Revamping Hardware Persistency Models

View-Based and Axiomatic Persistency Models for Intel-x86 and Armv8

Kyeongmin Cho
kyeongmin.cho@kaist.ac.kr
KAIST
Daejeon, Korea

Sung-Hwan Lee
sunghwan.lee@sf.snu.ac.kr
Seoul National University
Seoul, Korea

Azalea Raad
azalea@imperial.ac.uk
Imperial College London
London, United Kingdom

Jeehoon Kang
jeehoon.kang@kaist.ac.kr
KAIST
Daejeon, Korea

Abstract

Non-volatile memory (NVM) is a cutting-edge storage technology that promises the performance of DRAM with the durability of SSD. Recent work has proposed several *persistency models* for mainstream architectures such as Intel-x86 and Armv8, describing the order in which writes are propagated to NVM. However, these models have several limitations; most notably, they either lack operational models or do not support persistent synchronization patterns.

We close this gap by revamping the existing persistency models. First, inspired by the recent work on promising semantics, we propose a *unified operational style* for describing persistency using *views*, and develop view-based operational persistency models for Intel-x86 and Armv8, thus presenting the *first* operational model for Armv8 persistency. Next, we propose a *unified axiomatic style* for describing hardware persistency, allowing us to recast and repair the existing axiomatic models of Intel-x86 and Armv8 persistency. We prove that our axiomatic models are equivalent to the authoritative semantics reviewed by Intel and Arm engineers. We further prove that each axiomatic hardware persistency model is equivalent to its operational counterpart. Finally, we develop a persistent model checking algorithm and tool, and use it to verify several representative examples.

CCS Concepts: • **Theory of computation** → **Concurrency**; **Semantics and reasoning**; *Verification by model checking*.

Keywords: persistent memory, non-volatile random-access memory, NVRAM, persistency semantics, x86, Armv8

1 Introduction

Non-volatile memory (NVM) is an emerging class of storage technology that simultaneously provides (1) byte addressability, low latency, and high throughput as DRAM does; and (2) durability (data persistency across system crashes) and high capacity as SSD does. It is widely believed that NVM (a.k.a. persistent memory) will eventually supplant volatile memory [42], allowing efficient access to persistent data. This belief is backed by industrial support. Specifically, the two major architectures, Intel-x86 and Armv8 which together account for almost 100% of the desktop and mobile market, have extended their official specifications to support persistent programming [4, 20]. Intel has further released open-source NVM libraries such as PMDK [19], and manufactured its own line of NVM, Optane DC persistent memory [21], with an extended academic study evaluating its performance [24]. NVM is therefore expected to innovate high performance transactional systems [7, 18, 31, 35, 38, 53] and large-scale memory systems [37, 40, 50].

However, building correct transactional systems over persistent memory is difficult in part due to *relaxed* persistency: writes to NVM locations may not be persisted to memory in the program order due to micro-architectural optimizations such as out-of-order execution, store buffering, or caching protocols. For instance, consider the programs below:

(a) data := 42	
(b) commit := 1	(COMMITWEAK)

(a) data := 42	
(b) flush data	(COMMIT1)
(c) commit := 1	

Hereafter we assume all program variables in our examples are locations in NVM[1] initialized to 0; variable reads and writes are architecture-level load and store instructions, e.g., `mov` on Intel-x86 and `ldr`, `str` in Armv8; and that `flush` represents a persistency fence, e.g., `clflush` on Intel-x86 and `dc.cvap; dsb.sy` on Armv8.[2]

In both examples we aim to establish the invariant $I \stackrel{\triangle}{=}$ commit=1 $\Rightarrow$ data=42 even in case of an unexpected crash. In the case of CommitWeak without a persistency fence, we fail to establish I over mainstream architectures such as Intel-x86 and Armv8: the two stores may persist to NVM out of order, thereby allowing commit=1, data=0 upon crash recovery. By contrast, in the case of Commit1 the persistency fence at b ensures that the two stores persist in the intended (program) order, thereby establishing the invariant I. Microarchitecturally, `flush` data blocks until the previous store on data at a is persisted to NVM, thus ensuring that the store at c always persists after that of a. As such, persistency fences are expensive and should be used sparingly.

Relaxed persistency is further complicated in multi-threaded settings. Consider the following program with two threads:

$$
\begin{array}{l|l}
(a)\ \text{data := 42} & (b)\ \text{if (data != 0) \{} \\
 & (c)\ \quad\ \text{flush data} \qquad (\textsc{Commit2}) \\
 & (d)\ \quad\ \text{commit := 1 \}}
\end{array}
$$

This example differs from Commit1 in that data and commit are written to by different threads. Once again, if the fence at c were removed, the desired invariant I would no longer hold: although the store on data at a may be propagated (made visible) to the the right thread through cache coherence protocols, it may not be persisted to NVM prior to the crash. As before, the fence at c ensures that the store at a (which was propagated to the right thread before c) persists to NVM before the store at d, thus establishing I.

