_iico } E \implies (ew_1, ew_2) \in X.memory_order) \wedge$
 $(\forall ew \in (\text{mem_writes } E). \forall er \in (\text{mem_reads } E). \forall ef \in (\text{mfences } E).$
 $(ew, ef) \in \text{po_iico } E \wedge (ef, er) \in \text{po_iico } E \implies (ew, er) \in X.memory_order) \wedge$
 $(\forall e_1\ e_2 \in (\text{mem_accesses } E). \forall es \in (E.atomicity).$
 $(e_1 \in es \vee e_2 \in es) \wedge (e_1, e_2) \in \text{po_iico } E$
 $\implies$
 $(e_1, e_2) \in X.memory_order) \wedge$
 $(\forall es \in (E.atomicity). \forall e \in (\text{mem_accesses } E \setminus es).$
 $(\forall e' \in (es \cap \text{mem_accesses } E). (e, e') \in X.memory_order) \vee$
 $(\forall e' \in (es \cap \text{mem_accesses } E). (e', e) \in X.memory_order)) \wedge$
 $X.rfmap \in \text{reads_from_map_candidates } E \wedge$
 check_rfmap_written $E\ X \wedge$
 check_rfmap_initial $E\ X$

max_state_updates $E\ X\ l =$
 $\{\text{value_of } ew \mid ew \in \text{maximal_elements}$
 $\{\text{value_of } ew' \mid ew' \in \text{writes } E \wedge (\text{loc } ew' = \text{SOME } l)\}$
 $(\text{case } l \text{ of}$
 LOCATION_MEM $a \rightarrow X.memory_order$
 || LOCATION_REG $p\ r \rightarrow \text{po_iico } E)\}$

(check_final_state $E\ X\ \text{NONE} =$
 $\neg(\text{finite } E.events)) \wedge$
 (check_final_state $E\ X\ (\text{SOME } final_state) =$
 $\text{finite } E.events \wedge$
 $(\forall l.$
if (max_state_updates $E\ X\ l) = \{\}$ **then**
 $final_state\ l = X.initial_state\ l$
else
 $final_state\ l \in \text{max_state_updates } E\ X\ l))$

D.2 Abstract Machine Memory Model

clause_name $x = \mathbf{T}$

not_blocked $s\ p = (s.L = \text{NONE}) \vee (s.L = \text{SOME } p)$

no_pending $b\ a = \neg(\exists v'. \text{MEM } (a, v')\ b)$

The following is the primary HOL definition of the abstract machine transitions. The rules in Fig. 1 are typeset from an equivalent HOL definition in a slightly different style.

$$\begin{aligned}
& (\forall s \ a \ v \ p. \\
& \text{clause_name "read-mem"} \wedge \\
& \text{not_blocked } s \ p \wedge \\
& (s.M \ a = \text{SOME } v) \wedge \\
& \text{no_pending } (s.B \ p) a \\
& \implies \\
& \text{machine_trans } s \\
& \quad (\text{EVT } p \ (\text{ACCESS } R \ (\text{LOCATION_MEM } a) v)) \\
& \quad s) \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ a \ v \ p. \\
& \text{clause_name "read-buffer"} \wedge \\
& \text{not_blocked } s \ p \wedge \\
& (\exists b_1 \ b_2. (s.B \ p = b_1 ++ [(a, v)] ++ b_2) \wedge \text{no_pending } b_1 \ a) \\
& \implies \\
& \text{machine_trans } s \\
& \quad (\text{EVT } p \ (\text{ACCESS } R \ (\text{LOCATION_MEM } a) v)) \\
& \quad s) \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ r \ v \ p. \\
& \text{clause_name "read-reg"} \wedge \\
& (*\text{not_blocked } s \ p \ /\ *) \\
& (s.R \ p \ r = \text{SOME } v) \\
& \implies \\
& \text{machine_trans } s \\
& \quad (\text{EVT } p \ (\text{ACCESS } R \ (\text{LOCATION_REG } p \ r) v)) \\
& \quad s) \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ a \ v \ p \ s'. \\
& \text{clause_name "write-buffer"} \wedge \\
& (*\text{not_blocked } s \ p \ /\ *) \\
& (s' = \langle\langle R := s.R; \\
& \quad M := s.M; \\
& \quad B := (p \mapsto (a, v) \in s.B \ p) s.B; \\
& \quad L := s.L \rangle\rangle) \\
& \implies \\
& \text{machine_trans } s \\
& \quad (\text{EVT } p \ (\text{ACCESS } W \ (\text{LOCATION_MEM } a) v)) \\
& \quad s') \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ a \ v \ p \ b \ s'. \\
& \text{clause_name "write-mem"} \wedge \\
& \text{not_blocked } s \ p \wedge \\
& (s.B \ p = b ++ [(a, v)]) \wedge \\
& (s' = \langle\langle R := s.R; \\
& \quad M := (a \mapsto \text{SOME } v) s.M; \\
& \quad B := (p \mapsto b) s.B; \\
& \quad L := s.L \rangle\rangle) \\
& \implies \\
& \text{machine_trans } s \ \text{TAU } s') \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ r \ v \ p \ s'. \\
& \text{clause_name "write-reg"} \wedge \\
& (*\text{not_blocked } s \ p \ *) \\
& (s' = \langle\langle R := (p \mapsto (r \mapsto \text{SOME } v) (s.R \ p)) s.R;
\end{aligned}$$

$$\begin{aligned}
& M := s.M; \\
& B := s.B; \\
& L := s.L \rangle \rangle \\
\Rightarrow \\
\text{machine_trans } s & \quad (\text{EVT } p \text{ (ACCESS W (LOCATION_REG } p \text{ } r) v)) \\
& \quad s') \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ p. \\
& \text{clause_name "barrier"} \wedge \\
& (*\text{not_blocked } s \ p \ /\wedge *) \\
& (s.B \ p = []) \\
\Rightarrow \\
\text{machine_trans } s & \quad (\text{EVT } p \text{ (BARRIER MFENCE)) } s) \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ p \ b. \\
& \text{clause_name "nop"} \wedge \\
& (*\text{not_blocked } s \ p \ /\wedge *) \\
& b \neq \text{MFENCE} \\
\Rightarrow \\
\text{machine_trans } s & \quad (\text{EVT } p \text{ (BARRIER } b)) s) \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ p \ s'. \\
& \text{clause_name "lock"} \wedge \\
& (s.L = \text{NONE}) \wedge \\
& (s.B \ p = []) \wedge \\
& (s' = \langle \langle R := s.R; \\
& \quad M := s.M; \\
& \quad B := s.B; \\
& \quad L := \text{SOME } p \rangle \rangle) \\
\Rightarrow \\
\text{machine_trans } s & \quad (\text{LOCK } p) s') \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall s \ p \ s'. \\
& \text{clause_name "unlock"} \wedge \\
& (s.L = \text{SOME } p) \wedge \\
& (s.B \ p = []) \wedge \\
& (s' = \langle \langle R := s.R; \\
& \quad M := s.M; \\
& \quad B := s.B; \\
& \quad L := \text{NONE} \rangle \rangle) \\
\Rightarrow \\
\text{machine_trans } s & \quad (\text{UNLOCK } p) s')
\end{aligned}$$

$$\begin{aligned}
& \text{machine_state_to_state_constraint } s = \\
& \lambda l. \\
& \text{case } l \text{ of} \\
& \quad \text{LOCATION_MEM } a \rightarrow s.M \ a \\
& \quad \parallel \text{LOCATION_REG } p \ r \rightarrow s.R \ p \ r
\end{aligned}$$

$$\begin{aligned}
& \text{machine_final_state } path = \\
& \text{if finite } path \text{ then} \\
& \text{SOME (machine_state_to_state_constraint (last } path))} \\
& \text{else} \\
& \text{NONE}
\end{aligned}$$

$$\text{machine_init_state } sc =$$

```

 $\llbracket R := (\lambda p. sc \text{ (LOCATION\_REG } p \text{ )});$ 
 $M := (\lambda a. sc \text{ (LOCATION\_MEM } a));$ 
 $B := (\lambda p. []);$ 
 $L := \text{NONE} \rrbracket$ 

```

$\text{is_init } s = \exists sc. s = \text{machine_init_state } sc$

The definitions of okEpath and okMpath are given over event machines, in Appendix E.

E Proof Outlines

E.1 Event-annotated machine

We annotate the abstract machine with events to support the equivalence proof between the abstract machine and axiomatic presentations of x86-TSO.