Note that during normal (non-crashing) executions, under both Intel-x86 and Armv8 no thread can observe the undesirable behavior commit=1, data=0 even without the fence at c, underlining the difference between the *consistency order* (the order in which stores are propagated across threads) and the *persistency order* (the order in which stores are persisted to NVM). In general, relaxed concurrency models constrain the consistency order, while relaxed persistency models additionally constrain the persistency order, further compounding the complexity of relaxed concurrency.

In order to facilitate correct persistent programming with efficient use of persistency fences, existing work includes several persistency models [8, 17, 28, 30, 42, 46–48]. However, as we discuss below, these models have several shortcomings.

---

[1] As in [47], we assume all locations are durable locations in NVM.

[2] Armv8 recently introduced the `dc.cvadp` instruction that, unlike `dc.cvap`, guarantees persistence even in case of battery/hardware failures [4]. We focus on `dc.cvap` in this paper, but most discussions also apply to `dc.cvadp`.

Problem To our knowledge, no existing persistency model (except for PTSO_{syn} [28], discussed shortly below) satisfies all of the following properties simultaneously:

(A) **Describing mainstream architectures or languages**: For a persistency model to be widely used and applied, it should describe the persistency behavior of mainstream hardware/software platforms such as readily available architectures, e.g., Intel-x86 [20] and Armv8 [4], and ubiquitous languages, e.g., C/C++, over which several persistent libraries are implemented [18, 19]. Moreover, the model should be sufficiently relaxed that the behaviors observable on existing platforms are also allowed by the model. Otherwise, invariants that hold according to the model would be invalidated by executions on such platforms, rendering the model unsound for reasoning.

(B) **Supporting persistent synchronization patterns**: A persistency model should support common synchronization patterns used in practical implementations of persistent objects, e.g., transactions or file systems. For instance, a model should prohibit undesirable behaviors, e.g., commit=1, data=0 in Commit1 and Commit2 that capture the essence of practical implementations of transactional systems. In particular, the model should be sufficiently strict that unobservable behaviors on existing platforms are also forbidden by the model. Otherwise, admitting unobservable behaviors in the model makes it impossible to reason about such patterns. Moreover, a model should serve as an objective correctness criteria for new, more efficient designs of persistent objects, and in doing so, guide such new designs.

(C) **Operational**: An *operational* persistency model is desirable in that it enables stepping through an execution for debugging purposes. Moreover, operational models are more suitable for building high-level reasoning techniques such as program logics. By contrast, axiomatic models constrain the admitted behaviors through a set of axioms over full executions, making them undesirable for step-by-step reasoning (e.g., as in program logics).

The models of [8, 17, 30, 42, 46] do not satisfy (A). Specifically, [8, 17, 30, 42] present language-level persistency models put forward as academic *proposals*, and are not supported by mainstream programming languages. Similarly, [46] proposes a hardware persistency model, PTSO, by integrating buffered epoch persistency [42] with the TSO architecture of x86/SPARC [49]. However, PTSO is not supported by mainstream architectures of Intel-x86 and Armv8.

The PArmv8 model [48, "PARMv8"] describes the persistency semantics of the Armv8 architecture, but is not operational (C). Moreover, as we discuss shortly in §2, the PArmv8 model is too weak in that it violates multi-copy atomicity. Similarly, the Px86 model [47] describes the persistency semantics of the Intel-x86 architecture operationally

$$x86_{axiom}\ [3] \xleftrightarrow[\ (\S5)\ ]{equivalent} \mathbf{x86_{view}}\ (\S3)$$

adding persistency (x86) — *adding persistency* (x86view)

$$Px86\ [47] \xrightarrow{fixed} \mathbf{Px86_{axiom}}\ (\S4) \xleftrightarrow[\ (\S5)\ ]{equivalent} \mathbf{Px86_{view}}\ (\S3.5)$$

$$Armv8_{axiom}\ [44] \xleftrightarrow[\ [44]\ ]{equivalent} Armv8_{view}\ [44]$$