The event abstract machine stores with each memory and register location the event that last wrote there. It also keeps events in the write buffer.

```

 $\text{evt\_machine\_state} = \llbracket$ 
  (* Per processor registers, annotated with the event that last wrote it *)
   $eR : \text{proc} \rightarrow 'reg \rightarrow (\text{value} \# 'reg \text{ event option}) \text{ option};$ 
  (* main memory, annotated with the event that last wrote it *)
   $eM : \text{address} \rightarrow (\text{value} \# 'reg \text{ event option}) \text{ option};$ 
  (* Per processor FIFO write buffers *)
   $eB : \text{proc} \rightarrow 'reg \text{ event list};$ 
  (* Which processor holds the lock *)
   $eL : \text{proc} \text{ option}$ 
 $\rrbracket$ 

```

$\text{evt_no_pending } b \ a = \neg(\exists e. \text{MEM } e \ b \wedge (\text{loc } e = \text{SOME } (\text{LOCATION_MEM } a)))$

$\text{evt_not_blocked } s \ p = (s.eL = \text{NONE}) \vee (s.eL = \text{SOME } p)$

A $\text{TAUEVT } ew$ label reports that ew is the event moved from the write buffer to the store. A $\text{REVT } er \ ew_opt$ label does read event er and reports that ew_opt was the last write event on that location. WEVT and BEVT labels both simply perform the given event. The event set argument to UNLOCKE and LOCKE labels is ignored by the machine; it is used only for bookkeeping during a later proof.

```

 $\text{evt\_machine\_label} =$ 
 $\text{TAUEVT of 'reg event}$ 
 $| \text{REVT of 'reg event 'reg event option}$ 
 $| \text{WEVT of 'reg event}$ 
 $| \text{BEVT of 'reg event}$ 
 $| \text{LOCKE of proc 'reg event set}$ 
 $| \text{UNLOCKE of proc 'reg event set}$ 

```

The transition relation of the event machine is defined by the following rules. Read each clause as “the event-annotated machine has a transition from state s to state s' labelled l if each conjunct preceding $\implies$ is satisfied”.

```

 $(\forall s \ a \ v \ p \ er \ ew\_opt.$ 
 $\text{clause\_name "evt-read-mem"} \wedge$ 
 $\text{evt\_not\_blocked } s \ p \wedge$ 

```

$(\text{proc } er = p) \wedge$
 $(er.action = \text{ACCESS } R \text{ (LOCATION_MEM } a)v) \wedge$
 $(s.eM \ a = \text{SOME } (v, ew_opt)) \wedge$
 $\text{evt_no_pending } (s.eB \ p)a$
 $\implies$
 $\text{evt_machine_trans } s \text{ (REVT } er \ ew_opt)s) \wedge$

$(\forall s \ a \ v \ p \ er \ ew.$
 $\text{clause_name "evt-read-buffer"} \wedge$
 $\text{evt_not_blocked } s \ p \wedge$
 $(\text{proc } er = p) \wedge$
 $(er.action = \text{ACCESS } R \text{ (LOCATION_MEM } a)v) \wedge$
 $(ew.action = \text{ACCESS } W \text{ (LOCATION_MEM } a)v) \wedge$
 $(\exists b_1 \ b_2. (s.eB \ p = b_1 ++[ew] ++b_2) \wedge \text{evt_no_pending } b_1 \ a)$
 $\implies$
 $\text{evt_machine_trans } s \text{ (REVT } er \text{ (SOME } ew))s) \wedge$

$(\forall s \ r \ v \ p \ er \ ew_opt.$
 $\text{clause_name "evt-read-reg"} \wedge$
 $(*\text{evt_not_blocked } s \ p \ /\backslash *)$
 $(\text{proc } er = p) \wedge$
 $(er.action = \text{ACCESS } R \text{ (LOCATION_REG } p \ r)v) \wedge$
 $(s.eR \ p \ r = \text{SOME } (v, ew_opt))$
 $\implies$
 $\text{evt_machine_trans } s \text{ (REVT } er \ ew_opt)s) \wedge$

$(\forall s \ a \ v \ p \ ew \ s'.$
 $\text{clause_name "evt-write-buffer"} \wedge$
 $(*\text{evt_not_blocked } s \ p \ /\backslash *)$
 $(\text{proc } ew = p) \wedge$
 $(ew.action = \text{ACCESS } W \text{ (LOCATION_MEM } a)v) \wedge$
 $(s' = \llbracket eR := s.eR;$
 $\quad eM := s.eM;$
 $\quad eB := (p \mapsto ew \in s.eB \ p)s.eB;$
 $\quad eL := s.eL \rrbracket)$
 $\implies$
 $\text{evt_machine_trans } s \text{ (WEVT } ew)s') \wedge$

$(\forall s \ a \ v \ p \ ew \ b \ s'.$
 $\text{clause_name "evt-write-mem"} \wedge$
 $\text{evt_not_blocked } s \ p \wedge$
 $(\text{proc } ew = p) \wedge$
 $(ew.action = \text{ACCESS } W \text{ (LOCATION_MEM } a)v) \wedge$
 $(s.eB \ p = b ++[ew]) \wedge$
 $(s' = \llbracket eR := s.eR;$
 $\quad eM := (a \mapsto \text{SOME } (v, \text{SOME } ew))s.eM;$
 $\quad eB := (p \mapsto b)s.eB;$
 $\quad eL := s.eL \rrbracket)$
 $\implies$
 $\text{evt_machine_trans } s \text{ (TAUEVT } ew)s') \wedge$

$(\forall s \ r \ v \ p \ ew \ s'.$
 $\text{clause_name "evt-write-reg"} \wedge$
 $(*\text{evt_not_blocked } s \ p \ /\backslash *)$
 $(\text{proc } ew = p) \wedge$
 $(ew.action = \text{ACCESS } W \text{ (LOCATION_REG } p \ r)v) \wedge$
 $(s' = \llbracket eR := (p \mapsto (r \mapsto \text{SOME } (v, \text{SOME } ew))(s.eR \ p))s.eR;$

$$\begin{aligned}
& eM := s.eM; \\
& eB := s.eB; \\
& eL := s.eL \rangle) \\
\Rightarrow & \\
& \text{evt_machine_trans } s \text{ (WEVT } ew)s') \wedge \\
& (\forall s \ p \ eb. \\
& \text{clause_name "evt-barrier" } \wedge \\
& (*\text{evt_not_blocked } s \ p \ /\ \ *) \\
& (\text{proc } eb = p) \wedge \\
& (eb.action = \text{BARRIER MFENCE}) \wedge \\
& (s.eB \ p = [])) \\
\Rightarrow & \\
& \text{evt_machine_trans } s \text{ (BEVT } eb)s) \wedge \\
& (\forall s \ eb. \\
& \text{clause_name "evt-nop" } \wedge \\
& (*\text{not_blocked } s \ p \ /\ \ *) \\
& ((eb.action = \text{BARRIER SFENCE}) \vee (eb.action = \text{BARRIER LFENCE})) \\
\Rightarrow & \\
& \text{evt_machine_trans } s \text{ (BEVT } eb)s) \wedge \\
& (\forall s \ p \ s' \ es. \\
& \text{clause_name "evt-lock" } \wedge \\
& (s.eL = \text{NONE}) \wedge \\
& (s.eB \ p = [])) \wedge \\
& (s' = \langle eR := s.eR; \\
& \quad eM := s.eM; \\
& \quad eB := s.eB; \\
& \quad eL := \text{SOME } p \rangle) \\
\Rightarrow & \\
& \text{evt_machine_trans } s \text{ (LOCKE } p \ es)s') \wedge \\
& (\forall s \ p \ s' \ es. \\
& \text{clause_name "evt-unlock" } \wedge \\
& (s.eL = \text{SOME } p) \wedge \\
& (s.eB \ p = [])) \wedge \\
& (s' = \langle eR := s.eR; \\
& \quad eM := s.eM; \\
& \quad eB := s.eB; \\
& \quad eL := \text{NONE} \rangle) \\
\Rightarrow & \\
& \text{evt_machine_trans } s \text{ (UNLOCKE } p \ es)s') \\
& \text{evt_machine_state_to_state_constraint } s = \\
& \lambda l. \\
& \text{case } l \text{ of} \\
& \quad \text{LOCATION_MEM } a \rightarrow \text{OPTION_MAP FST } (s.eM \ a) \\
& \quad \parallel \text{LOCATION_REG } p \ r \rightarrow \text{OPTION_MAP FST } (s.eR \ p \ r)
\end{aligned}$$

A path through the machine should only be considered if it reasonably corresponds to an event structure. The `okEpath` predicate expresses the necessary conditions. (`nth_label` and `PL` are from HOL's library for paths through a LTS. `nth_label` gets the `nth` label from the head of the path, and `PL` returns a (downward closed) set of all numbers that can be used to index into the path, which can be finite or infinite. `first` returns the first state in a path.)

(`get_orig_event (REVT  $e$   $\_$ ) = SOME  $e$ )  $\wedge$`

$(\text{get_orig_event } (\text{WEvt } e) = \text{SOME } e) \wedge$
 $(\text{get_orig_event } (\text{BEvt } e) = \text{SOME } e) \wedge$
 $(\text{get_orig_event } _ = \text{NONE})$

$\text{locked_segment } \text{path } i \ j \ p =$
 $j + 1 \in \text{PL } \text{path} \wedge$
 $i < j \wedge$
 $(\exists es. \text{nth_label } i \ \text{path} = \text{LOCKE } p \ es) \wedge$
 $(\exists es. \text{nth_label } j \ \text{path} = \text{UNLOCKE } p \ es) \wedge$
 $(\forall k \ es. i < k \wedge k < j \implies \text{nth_label } k \ \text{path} \neq \text{UNLOCKE } p \ es)$

$\text{okEpath } E \ \text{path} =$
 $(\text{* The REvt, WEvt and BEvt labels are exactly the set of events *})$
 $(E.\text{events} = \{e \mid \exists i. i + 1 \in \text{PL } \text{path} \wedge (\text{get_orig_event } (\text{nth_label } i \ \text{path}) = \text{SOME } e)\}) \wedge$
 $(\text{* No REvt, WEvt, or BEvt appears twice as a label *})$
 $(\forall i \ j \ e_1 \ e_2.$
 $\quad i + 1 \in \text{PL } \text{path} \wedge j + 1 \in \text{PL } \text{path} \wedge$
 $\quad (\text{get_orig_event } (\text{nth_label } i \ \text{path}) = \text{SOME } e_1) \wedge (\text{get_orig_event } (\text{nth_label } j \ \text{path}) = \text{SOME } e_2) \wedge$
 $\quad (e_1 = e_2)$
 $\quad \implies$
 $\quad (i = j)) \wedge$
 $(\text{* The REvt, WEvt, and BEvt parts of the trace follow po_iico *})$
 $(\forall (e_1, e_2) \in (\text{po_iico } E). \exists i \ j.$
 $\quad i < j \wedge j + 1 \in \text{PL } \text{path} \wedge$
 $\quad (\text{get_orig_event } (\text{nth_label } i \ \text{path}) = \text{SOME } e_1) \wedge (\text{get_orig_event } (\text{nth_label } j \ \text{path}) = \text{SOME } e_2)) \wedge$
 $(\text{* atomic sets of events are properly bracketed by lock/unlock pairs *})$
 $(\forall es \in (E.\text{atomicity}).$
 $\quad \exists i \ j \ p.$
 $\quad \text{locked_segment } \text{path } i \ j \ p \wedge$
 $\quad (\{e \mid e \in es \wedge e \in \text{mem_accesses } E\}$
 $\quad =$
 $\quad \{e \mid \exists k. i < k \wedge k < j \wedge$
 $\quad \quad (\text{get_orig_event } (\text{nth_label } k \ \text{path}) = \text{SOME } e) \wedge$
 $\quad \quad e \in \text{mem_accesses } E \wedge$
 $\quad \quad (\text{proc } e = p)\})$

Finally, we require all paths to satisfy the okMpath predicate which ensures the head of the path is in a starting state and that no write event stays in the buffer forever (okpath is part of HOL's path library and simply says that path is a path through the given LTS).