adding persistency — *adding persistency*

$$PArmv8\ [48] \xrightarrow{fixed} \mathbf{PArmv8_{axiom}}\ (\S6.3) \xleftrightarrow[\ (\S6.4)\ ]{equivalent} \mathbf{PArmv8_{view}}\ (\S6.2)$$
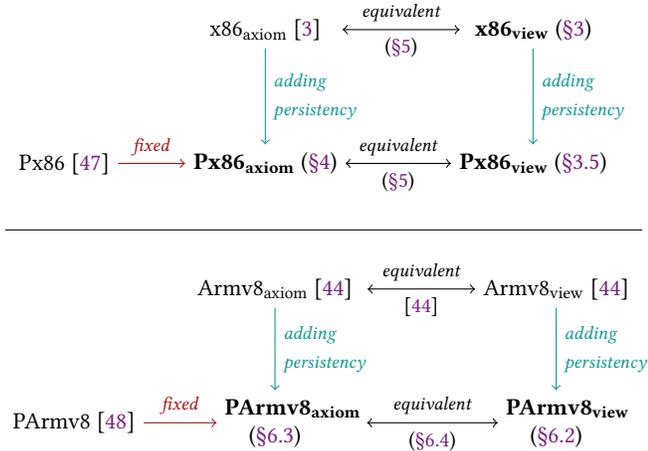
Figure 1. Relationship among Intel-x86 and Armv8 models

and axiomatically.[3] However, as we discuss shortly, Px86 is too relaxed and does not always support persistent synchronization patterns in the presence of I/O (B).

Khyzha and Lahav [28] recently developed the $PTSO_{syn}$ model for Intel-x86 that fixes the Px86 problem regarding I/O and satisfies (A)–(C). However, they do not discuss this problem as they have a different motivation, i.e., presenting a model that better matches the developers' intuition [28, §1]. We discuss $PTSO_{syn}$ in more detail later in §8.

Our Solution, Contributions and Outline We propose a *unified operational style* for describing relaxed persistency using *views*, and develop view-based persistency models of Intel-x86/Armv8 that satisfy all three (A)–(C) properties. In doing so, we develop the *first* operational model for Armv8 persistency. Our operational models highlight 2 flaws in the existing (axiomatic) persistency models of Intel-x86 and Armv8. To remedy this, we develop a *unified axiomatic style* for persistency, adapt the the existing Intel-x86/Armv8 persistency models to our unified style, and repair their flaws.

The remainder of this paper is organized as follows:

- We discuss the shortcomings of the existing persistency models of Intel-x86/Armv8 and present an intuitive account of our solution as view-based models (§2).
- We develop $x86_{view}$, a new view-based model for Intel-x86 concurrency (§3).
- We develop $Px86_{view}$ (§3.5) and $PArmv8_{view}$ (§6.2), respectively extending the $x86_{view}$ and $Armv8_{view}$ [44] models to account for persistency.
- We present $Px86_{axiom}$ (§4) and $PArmv8_{axiom}$ (§6.3), our axiomatic models of Intel-x86 and Armv8 persistency

that simplify and repair the state-of-the-art models of the respective architectures [47, 48]. We prove that our axiomatic models are equivalent to the authoritative semantics reviewed by Intel and Arm engineers, modulo our proposed fixes (§4.4 and §6.3). Our proposed fix in $PArmv8_{axiom}$ has been reviewed by Arm engineers.

- We prove that $Px86_{view}$ and $PArmv8_{view}$ are equivalent to $Px86_{axiom}$ and $PArmv8_{axiom}$, respectively. The equivalence proof is mechanized in Coq (§5 and §6.4).
- We develop a stateless model checker for persistency and use it to verify several representative examples under $PArmv8_{view}$ (§7). We conclude with related and future work (§8).

We present an overview of the concurrency and persistency models we present in this paper in Fig. 1, summarizing their relationship with existing models in the literature.

2 Overview

We discuss the shortcomings of Px86 and PArmv8 as regards to (B) (§2.1 and 2.2). We then present an intuitive account of our key idea to provide a persistency model that satisfies all three desired properties in (A)–(C) simultaneously (§2.3).

2.1 The Px86 Model and Synchronous Flushes

The Px86 model [47] is too weak in that its instruction for propagating stores to NVM behaves *asynchronously*: executing `clflush` under Px86 does not block execution, and merely guarantees that the pending stores on the given location will be persisted to NVM at some future point. For instance, if the generic `flush` data instruction in Commit1 is replaced with its Intel-x86 analogue, `clflush` data, then once `clflush` data is executed, there is no guarantee under Px86 that the earlier stores on data (including that at *a*) are persisted to NVM; rather (1) these stores will be persisted to NVM at some future point; and (2) they will be persisted to NVM *before all future stores* (including that at *c*). In other words, the persistency ordering guarantees of `clflush` in (2) allows us to establish the desired invariant *I*, even though the effect of `clflush` data may not immediately take place.