$\text{evt_machine_init_state } sc =$
 $\langle\langle \text{eR} := (\lambda p \ r. \text{OPTION_MAP } (\lambda v. (v, \text{NONE})) (sc \ (\text{LOCATION_REG } p \ r))) ;$
 $\text{eM} := (\lambda a. \text{OPTION_MAP } (\lambda v. (v, \text{NONE})) (sc \ (\text{LOCATION_MEM } a))) ;$
 $\text{eB} := (\lambda p. []);$
 $\text{eL} := \text{NONE} \rangle\rangle$

$\text{evt_is_init } s = \exists sc. s = \text{evt_machine_init_state } sc$

$\text{okMpath } \text{path} =$
 $\text{evt_is_init } (\text{first } \text{path}) \wedge$
 $\text{okpath } \text{evt_machine_trans } \text{path} \wedge$
 $\forall i \ e.$
 $i + 1 \in \text{PL } \text{path} \wedge (\text{nth_label } i \ \text{path} = \text{WEvt } e) \wedge \text{is_mem_access } e$
 $\implies$
 $\exists j. j + 1 \in \text{PL } \text{path} \wedge i < j \wedge (\text{nth_label } j \ \text{path} = \text{TAUEvt } e)$

To show that the two machines are equivalent, we define how labels and states are erased. `erase_label` simply removes the annotating information. `erase_state` is phrased as a relation, because, when erasing the write buffers, their type does not guarantee that an address/value pair can be extracted (although it will be the case for all states on path satisfying `okMpath`). The `annotated_labels` function returns the set of all labels that could be annotated versions of its input, given any potential write annotations.

$$\begin{aligned} &(\text{erase_label } (\text{TAUEVT } _) = \text{TAU}) \wedge \\ &(\text{erase_label } (\text{REVT } e _) = \text{EVT } (\text{proc } e)e.\text{action}) \wedge \\ &(\text{erase_label } (\text{WEVT } e) = \text{EVT } (\text{proc } e)e.\text{action}) \wedge \\ &(\text{erase_label } (\text{BEVT } e) = \text{EVT } (\text{proc } e)e.\text{action}) \wedge \\ &(\text{erase_label } (\text{LOCKE } p \text{ es}) = \text{LOCK } p) \wedge \\ &(\text{erase_label } (\text{UNLOCKE } p \text{ es}) = \text{UNLOCK } p) \end{aligned}$$

$$\begin{aligned} \text{erase_state } s \ s' = & \\ &(\forall p \ r. s'.R \ p \ r = \text{OPTION_MAP FST } (s.eR \ p \ r)) \wedge \\ &(\forall a. s'.M \ a = \text{OPTION_MAP FST } (s.eM \ a)) \wedge \\ &(\forall p. (\text{LENGTH } (s'.B \ p) = \text{LENGTH } (s.eB \ p)) \wedge \\ &\quad \forall n. n < \text{LENGTH } (s.eB \ p) \implies \\ &\quad \quad \exists e \ a \ v. (\text{EL } n \ (s.eB \ p) = e) \wedge \\ &\quad \quad \quad (\text{proc } e = p) \wedge \\ &\quad \quad \quad (e.\text{action} = \text{ACCESS W } (\text{LOCATION_MEM } a)v) \wedge \\ &\quad \quad \quad (\text{EL } n \ (s'.B \ p) = (a, v))) \wedge \\ &(s'.L = s.eL) \end{aligned}$$

$$\begin{aligned} &(\text{annotated_labels } \text{TAU } ew \ e_opt = \{\text{TAUEVT } ew\}) \wedge \\ &(\text{annotated_labels } (\text{EVT } p \ (\text{ACCESS } R \ l \ v))ew \ e_opt = \\ &\quad \{\text{REVT } e \ e_opt \mid e \mid (e.\text{action} = \text{ACCESS } R \ l \ v) \wedge (p = \text{proc } e)\}) \wedge \\ &(\text{annotated_labels } (\text{EVT } p \ (\text{ACCESS } W \ l \ v))ew \ e_opt = \\ &\quad \{\text{WEVT } e \mid (e.\text{action} = \text{ACCESS } W \ l \ v) \wedge (p = \text{proc } e)\}) \wedge \\ &(\text{annotated_labels } (\text{EVT } p \ (\text{BARRIER } b))ew \ e_opt = \\ &\quad \{\text{BEVT } e \mid (e.\text{action} = \text{BARRIER } b) \wedge (p = \text{proc } e)\}) \wedge \\ &(\text{annotated_labels } (\text{LOCK } p)ew \ e_opt = \\ &\quad \{\text{LOCKE } p \text{ es} \mid \text{es} \mid \mathbf{T}\}) \wedge \\ &(\text{annotated_labels } (\text{UNLOCK } p)ew \ e_opt = \\ &\quad \{\text{UNLOCKE } p \text{ es} \mid \text{es} \mid \mathbf{T}\}) \end{aligned}$$

The following erasure theorems ensure that the machines are equivalent.

Theorem 5

$$\forall sc. \text{erase_state } (\text{evt_machine_init_state } sc)(\text{machine_init_state } sc)$$

Proof sketch (full proof in HOL):

Immediate from the definitions. □

Theorem 6

$$\begin{aligned} &\forall s \ s'. \\ &\text{erase_state } s \ s' \\ &\implies \\ &(\text{evt_machine_state_to_state_constraint } s = \\ &\text{machine_state_to_state_constraint } s') \end{aligned}$$

Proof sketch (full proof in HOL):

Immediate from the definitions. □

Theorem 7

$\forall l \ l' \ ew \ er_opt.$
 $l \in \text{annotated_labels} \ l' \ ew \ er_opt$
 $\implies$
 $(l' = \text{erase_label } l)$

Proof sketch (full proof in HOL):
 Immediate from the definitions. □

Theorem 8

$\forall s_1 \ l \ s_2 \ s'_1.$
 $\text{evt_machine_trans } s_1 \ l \ s_2 \wedge$
 $\text{erase_state } s_1 \ s'_1$
 $\implies$
 $\exists s'_2.$
 $\text{erase_state } s_2 \ s'_2 \wedge$
 $\text{machine_trans } s'_1 \ (\text{erase_label } l) s'_2$

Proof sketch (full proof in HOL):
 By case analysis on which transition the event machine takes. □

Theorem 9

$\forall s'_1 \ l' \ s'_2 \ s_1.$
 $\text{machine_trans } s'_1 \ l' \ s'_2 \wedge$
 $\text{erase_state } s_1 \ s'_1$
 $\implies$
 $\exists ew \ ew_opt.$
 $\forall l. l \in \text{annotated_labels} \ l' \ ew \ ew_opt \implies$
 $\exists s_2. \text{erase_state } s_2 \ s'_2 \wedge \text{evt_machine_trans } s_1 \ l \ s_2$

For this theorem, some of the annotation is “output” coming from the event machine itself: any possible ew and ew_opt events that should be recorded on TAUEVT or REVT labels. However, the use of annotated_labels ensures that the event machine will be able to progress with any other annotating information (such as the event’s iiid or eiid).

Proof sketch (full proof in HOL):
 By case analysis on which transition the machine takes. □

E.2 Linear valid executions

The following definition of $\text{linear_valid_execution}$ is equivalent to the definition of valid_execution . It differs only by requiring a linear order over memory access events, instead of a partial order, and by omitting two properties that are redundant given a linear ordering (those are: that memory writes have a linear order, and that only finitely many memory reads are unrelated to a same-location memory write).

$\text{linear_valid_execution } E \ X =$
 $\text{linear_order } X.\text{memory_order} \ (\text{mem_accesses } E) \wedge$
 $\text{finite_prefixes } X.\text{memory_order} \ (\text{mem_accesses } E) \wedge$
 $(\forall er \in (\text{mem_reads } E). \forall e \in (\text{mem_accesses } E). (er, e) \in \text{po_iico } E \implies (er, e) \in X.\text{memory_order}) \wedge$
 $(\forall ew_1 \ ew_2 \in (\text{mem_writes } E). (ew_1, ew_2) \in \text{po_iico } E \implies (ew_1, ew_2) \in X.\text{memory_order}) \wedge$
 $(\forall ew \in (\text{mem_writes } E). \forall er \in (\text{mem_reads } E). \forall ef \in (\text{mfences } E).$
 $\quad (ew, ef) \in \text{po_iico } E \wedge (ef, er) \in \text{po_iico } E \implies (ew, er) \in X.\text{memory_order}) \wedge$
 $(\forall e_1 \ e_2 \in (\text{mem_accesses } E). \forall es \in (E.\text{atomicity}).$
 $\quad (e_1 \in es \vee e_2 \in es) \wedge (e_1, e_2) \in \text{po_iico } E$
 $\quad \implies$
 $\quad (e_1, e_2) \in X.\text{memory_order}) \wedge$
 $(\forall es \in (E.\text{atomicity}). \forall e \in (\text{mem_accesses } E \setminus es).$

$$\begin{aligned}
& (\forall e' \in (es \cap \text{mem_accesses } E). (e, e') \in X.\text{memory_order}) \vee \\
& (\forall e' \in (es \cap \text{mem_accesses } E). (e', e) \in X.\text{memory_order})) \wedge \\
& X.\text{rfmap} \in \text{reads_from_map_candidates } E \wedge \\
& \text{check_rfmap_written } E X \wedge \\
& \text{check_rfmap_initial } E X
\end{aligned}$$

Theorem 10

$$\forall E X. \text{linear_valid_execution } E X \implies \text{valid_execution } E X$$

Proof sketch (full proof in HOL):

Immediate from the definitions. □

Theorem 11

$$\begin{aligned}
& \forall E X \text{ memory_order}'. \\
& \text{valid_execution } E X \wedge \\
& X.\text{memory_order} \subseteq \text{memory_order}' \wedge \\
& \text{linear_order } \text{memory_order}' (\text{mem_accesses } E) \wedge \\
& \text{finite_prefixes } \text{memory_order}' (\text{mem_accesses } E) \wedge \\
& (\forall er \in (\text{mem_reads } E). \forall ew \in (\text{mem_writes } E). \\
& (\text{loc } er = \text{loc } ew) \wedge (ew, er) \in \text{memory_order}' \implies (ew, er) \in X.\text{memory_order}) \\
& \implies \\
& \text{linear_valid_execution } E (X \oplus \text{memory_order} := \text{memory_order}')
\end{aligned}$$

This theorem says that any linearization of memory order is still a valid memory ordering, as long as it obeys the following additional condition. If ew is a memory write and er is a memory read from the same location, then if the original memory ordering does not relate ew and er , the linearization must make er come before ew . Otherwise, in the linearization, er might see ew 's write, but this should only happen if ew precedes er in the original memory order.

Proof sketch (full proof in HOL):

By unfolding all of the relevant definitions. □

Theorem 12

$$\begin{aligned}
& \forall E X. \\
& \text{well_formed_event_structure } E \wedge \\
& \text{valid_execution } E X \\
& \implies \\
& \exists \text{memory_order}'. \\
& X.\text{memory_order} \subseteq \text{memory_order}' \wedge \\
& \text{linear_order } \text{memory_order}' (\text{mem_accesses } E) \wedge \\
& \text{finite_prefixes } \text{memory_order}' (\text{mem_accesses } E) \wedge \\
& (\forall er \in (\text{mem_reads } E). \forall ew \in (\text{mem_writes } E). \\
& (\text{loc } er = \text{loc } ew) \wedge (ew, er) \in \text{memory_order}' \implies (ew, er) \in X.\text{memory_order})
\end{aligned}$$

Proof sketch (full proof in HOL):

Define the following relation, `complete_memory_order`, that extends memory order with any missing same-location memory read to memory write dependencies. Note that the definition below also adds in enough other edges to remain transitive.

$$\begin{aligned}
& \text{complete_memory_order } E \text{ memory_order} = \\
& \text{memory_order} \cup \\
& \{(e_1, e_2) \mid \exists ew \text{ er}. (ew, er) \notin \text{memory_order} \wedge (er, ew) \notin \text{memory_order} \wedge \\
& \quad ew \in \text{mem_writes } E \wedge er \in \text{mem_reads } E \wedge (\text{loc } ew = \text{loc } er) \wedge \\
& \quad (e_1, er) \in \text{memory_order} \wedge (ew, e_2) \in \text{memory_order}\}
\end{aligned}$$

We have the following theorem (proved in HOL) that allows us to linearize partial orders, maintaining the finite prefixes property.

$$\forall r \ s. \text{countable } s \wedge \text{partial_order } r \ s \wedge \text{finite_prefixes } r \ s \implies \exists r'. \text{linear_order } r' \ s \wedge \text{finite_prefixes } r' \ s \wedge r \subseteq r'$$

The rest of the theorem proceeds by unfolding definitions, noting that `valid_execution` requires `X.memory_order` to have finite prefixes, which allows us to prove that `complete_memory_order` has finite prefixes as well. □

E.3 Abstract machine model/axiomatic model equivalence

E.3.1 Abstract machine validity

The `path_to_X` function produces an execution witness whose (linear) memory ordering comes from the labels of its input path, whose reads-from map comes from the annotations on the labels, and whose initial state corresponds to the first state in the path.

```
(get_mem_event (TAUEVT e) = SOME e) ∧
(get_mem_event (REVT e _) =
  if is_mem_access e then
    SOME e
  else
    NONE) ∧
(get_mem_event _ = NONE)
```

```
path_to_X path =
  ⟨⟦ memory_order :=
    {(e1, e2) |
      ∃ j i l1 l2.
        j + 1 ∈ PL path ∧ i ≤ j ∧
        (nth_label i path = l1) ∧ (nth_label j path = l2) ∧
        (get_mem_event l1 = SOME e1) ∧ (get_mem_event l2 = SOME e2)};
    rfmap := {(ew, er) | (ew, er) | ∃ i. i + 1 ∈ PL path ∧ (nth_label i path = REVT er (SOME ew))};
    initial_state := evt_machine_state_to_state_constraint (first path)⟩⟩
```

Theorem 13

```
∀ E path.
  well_formed_event_structure E ∧
  okEpath E path ∧
  okMpath path
  ⟹
  linear_valid_execution E (path_to_X path)
```

Proof sketch (full proof in HOL):

By unfolding the definitions in `linear_valid_execution` and `okEpath`, and by using the following theorems that express how the labels of a valid path through the `lts` relate to each other. Each of these theorems is a “safety” style theorem that is proved (in HOL) by induction on the location of the largest label index mentioned. Some of them require nested inductions, and some rely on the being no duplicate writes (which `okEpath` ensures).

```
no_dup_writes path =
  ∀ i j ew.
    SUC i ∈ PL path ∧
    SUC j ∈ PL path ∧
    (nth_label i path = WEVT ew) ∧
    (nth_label j path = WEVT ew)
  ⟹
```

$(i = j)$

Each TAU EVT comes from a WEVT.

Theorem 14

$\forall path\ j\ e.$
 $okMpath\ path \wedge$
 $j + 1 \in PL\ path \wedge$
 $(nth_label\ j\ path = TAU EVT\ e)$
 $\implies$
 $\exists i. i < j \wedge (nth_label\ i\ path = WEVT\ e) \wedge is_mem_access\ e$

Each same processor pair of TAU EVTs comes from a pair of WEVTs in order.

Theorem 15

$\forall path\ i\ j\ ew_1\ ew_2.$
 $okMpath\ path \wedge$
 $j + 1 \in PL\ path \wedge$
 $i < j \wedge$
 $(proc\ ew_1 = proc\ ew_2) \wedge$
 $(nth_label\ i\ path = TAU EVT\ ew_1) \wedge$
 $(nth_label\ j\ path = TAU EVT\ ew_2)$
 $\implies$
 $\exists k\ l. k < i \wedge l < j \wedge k < l \wedge$
 $(nth_label\ k\ path = WEVT\ ew_1) \wedge (nth_label\ l\ path = WEVT\ ew_2)$

REVT labels have reads on them.

Theorem 16

$\forall path\ i\ e_1\ e_2.$
 $okMpath\ path \wedge$
 $i + 1 \in PL\ path \wedge$
 $(nth_label\ i\ path = REVT\ e_1\ e_2)$
 $\implies$
 $\exists l\ v. e_1.action = ACCESS\ R\ l\ v$

WEVT labels have writes on them.

Theorem 17

$\forall path\ i\ e.$
 $okMpath\ path \wedge$
 $i + 1 \in PL\ path \wedge$
 $(nth_label\ i\ path = WEVT\ e)$
 $\implies$
 $\exists l\ v. e.action = ACCESS\ W\ l\ v$

BEVT labels have barriers on them.

Theorem 18

$\forall path\ i\ e.$
 $okMpath\ path \wedge$
 $i + 1 \in PL\ path \wedge$
 $(nth_label\ i\ path = BEVT\ e)$
 $\implies$
 $\exists f. e.action = BARRIER\ f$

Each same processor pair of WEVT labels is followed by a pair of TAU EVT labels in order. This relies on the liveness property in okMpath.

Theorem 19

$$\begin{aligned}
& \forall path \ i \ j \ ew_1 \ ew_2. \\
& \text{okMpath } path \wedge \\
& j + 1 \in \text{PL } path \wedge \\
& i < j \wedge \\
& (\text{proc } ew_1 = \text{proc } ew_2) \wedge \\
& (\text{nth_label } i \ path = \text{WEVT } ew_1) \wedge \\
& (\text{nth_label } j \ path = \text{WEVT } ew_2) \wedge \\
& \text{is_mem_access } ew_1 \wedge \\
& \text{is_mem_access } ew_2 \wedge \\
& \text{no_dup_writes } path \\
& \implies \\
& \exists k \ l. l + 1 \in \text{PL } path \wedge k < l \wedge (\text{nth_label } k \ path = \text{TAUEVT } ew_1) \wedge (\text{nth_label } l \ path = \text{TAUEVT } ew_2)
\end{aligned}$$

TAUEVTs occur before a label that would require an empty buffer.