The asynchronous behavior of `clflush` is observable in the presence of *external operations* as they narrow down possible crash points through additional observations. For instance, consider the variant of Commit1 below where we replace the store to commit with an analogous I/O operation that writes "commit" to `file` on disk:

```
(a) data := 42    (d) if (flag != 0) {
(b) flush data    (e)   log(file, "commit")}    (CommitE)
(c) flag := 1
```

Let us write *C* to denote that `file` contains "commit". Under Px86 it is possible to observe the post-crash state $S : \text{data}=0 \wedge C$; i.e., when the I/O operation is executed, the asynchronous effect of `clflush` may not have taken place yet.

---

[3]In [47] the authors introduce two persistency models for Intel-x86: $Px86_{man}$ which formalizes the ambiguous and under-specified behavior described in the Intel reference manual [20], and $Px86_{sim}$ which simplifies and strengthens $Px86_{man}$ to capture the architectural intent envisaged by Intel engineers. In this paper we focus on the $Px86_{sim}$ model and simply refer to it as Px86.

Figure 2. A view-based execution of Commit1

As such, to support persistency synchronization patterns such as CommitE in the presence of external operations under Px86, we must strengthen Px86 by modeling the behavior of clflush *synchronously*. Let us write SPx86 for a strengthening of Px86 in which clflush instructions are executed synchronously, i.e., they block until all pending stores on the location are persisted to NVM. In the absence of external operations such as I/O or network messages, the asynchronous behavior of clflush cannot be observed, i.e., SPx86 is indistinguishable from Px86. By contrast, in the presence of external operations, only SPx86 satisfies an invariant analogous to that of Commit1: $C \Rightarrow$ data=42; i.e., once the I/O operation is executed, the synchronous effect of clflush must have taken place and S cannot be observed.

2.2 The PArmv8 Model and Multi-Copy Atomicity

The PArmv8 model [48] is too weak in that it violates the principles of *multi-copy atomicity* (MCA) which ensures that a write by one thread is made visible to all other threads simultaneously. Although Armv8 was originally non-MCA, it was recently simplified to observe MCA [43]. However, the persistency extension of Armv8 in PArmv8 violates MCA by allowing the following behavior regarding persistency:

$$
\begin{array}{ll}
(a)\ \mathsf{y} := 1 & (f)\ \mathsf{x} := 1 \\
(b)\ \mathsf{dsb.sy} & (g)\ \mathsf{dsb.sy} \\
(c)\ \mathsf{flushopt}\ \mathsf{x} & (h)\ \mathsf{flushopt}\ \mathsf{y} \qquad (\textsc{FlushMCA}) \\
(d)\ \mathsf{dsb.sy} & (i)\ \mathsf{dsb.sy} \\
(e)\ \mathsf{z} := 1 & (j)\ \mathsf{w} := 1
\end{array}
$$

Executing flushopt x persists all pending stores on the same cache line as x *asynchronously*.[4] Moreover, if flushopt x is followed by a data synchronization barrier, dsb.sy, its effects take place *synchronously*; i.e., executing dsb.sy awaits the completion of all earlier flushopt by the same thread.

We argue that MCA should preclude the post-crash state $S : \mathsf{z} = \mathsf{w} = 1 \ \wedge \ \mathsf{x} = \mathsf{y} = 0$. First, to observe $\mathsf{z} = \mathsf{w} = 1$ after a crash, the two threads should have fully executed to the end. Second, to observe $\mathsf{y} = 0$ after a crash, (a) should not have been made visible to (h) prior to the crash, and thus (h) must be ordered before (a). Third, (g) must be ordered before (h) and (a) before (b) because (g) and (b) are fences. Transitively, (g) must be ordered before (h), (a), and then

(b). As such, (f) should be visible to (c), thus ensuring $\mathsf{x} = 1$ after the crash and precluding the behavior in S.[5]

To ensure MCA for persistency, we must thus strengthen PArmv8 by enforcing an order between a flush (e.g., c) and a write on the same location that is not persisted by the flush (e.g., f). Let us write SPArmv8 for such a strengthening of PArmv8. Under SPArmv8, if $\mathsf{x} = 0$ after a crash, then (c) is ordered before (f); (a) is ordered before (h); $\mathsf{y} = 1$ is persisted to the NVM; and thus S cannot be observed.

Upon discussing FlushMCA with engineers at Arm, they confirmed that this non-MCA behavior is indeed prohibited and our proposal in §6.3 is the correct interpretation of Arm architecture reference manual [4].