Theorem 20

$$\begin{aligned}
& \forall path \ i \ j \ ef \ ew \ p \ es. \\
& \text{okMpath } path \wedge \\
& j + 1 \in \text{PL } path \wedge \\
& i < j \wedge \\
& (((\text{nth_label } j \ path = \text{BEVT } ef) \wedge (\text{proc } ef = \text{proc } ew) \wedge (ef.action = \text{BARRIER MFENCE})) \vee \\
& (\text{nth_label } j \ path = \text{UNLOCKE } (\text{proc } ew)es) \vee \\
& (\text{nth_label } j \ path = \text{LOCKE } (\text{proc } ew)es)) \wedge \\
& (\text{nth_label } i \ path = \text{WEVT } ew) \wedge \\
& \text{is_mem_access } ew \\
& \implies \\
& \exists k. k < j \wedge (\text{nth_label } k \ path = \text{TAUEVT } ew)
\end{aligned}$$

A REVT label reads from the most recent TAUEVT label if there are no intermediate same processor, same location WEVT labels. Otherwise it reads from the most recent of those.

Theorem 21

$$\begin{aligned}
& \forall path \ j \ er \ ew. \\
& \text{okMpath } path \wedge \\
& j + 1 \in \text{PL } path \wedge \\
& \text{no_dup_writes } path \wedge \\
& (\text{nth_label } j \ path = \text{REVT } er \ (\text{SOME } ew)) \\
& \implies \\
& (\text{loc } er = \text{loc } ew) \wedge \\
& (\text{value_of } er = \text{value_of } ew) \wedge \\
& (((\text{proc } er = \text{proc } ew) \wedge \\
& \exists i. \\
& i < j \wedge \\
& (\text{nth_label } i \ path = \text{WEVT } ew) \wedge \\
& (\forall k. i < k \wedge k < j \implies \text{nth_label } k \ path \neq \text{TAUEVT } ew) \wedge \\
& (\forall k \ ew'. i < k \wedge k < j \wedge (\text{nth_label } k \ path = \text{WEVT } ew') \wedge (\text{proc } er = \text{proc } ew') \implies \\
& \quad \text{loc } ew \neq \text{loc } ew')) \vee \\
& (\exists i. \\
& i < j \wedge \\
& (\text{nth_label } i \ path = \text{TAUEVT } ew) \wedge \\
& (\forall k \ ew'. k < j \wedge (\text{nth_label } k \ path = \text{WEVT } ew') \wedge (\text{proc } er = \text{proc } ew') \wedge \\
& \quad (\text{loc } ew' = \text{loc } er) \implies \\
& \quad \exists l. k < l \wedge l < j \wedge (\text{nth_label } l \ path = \text{TAUEVT } ew')) \wedge \\
& (\forall k \ ew'. i < k \wedge k < j \wedge (\text{nth_label } k \ path = \text{TAUEVT } ew') \implies \text{loc } ew \neq \text{loc } ew'))))
\end{aligned}$$

A REVT that reads from the initial state is not preceded by write labels to its location.

Theorem 22

$$\begin{aligned}
& \forall path \ i \ er. \\
& \text{okMpath } path \wedge \\
& i + 1 \in \text{PL } path \wedge \\
& (\text{nth_label } i \ path = \text{REVT } er \ \text{NONE}) \\
& \implies \\
& (\text{value_of } er = \text{evt_machine_state_to_state_constraint } (\text{first } path)(\text{the } (\text{loc } er))) \wedge \\
& \forall j \ ew. \\
& j < i \wedge \\
& ((\text{nth_label } j \ path = \text{TAUEVT } ew) \vee ((\text{nth_label } j \ path = \text{WEVT } ew) \wedge (\text{proc } er = \text{proc } ew))) \\
& \implies \\
& \text{loc } ew \neq \text{loc } er
\end{aligned}$$

Memory read and write labels inside of a locked segment are on the processor that is locked.

Theorem 23

$$\begin{aligned}
& \forall path \ i \ j \ k \ p \ e \ e'. \\
& \text{okMpath } path \wedge \\
& \text{locked_segment } path \ i \ k \ p \wedge \\
& i < j \wedge \\
& j < k \wedge \\
& \text{is_mem_access } e \wedge \\
& ((\text{nth_label } j \ path = \text{REVT } e \ e') \vee \\
& (\text{nth_label } j \ path = \text{TAUEVT } e)) \\
& \implies \\
& (\text{proc } e = p)
\end{aligned}$$

WEVT and TAUEVT labels with the same write occur atomically with respect to locks on the write's processor.

Theorem 24

$$\begin{aligned}
& \forall path \ i \ j \ ew. \\
& \text{okMpath } path \wedge \\
& \text{locked_segment } path \ i \ j \ (\text{proc } ew) \wedge \\
& \text{is_mem_access } ew \\
& \implies \\
& ((\exists k. i < k \wedge k < j \wedge (\text{nth_label } k \ path = \text{WEVT } ew)) = \\
& (\exists k. i < k \wedge k < j \wedge (\text{nth_label } k \ path = \text{TAUEVT } ew)))
\end{aligned}$$

□

E.3.2 Abstract machine completeness

The proof that the abstract machine can perform any valid execution is currently partly in HOL and partly by hand. We intend to fully mechanize the parts that currently by hand.

We first need several additional definitions for grouping and operating on sets of labels.

$$\begin{aligned}
\text{memL } E \ X = & \\
& \{\text{TAUEVT } e \mid e \in \text{mem_writes } E\} \cup \\
& \{\text{REVT } er \ \text{NONE} \mid er \in \text{mem_reads } E \wedge er \notin \text{range } X.\text{rfmap}\} \cup \\
& \{\text{REVT } er \ (\text{SOME } ew) \mid er \in \text{mem_reads } E \wedge (ew, er) \in X.\text{rfmap}\}
\end{aligned}$$

$$\begin{aligned}
\text{localL } E \ X = & \\
& \{\text{REVT } er \ \text{NONE} \mid er \in \text{reads } E \wedge er \notin \text{range } X.\text{rfmap}\} \cup \\
& \{\text{REVT } er \ (\text{SOME } ew) \mid er \in \text{reads } E \wedge (ew, er) \in X.\text{rfmap}\} \cup \\
& \{\text{WEVT } e \mid e \in \text{writes } E\} \cup \\
& \{\text{BEVT } e \mid e \in \text{fences } E\}
\end{aligned}$$

$$\text{proc_es } es = \{\text{proc } e \mid e \in es\}$$

$$\begin{aligned} \text{lockL } E X = & \\ \{ \text{LOCKE } p \text{ es} \mid & \text{es} \in E.\text{atomicity} \wedge (p \in \text{proc_es } \text{es}) \} \cup \\ \{ \text{UNLOCKE } p \text{ es} \mid & \text{es} \in E.\text{atomicity} \wedge (p \in \text{proc_es } \text{es}) \} \end{aligned}$$

$$\begin{aligned} \text{allL } E X = & \\ \text{memL } E X \cup & \\ \text{localL } E X \cup & \\ \text{lockL } E X & \end{aligned}$$

$$\begin{aligned} (\text{Le } (\text{TAUEVT } e) = \text{SOME } e) \wedge & \\ (\text{Le } (\text{REVT } e \text{ } _) = \text{SOME } e) \wedge & \\ (\text{Le } (\text{WEVT } e) = \text{SOME } e) \wedge & \\ (\text{Le } (\text{BEVT } e) = \text{SOME } e) \wedge & \\ (\text{Le } _ = \text{NONE}) & \end{aligned}$$

$$\begin{aligned} (\text{Les } (\text{LOCKE } p \text{ es}) = \text{SOME } \text{es}) \wedge & \\ (\text{Les } (\text{UNLOCKE } p \text{ es}) = \text{SOME } \text{es}) \wedge & \\ (\text{Les } _ = \text{NONE}) & \end{aligned}$$

Our overall strategy is to characterize an ordering on labels `label_order` such that any sequence of labels that follows this ordering has a path through the machine (that is also an `okEpath` and `okMpath`). We then show that the properties in `linear_valid_execution` ensure that such a sequence exists. The first part of the proof is not currently mechanized in HOL, but is proved by hand. We have proved in HOL that there exists a partial ordering over labels `lo` such that `label_order` $\subseteq$ `lo` (which via hand proof has finite prefixes), and therefore there exists a sequence consistent with `label_order` (via the linearization theorem mentioned in Section E.2).

$$\begin{aligned} \text{label_order } E X = & \\ \{ (l, l') \mid l \in \text{memL } E X \wedge l' \in \text{memL } E X \wedge & \\ (\text{the } (\text{Le } l), \text{the } (\text{Le } l')) \in X.\text{memory_order} \} \cup & \\ \{ (l, l') \mid l \in \text{localL } E X \wedge l' \in \text{localL } E X \wedge & \\ (\text{the } (\text{Le } l), \text{the } (\text{Le } l')) \in \text{po_iico } E \} \cup & \\ \{ (\text{WEVT } e, \text{TAUEVT } e) \mid \text{WEVT } e \in \text{allL } E X \wedge \text{TAUEVT } e \in \text{allL } E X \} \cup & \\ \{ (\text{TAUEVT } e, \text{BEVT } e') \mid \text{TAUEVT } e \in \text{allL } E X \wedge \text{BEVT } e' \in \text{allL } E X \wedge & \\ e' \in \text{mfences } E \wedge (e, e') \in \text{po_iico } E \} \cup & \\ \{ (\text{TAUEVT } e, \text{LOCKE } p \text{ es}) \mid \text{TAUEVT } e \in \text{allL } E X \wedge \text{LOCKE } p \text{ es} \in \text{allL } E X \wedge & \\ e \notin \text{es} \wedge \exists e'. e' \in \text{es} \wedge e' \in \text{mem_accesses } E \wedge (e, e') \in \text{po_iico } E \} \cup & \\ \{ (\text{LOCKE } p \text{ es}, l) \mid \text{LOCKE } p \text{ es} \in \text{allL } E X \wedge l \in \text{allL } E X \wedge & \\ \exists e. (\text{Le } l = \text{SOME } e) \wedge e \in \text{es} \wedge e \in \text{mem_accesses } E \} \cup & \\ \{ (l, \text{UNLOCKE } p \text{ es}) \mid l \in \text{memL } E X \wedge \text{UNLOCKE } p \text{ es} \in \text{allL } E X \wedge & \\ \exists e. (\text{Le } l = \text{SOME } e) \wedge e \in \text{es} \} \cup & \\ \{ (\text{UNLOCKE } p \text{ es}, l) \mid \text{UNLOCKE } p \text{ es} \in \text{allL } E X \wedge l \in \text{allL } E X \wedge & \\ \exists e e'. (\text{Le } l = \text{SOME } e) \wedge e' \in \text{es} \wedge e \notin \text{es} \wedge & \\ (e', e) \in X.\text{memory_order} \wedge & \\ (l \in \text{memL } E X \vee (e \in \text{mem_accesses } E \wedge (\text{proc } e = p))) \} \cup & \\ \{ (l, \text{LOCKE } p \text{ es}) \mid l \in \text{memL } E X \wedge \text{LOCKE } p \text{ es} \in \text{allL } E X \wedge & \\ \exists e e'. (\text{Le } l = \text{SOME } e) \wedge e' \in \text{es} \wedge e \notin \text{es} \wedge & \\ (e, e') \in X.\text{memory_order} \} \cup & \\ \{ (\text{UNLOCKE } p \text{ es}, \text{LOCKE } p' \text{ es}') \mid \text{UNLOCKE } p \text{ es} \in \text{allL } E X \wedge \text{LOCKE } p' \text{ es}' \in \text{allL } E X \wedge & \\ \text{es} \neq \text{es}' \wedge \exists e e'. e \in \text{es} \wedge e' \in \text{es}' \wedge (e, e') \in X.\text{memory_order} \} \cup & \\ \{ (\text{WEVT } e, \text{WEVT } e') \mid \text{WEVT } e \in \text{allL } E X \wedge \text{WEVT } e' \in \text{allL } E X \wedge & \\ (e, e') \in X.\text{memory_order} \wedge (\text{proc } e = \text{proc } e') \} & \end{aligned}$$

`lo1`–`lo4` correspond to clauses of `label_order`, and `lo_events` composes them to form a partial order. The proof that `lo_events` is a transitive and antisymmetric uses several lemmas that collapse particular relation compositions into shorter ones, and then relies on the transitivity and antisymmetry of the underlying memory and program orders. For example,

$$\forall E X l_1 l_2 l_3 l_4. \text{well_formed_event_structure } E \wedge \text{linear_valid_execution } E X \wedge (l_1, l_2) \in \text{lo4 } E X \wedge (l_2, l_3) \in \text{lo2 } E X \wedge (l_3, l_4) \in \text{lo3 } E X \implies (l_1, l_4) \in \text{lo1 } E X$$

$$\begin{aligned} \text{lo1 } E X = \\ \{(l, l') \mid l \in \text{memL } E X \wedge l' \in \text{memL } E X \wedge \\ (\mathbf{the} (\text{Le } l), \mathbf{the} (\text{Le } l')) \in X.\text{memory_order}\} \end{aligned}$$

$$\begin{aligned} \text{lo2 } E X = \\ \{(l, l') \mid l \in \text{localL } E X \wedge l' \in \text{localL } E X \wedge \\ (\mathbf{the} (\text{Le } l), \mathbf{the} (\text{Le } l')) \in \text{po_iico } E\} \cup \\ \{(\text{WEVT } e, \text{WEVT } e') \mid \text{WEVT } e \in \text{allL } E X \wedge \text{WEVT } e' \in \text{allL } E X \wedge \\ (e, e') \in X.\text{memory_order} \wedge (\text{proc } e = \text{proc } e')\} \cup \\ \{(l, \text{WEVT } e) \mid l \in \text{localL } E X \wedge \text{WEVT } e \in \text{allL } E X \wedge \\ \exists e'. \text{WEVT } e' \in \text{localL } E X \wedge (\mathbf{the} (\text{Le } l), e') \in \text{po_iico } E \wedge \\ (e', e) \in X.\text{memory_order} \wedge (\text{proc } e' = \text{proc } e)\} \end{aligned}$$

$$\begin{aligned} \text{lo3 } E X = \\ \{(\text{WEVT } e, \text{TAUEVT } e) \mid \text{WEVT } e \in \text{allL } E X \wedge \text{TAUEVT } e \in \text{allL } E X\} \end{aligned}$$

$$\begin{aligned} \text{lo4 } E X = \\ \{(\text{TAUEVT } e, \text{BEVT } e') \mid \text{TAUEVT } e \in \text{allL } E X \wedge \text{BEVT } e' \in \text{allL } E X \wedge \\ e' \in \text{mfences } E \wedge (e, e') \in \text{po_iico } E\} \end{aligned}$$

$$\begin{aligned} \text{lo_events } E X = \\ \text{lo1 } E X \cup \text{lo2 } E X \cup \\ \text{lo1 } E X \circ \text{lo2 } E X \cup \\ \text{lo2 } E X \circ \text{lo1 } E X \cup \\ \text{lo2 } E X \circ \text{lo1 } E X \circ \text{lo2 } E X \cup \\ \text{lo2 } E X \circ \text{lo3 } E X \circ \text{lo1 } E X \cup \\ \text{lo2 } E X \circ \text{lo3 } E X \circ \text{lo1 } E X \circ \text{lo2 } E X \cup \\ \text{lo1 } E X \circ \text{lo4 } E X \circ \text{lo2 } E X \cup \\ \text{lo2 } E X \circ \text{lo1 } E X \circ \text{lo4 } E X \circ \text{lo2 } E X \cup \\ \text{lo2 } E X \circ \text{lo3 } E X \circ \text{lo1 } E X \circ \text{lo4 } E X \circ \text{lo2 } E X \end{aligned}$$

lo5 defines a partial order over just the lock labels, using the above proof technique of collapsing relational compositions.

$$\begin{aligned} \text{lo5 } E X = \\ \{(\text{LOCKE } p \text{ es}, \text{LOCKE } p \text{ es}) \mid \text{LOCKE } p \text{ es} \in \text{lockL } E X\} \cup \\ \{(\text{UNLOCKE } p \text{ es}, \text{UNLOCKE } p \text{ es}) \mid \text{LOCKE } p \text{ es} \in \text{lockL } E X\} \cup \\ \{(l1, l2) \mid l1 \in \text{lockL } E X \wedge l2 \in \text{lockL } E X \wedge \\ \mathbf{the} (\text{Les } l1) \neq \mathbf{the} (\text{Les } l2) \wedge \\ \exists e_1 e_2. e_1 \in \mathbf{the} (\text{Les } l1) \wedge e_2 \in \mathbf{the} (\text{Les } l2) \wedge \\ (e_1, e_2) \in X.\text{memory_order}\} \cup \\ \{(\text{LOCKE } p \text{ es}, \text{UNLOCKE } p \text{ es}) \mid \text{LOCKE } p \text{ es} \in \text{lockL } E X \wedge \text{UNLOCKE } p \text{ es} \in \text{lockL } E X\} \end{aligned}$$

lo6–lo8 express how lock labels relate to other labels, to enforce the atomicity guarantees required of a locked instruction.

$$\begin{aligned} \text{lo6 } E X = \\ \{(\text{LOCKE } p \text{ es}, l) \mid \text{LOCKE } p \text{ es} \in \text{allL } E X \wedge l \in \text{allL } E X \wedge \\ \exists e. (\text{Le } l = \text{SOME } e) \wedge e \in \text{es} \wedge e \in \text{mem_accesses } E\} \end{aligned}$$

$$\begin{aligned} \text{lo7 } E X = \\ \{(l, \text{UNLOCKE } p \text{ es}) \mid l \in \text{memL } E X \wedge \text{UNLOCKE } p \text{ es} \in \text{allL } E X \wedge \\ \exists e. (\text{Le } l = \text{SOME } e) \wedge e \in \text{es}\} \end{aligned}$$

$$\begin{aligned} \text{lo8 } E X = & \{ (\text{UNLOCKE } p \text{ es}, l) \mid \text{UNLOCKE } p \text{ es} \in \text{allL } E X \wedge l \in \text{allL } E X \wedge \\ & \exists e e'. (l = \text{SOME } e) \wedge e' \in \text{es} \wedge e \notin \text{es} \wedge \\ & (e', e) \in X.\text{memory_order} \wedge \\ & (l \notin \text{memL } E X \implies ((\text{proc } e = p) \wedge e \in \text{mem_accesses } E)) \} \end{aligned}$$

$$\text{lo68 } E X = \text{lo6 } E X \cup \text{lo8 } E X$$

$$\text{lo79 } E X = \text{lo7 } E X \cup \text{lo9 } E X$$

Finally, lo defines a partial order over all labels in allL $E X$.

$$\begin{aligned} \text{lo } E X = & \text{lo_events } E X \cup \\ & \text{lo5 } E X \cup \\ & \text{lo5 } E X \circ \text{lo68 } E X \circ \text{lo_events } E X \cup \\ & \text{lo5 } E X \circ \text{lo68 } E X \circ \text{lo_events } E X \cup \\ & \text{lo_events } E X \circ \text{lo79 } E X \circ \text{lo5 } E X \cup \\ & \text{lo_events } E X \circ \text{lo79 } E X \circ \text{lo5 } E X \circ \text{lo68 } E X \circ \text{lo_events } E X \end{aligned}$$

Theorem 25

$$\begin{aligned} & \forall E X. \\ & \text{well_formed_event_structure } E \wedge \text{linear_valid_execution } E X \\ & \implies \\ & \text{partial_order } (\text{lo } E X)(\text{allL } E X) \end{aligned}$$

Proof sketch (full proof in HOL): The proof follows the methodology introduced above of collapsing certain relation compositions to other, shorter ones. The most intricate cases involve showing that there is no cycle formed between lo_events and lo6–lo9, since lock labels are not directly apparent in the execution witness. $\square$

Theorem 26

$$\begin{aligned} & \forall E X. \\ & \text{well_formed_event_structure } E \wedge \text{linear_valid_execution } E X \\ & \implies \\ & \text{label_order } E X \subseteq \text{lo } E X \end{aligned}$$

Proof sketch (full proof in HOL): Immediate by expanding the definitions. $\square$

Theorem 27

$$\begin{aligned} & \forall E X. \text{well_formed_event_structure } E \wedge \text{linear_valid_execution } E X \\ & \implies \\ & \exists \text{path}. \text{okEpath } E \text{ path} \wedge \text{okMpath } \text{path} \end{aligned}$$

Proof sketch: Create a sequence from label_order and assume that the machine has gone i steps through the sequence, and show that it can take one more by cases on the next label. $\square$

E.4 Executable checker

The executable checker essentially replaces sets with lists and quantification with iteration over the lists.

$$\begin{aligned} \text{ch_event_structure} = & \llbracket \text{ch_procs} : \text{proc list}; \\ & \text{ch_events} : ('reg \text{ event})\text{list}; \\ & \text{ch_intra_causality} : ('reg \text{ event})\text{ch_reln}; \\ & \text{ch_atomicity} : ('reg \text{ event})\text{list list} \rrbracket \end{aligned}$$

$$\text{ch_execution_witness} = \llbracket (* \text{ the memory order is the transitive closure of the pairs in ch_memory_order } *) \rrbracket$$

```

ch_memory_order : ('reg event)ch_reln;
ch_rfmap : 'reg event → 'reg event option;
ch_initial_state : 'reg location → value option]

```

```

cis_mem_access e =
case e.action of
  ACCESS d (LOCATION_MEM a)v → T
  || _ → F

```

```

is_mem_read e =
case e.action of
  ACCESS R (LOCATION_MEM a)v → T
  || _ → F

```

```

is_mem_write e =
case e.action of
  ACCESS W (LOCATION_MEM a)v → T
  || _ → F

```

```

is_read e =
case e.action of
  ACCESS R l v → T
  || _ → F

```

```

is_write e =
case e.action of
  ACCESS W l v → T
  || _ → F

```

```

is_barrier e =
case e.action of
  BARRIER MFENCE → T
  || _ → F

```

```

check_po_iico intra e1 e2 =
if proc e1 = proc e2 then
  if e1.iuid.poi < e2.iuid.poi then
    T
  else if e1.iuid.poi = e2.iuid.poi then
    e1 ≠ e2 ∧ MEM (e1, e2)intra
  else
    F
else
  F

```

```

check_po_iico_in_mo intra mo e1 e2 =
if check_po_iico intra e1 e2 then
  MEM (e1, e2)mo
else
  T

```

```

barrier_separated intra barriers e1 e2 =

```

```

(proc  $e_1 = \text{proc } e_2$ )  $\wedge$ 
EXISTS ( $\lambda eb. \text{check\_po\_iico } \text{intra } e_1 \text{ } eb \wedge \text{check\_po\_iico } \text{intra } eb \text{ } e_2$ )
  barriers

```

```

previous_writes1  $er \text{ } r =$ 
MAP FST (FILTER ( $\lambda(ew, er'). (er' = er) \wedge \text{is\_write } ew \wedge (\text{loc } ew = \text{loc } er)$ ))  $r$ 

```

```

previous_writes2  $er \text{ intra } es =$ 
FILTER ( $\lambda ew. (\text{loc } ew = \text{loc } er) \wedge \text{check\_po\_iico } \text{intra } ew \text{ } er$ )  $es$ 

```

```

check_maximal1  $x \text{ } xs \text{ } r =$ 
MEM  $x \text{ } xs \wedge$ 
EVERY ( $\lambda x'. \text{if } x \neq x' \text{ then } \neg(\text{MEM } (x, x') r) \text{ else } \mathbf{T}$ )  $xs$ 

```

```

check_maximal2  $x \text{ } xs \text{ intra } =$ 
MEM  $x \text{ } xs \wedge$ 
EVERY ( $\lambda x'. \text{if } x \neq x' \text{ then } \neg(\text{check\_po\_iico } \text{intra } x \text{ } x') \text{ else } \mathbf{T}$ )  $xs$ 

```

```

(cross[_]_ = [])  $\wedge$ 
(cross (( $x, y$ )  $\in r$ )  $r' = \text{MAP } (\lambda(x', y'). (x, y')) r' ++ \text{cross } r \text{ } r'$ )

```

```

tinsert ( $x, y$ )  $r =$ 
let  $left = \text{FILTER } (\lambda(x', y'). y' = x) r$  in
let  $right = \text{FILTER } (\lambda(x', y'). x' = y) r$  in
( $x, y$ )  $\in r ++$ 
MAP ( $\lambda(x', y'). (x', y)$ )  $left ++$ 
MAP ( $\lambda(x', y'). (x, y')$ )  $right ++$ 
cross  $left \text{ } right$ 

```

```

(tclosel[_]_  $acc = acc$ )  $\wedge$ 
(tclosel (( $x, y$ )  $\in r$ )  $acc = \text{tclosel } r \text{ (tinsert } (x, y) acc)$ )

```

```

check_valid_execution  $E \text{ } X =$ 
let  $mo = \text{tclosel (FILTER } (\lambda(e_1, e_2). e_1 \neq e_2) X.ch\_memory\_order)$  in
let  $writes = \text{FILTER is\_write } E.ch\_events$  in
let  $reads = \text{FILTER is\_read } E.ch\_events$  in
let  $mwrites = \text{FILTER is\_mem\_write } writes$  in
let  $mreads = \text{FILTER is\_mem\_read } reads$  in
let  $barriers = \text{FILTER is\_barrier } E.ch\_events$  in
let  $intra = \text{tclosel } E.ch\_intra\_causality$  in
(* partial order *)
EVERY ( $\lambda(e_1, e_2). e_1 \neq e_2$ )  $mo \wedge$ 
EVERY ( $\lambda(e_1, e_2). \text{cis\_mem\_access } e_1 \wedge \text{cis\_mem\_access } e_2 \wedge$ 
  MEM  $e_1 \text{ } E.ch\_events \wedge \text{MEM } e_2 \text{ } E.ch\_events$ )  $X.ch\_memory\_order \wedge$ 
(* linear order on mwrites *)
EVERY ( $\lambda e_1. \text{EVERY } (\lambda e_2. \text{if } e_1 \neq e_2 \text{ then}$ 
  MEM ( $e_1, e_2$ )  $mo \vee \text{MEM } (e_2, e_1) mo$ 
  else
    T)
  mwrites)
  mwrites  $\wedge$ 
(* po_iico in memory_order *)
EVERY ( $\lambda er_1. \text{EVERY } (\lambda er_2. \text{check\_po\_iico\_in\_mo } \text{intra } mo \text{ } er_1 \text{ } er_2) mreads$ )

```

```

    mreads ∧
    EVERY (λer. EVERY (λew. check_po_iico_in_mo intra mo er ew) mwrites)
    mreads ∧
    EVERY (λew1. EVERY (λew2. check_po_iico_in_mo intra mo ew1 ew2) mwrites)
    mwrites ∧
    EVERY (λew.
      EVERY (λer.
        if barrier_separated intra barriers ew er ∨
          EXISTS (λes. MEM ew es ∨ MEM er es) E.ch_atomicity then
            check_po_iico_in_mo intra mo ew er
        else
          T)
      mreads)
    mwrites ∧
    (* atomicity *)
    EVERY (λes.
      EVERY (λe.
        if ¬(MEM e es) then
          EVERY (λe'.
            if is_mem_read e' ∨ is_mem_write e' then
              MEM (e, e') mo
            else
              T)
            es ∨
            EVERY (λe'.
              if is_mem_read e' ∨ is_mem_write e' then
                MEM (e', e) mo
              else
                T)
            es
          else
            T)
        (mreads ++ mwrites))
    E.ch_atomicity ∧
    (* rfmc *)
    EVERY (λer.
      case X.ch_rfmap er of
        SOME ew →
          is_read er ∧ is_write ew ∧ MEM ew E.ch_events ∧
          (loc er = loc ew) ∧ (value_of er = value_of ew)
        || NONE → T)
    E.ch_events ∧
    (* rfmap written and initial*)
    EVERY (λer.
      case X.ch_rfmap er of
        SOME ew →
          if is_mem_write ew then
            check_maximal1 ew (previous_writes1 er mo ++
              previous_writes2 er intra writes) mo
          else
            check_maximal2 ew (previous_writes2 er intra writes) intra
        || NONE →
          (case loc er of
            SOME l →
              (value_of er = X.ch_initial_state l) ∧
              (previous_writes1 er mo = []) ∧
              (previous_writes2 er intra writes = []))

```

```

    || NONE  $\rightarrow$  F))
reads

check_set_eq  $es_1\ es_2$  = EVERY ( $\lambda e.$  MEM  $e\ es_2$ ) $es_1$   $\wedge$  EVERY ( $\lambda e.$  MEM  $e\ es_1$ ) $es_2$ 

check_well_formed_event_structure  $E$  =
let  $intra$  = tclose (FILTER ( $\lambda(e_1, e_2).$   $e_1 \neq e_2$ ) $E.ch\_intra\_causality$ )[ ] in
EVERY ( $\lambda e.$  MEM (proc  $e$ ) $E.ch\_procs$ ) $E.ch\_events$   $\wedge$ 
EVERY ( $\lambda e_1.$ 
  EVERY ( $\lambda e_2.$ 
    if ( $e_1.iid = e_2.iid$ )  $\wedge$  ( $e_1.eid = e_2.eid$ ) then
       $e_1 = e_2$ 
    else T)
     $E.ch\_events$ )
   $E.ch\_events$   $\wedge$ 
EVERY ( $\lambda(e_1, e_2).$  MEM  $e_1\ E.ch\_events$   $\wedge$  MEM  $e_2\ E.ch\_events$ ) $E.ch\_intra\_causality$   $\wedge$ 
EVERY ( $\lambda(e_1, e_2).$   $e_1 \neq e_2$ ) $intra$   $\wedge$ 
EVERY ( $\lambda(e_1, e_2).$   $e_1.iid = e_2.iid$ ) $intra$   $\wedge$ 
 $\neg$  MEM[ ] $E.ch\_atomicity$   $\wedge$ 
EVERY ( $\lambda es.$  EVERY ( $\lambda e.$  MEM  $e\ E.ch\_events$ ) $es$ ) $E.ch\_atomicity$   $\wedge$ 
EVERY ( $\lambda es_1.$ 
  EVERY ( $\lambda es_2.$ 
    if  $\neg$  check_set_eq  $es_1\ es_2$  then
      EVERY ( $\lambda e_1.$  EVERY ( $\lambda e_2.$   $e_1 \neq e_2$ ) $es_1$ ) $es_2$ 
    else
      T)
     $E.ch\_atomicity$ )
   $E.ch\_atomicity$   $\wedge$ 
EVERY ( $\lambda es_1.$  EVERY ( $\lambda e_1.$  EVERY ( $\lambda e_2.$   $e_1.iid = e_2.iid$ ) $es_1$ ) $es_1$ ) $E.ch\_atomicity$   $\wedge$ 
EVERY ( $\lambda e.$  case loc  $e$  of SOME (LOCATION_REG  $p\ r$ )  $\rightarrow p = \text{proc } e$  ||  $\_ \rightarrow$  T) $E.ch\_events$   $\wedge$ 
EVERY ( $\lambda(e_1, e_2).$   $\neg$  is_mem_write  $e_1$ ) $intra$   $\wedge$ 
EVERY ( $\lambda e_1.$ 
  EVERY ( $\lambda e_2.$ 
    if is_write  $e_1$   $\wedge$   $e_1 \neq e_2$   $\wedge$  (is_write  $e_2$   $\vee$  is_read  $e_2$ )  $\wedge$ 
      ( $e_1.iid = e_2.iid$ )  $\wedge$  (loc  $e_1 =$  loc  $e_2$ ) then
        MEM ( $e_1, e_2$ ) $intra$   $\vee$  MEM ( $e_2, e_1$ ) $intra$ 
      else
        T)
       $E.ch\_events$ )
     $E.ch\_events$   $\wedge$ 
EVERY ( $\lambda es.$ 
  EVERY ( $\lambda e_1.$ 
    EVERY ( $\lambda e_2.$ 
      if  $e_1.iid = e_2.iid$  then MEM  $e_2\ es$  else T)
       $E.ch\_events$ )
     $es$ )
   $E.ch\_atomicity$   $\wedge$ 
EVERY ( $\lambda es.$  EXISTS ( $\lambda e.$  is_mem_read  $e$ ) $es$ ) $E.ch\_atomicity$ 

```

The `chE_to_E` and `chX_to_X` functions convert from the list-based representation of event structures and execution witnesses to the set-based one.

```

chE_to_E  $E$  =
 $\llbracket$   $procs := \text{set } E.ch\_procs;$ 
   $events := \text{set } E.ch\_events;$ 
   $intra\_causality :=$ 
    ( $\text{set } E.ch\_intra\_causality$ ) $^+ \cup \{(e, e) \mid e \in (\text{set } E.ch\_events)\};$ 

```

$atomicity := \text{set } (\text{MAP } \text{set } E.ch_atomicity))$

$chX_to_X \ E \ X =$
 $\llbracket \text{memory_order} := (\text{set } X.ch_memory_order)^+ \cup$
 $\{(e, e) \mid e \in \text{mem_accesses } (chE_to_E \ E)\};$
 $rfmap := \{(ew, er) \mid \text{MEM } er \ E.ch_events \wedge (X.ch_rfmap \ er = \text{SOME } ew)\};$
 $initial_state := X.ch_initial_state \rrbracket$

Theorem 28

$\forall E. \text{check_well_formed_event_structure } E = \text{well_formed_event_structure } (chE_to_E \ E)$

Proof sketch (full proof in HOL):

By expanding the definitions and using various theorems about EVERY, EXISTS and MEM from HOL's list library. □

Theorem 29

$\forall E \ X.$
 $\text{well_formed_event_structure } (chE_to_E \ E)$
 $\implies$
 $(\text{check_valid_execution } E \ X = \text{valid_execution } (chE_to_E \ E)(chX_to_X \ E \ X))$

Proof sketch (full proof in HOL):

By expanding the definitions and using various theorems about EVERY, EXISTS and MEM from HOL's list library. □

F Change History

- r1068, 2008-12-19: First version of the axiomatic model.
- r1105, 2009-01-07: Fixed the definition of valid execution for infinite executions, adding the condition that only a finite number of same-location reads can be unrelated to a write to that location. Fixed the definition of check_rfmap_written for reads from registers.
- *Revision* : 1746, 2009-03-25: Technical Report version, including the x86-TSO abstract machine memory model, the equivalence result, the verified checker, discussion of litmus tests, and proof outlines.

References

- [1] *AMD64 Architecture Programmer's Manual (3 vols)*. Advanced Micro Devices, Sept. 2007. rev. 3.14.
- [2] *Intel 64 and IA-32 Architectures Software Developer's Manual (5 vols)*. Intel Corporation, Nov. 2008. rev. 29.
- [3] S. Adve and K. Gharachorloo. Shared memory consistency models: A tutorial. *IEEE Computer*, 29(12):66–76, Dec 1996.
- [4] D. Aspinall and J. Sevcik. Formalising Java's data race free guarantee. In *Proc. TPHOLs, LNCS*, 2007.
- [5] H.-J. Boehm and S. Adve. Foundations of the C++ concurrency memory model. In *Proc. PLDI*, pages 68–78, 2008.
- [6] G. Boudol and G. Petri. Relaxed memory models: an operational approach. In *Proc. POPL*, pages 392–403, 2009.

- [7] S. Burckhardt and M. Musuvathi. Effective program verification for relaxed memory models. Technical Report MSR-TR-2008-12, Microsoft Research, 2008. Conference version in Proc. CAV 2008, LNCS 5123.
- [8] S. Burckhardt, M. Musuvathi, and V. Singh. Verifying compiler transformations for concurrent programs, Jan. 2009. Technical report MSR-TR-2008-171.
- [9] D. Dice. Java memory model concerns on Intel and AMD systems. http://blogs.sun.com/dave/entry/java_memory_model_concerns_on, Jan. 2008.
- [10] S. Hangal, D. Vahia, C. Manovit, J.-Y. J. Lu, and S. Narayanan. TSOtool: A program for verifying memory systems using the memory consistency model. In *Proc. ISCA*, pages 114–123, 2004.
- [11] The HOL 4 system. <http://hol.sourceforge.net/>.
- [12] Intel. Intel 64 architecture memory ordering white paper, 2007. SKU 318147-001.
- [13] L. Higham, J. Kawash, and N. Verwaal. Defining and comparing memory consistency models. In *PDCS*, 1997. Full version as TR #98/612/03, U. Calgary.
- [14] P. Loewenstein. Personal communication, Nov. 2008.
- [15] P. N. Loewenstein, S. Chaudhry, R. Cypher, and C. Manovit. Multiprocessor memory model verification. In *Proc. AFM (Automated Formal Methods)*, Aug. 2006. FLoC workshop. <http://fm.csl.sri.com/AFM06/>.
- [16] V. M. Luchangco. *Memory consistency models for high-performance distributed computing*. PhD thesis, MIT, 2001.
- [17] A. Roy, S. Zeisset, C. J. Fleckenstein, and J. C. Huang. Fast and generalized polynomial time memory consistency verification. In *CAV*, pages 503–516, 2006.
- [18] A. Roychoudhury. Formal reasoning about hardware and software memory models. In *Proc. ICFEM*, pages 423–434, 2002.
- [19] V. Saraswat, R. Jagadeesan, M. Michael, and C. von Praun. A theory of memory models. In *Proc. PPoPP*, 2007.
- [20] S. Sarkar, P. Sewell, F. Zappa Nardelli, S. Owens, T. Ridge, T. Braibant, M. Myreen, and J. Alglave. The semantics of x86-CC multiprocessor machine code. In *Proc. POPL 2009*, Jan. 2009.
- [21] P. S. Sindhu, J.-M. Frailong, and M. Cekleov. Formal specification of memory models. In *Scalable Shared Memory Multiprocessors*, pages 25–42. Kluwer, 1991.
- [22] S. Owens, S. Sarkar, and P. Sewell. The x86-TSO model, 2009. www.cl.cam.ac.uk/users/pes20/weakmemory/.
- [23] SPARC International, Inc. The SPARC architecture manual, v. 8. Revision SAV080SI9308. <http://www.sparc.org/standards/V8.pdf>, 1992.
- [24] Y. Yang, G. Gopalakrishnan, G. Lindstrom, and K. Slind. Nemos: A framework for axiomatic and executable specifications of memory consistency models. In *IPDPS*, 2004.