2.3 Our Solution: View-Based Operational Models

We present view-based operational models for the relaxed persistency behavior of Intel-x86/Armv8 architectures that satisfy all three properties in (A)–(C). We build our model over the view-based model of Armv8/RISC-V relaxed-memory *concurrency* [44]. Intuitively, view-based models [27, 34, 44] combine two key ideas: (1) recording the entire *store history* in the memory and allowing threads to read old values; and (2) imposing ordering constraints with per-thread *views* representing the set of stores propagated to each thread and thus constraining the outcomes of future loads and stores by a thread. Here, we further introduce the notion of *persistency views* for each location l, denoting the set of stores on l that have persisted to NVM and thus will survive a crash.

We next illustrate these ideas through a view-based execution of Commit1 in Fig. 2, comprising a single thread *tid*. At each execution stage, the store history is recorded in memory as an indexed (timestamped, e.g., @1) list of stores; the view of *tid* records (the timestamp of) the latest store propagated to *tid* (the *tid*-labelled arrow); and the persistency view of each location l records (the timestamp of) the latest store on l that has persisted to NVM (the NVM[l]-labelled arrows).

The initial memory is $M = []$, denoting the empty history (no stores have executed), depicted as init at timestamp 0 (@0); the *tid* view is $v = @0$ (no stores have propagated to *tid*); and the persistency view of each location l is $v_{\mathrm{NVM}}[l] = @0$ (no stores on l have persisted to NVM). Subsequently:

---

[4]For the sake of uniformity with our respective Intel-x86 models, we write flushopt x in lieu of the Armv8 instruction dc cvap x.

[5]The reader may have noted that this behavior is forbidden even if the dsb.sy at (b) and (g) are replaced with the weaker dmb.sy. We opt for dsb.sy to simplify the example by using only one kind of fence.

(a) Executing data := 42 appends its store to memory ($M =$ [⟨data := 42⟩@1]), and advances the *tid* view ($v = @1$): the store is executed by and thus propagated to *tid*.

(b) Executing flush data joins the persistency view of data with the *tid* view ($v_{\text{NVM}}[\text{data}] = @1$), ensuring that the latest data store propagated to *tid* is persisted to NVM.

(c) Analogously, executing commit := 1 yields $v = @2$ and $M = [⟨\text{data} := 42⟩@1, ⟨\text{commit} := 1⟩@2]$.

The post-crash outcomes (NVM contents) are then determined by the persistency views. Concretely, after a crash each NVM location l may contain a value written by a store whose timestamp is at least $v_{\text{NVM}}[l]$. For instance, if a crash occurs after executing flush data, then in the post-crash state $v_{\text{NVM}}[\text{data}] = @1$ and thus data = 42@1; i.e., data := 42 must have persisted to NVM, establishing invariant I.

We next describe an execution of Commit2, where tid_1 and tid_2 denote the left and right threads, respectively. Initially, the memory is $M = []$; the persistency view is $v_{\text{NVM}} = \lambda l.@0$; and the tid_i view is $v_i = @0$ (for $i \in \{1, 2\}$). Then:

(a) Executing data := 42 yields $M=[⟨\text{data} := 42⟩@1]$, $v_1=@1$.

(b) Thread tid_2 may then load data = 42 as its view timestamp ($v_2 = @0$) is less than @1 of data := 42. After loading data = 42, the tid_2 view is joined with @1: $v_2 = @1$.

(c) Executing flush data yields $v_{\text{NVM}}[\text{data}] = @1$.

(d) Executing commit := 1 results in $M = [⟨\text{data} := 42⟩@1, ⟨\text{commit} := 1⟩@2]$ and $v_2 = @2$.

As with Commit1, the invariant I holds in case of a crash.

Our models indeed satisfy all desired properties. (A) Our models capture the persistency behavior of the mainstream Armv8 and Intel-x86 architectures. Specifically, we prove that our models are equivalent (modulo fixes) to the axiomatic models of [47, 48] reviewed by Intel/Arm engineers. Our equivalence proof is mechanized in Coq [1] and is publicly available [10]. (B) Our models support persistent synchronization patterns such as those of Commit1 and Commit2. (C) Our models are operational as with the existing family of view-based models [27, 44]. Furthermore, to support reasoning about programs over our models, we develop a stateless model checking algorithm and tool for persistency verification, and use it to verify several representative examples under PArmv8$_{\text{view}}$.[6] Our model checking tool and verified examples are open-source and publicly available [10].

---

[6]As a proof of concept, we focus on model checking only Armv8 persistency. This is sufficient to showcase the feasibility of model checking for hardware persistency since Armv8 is more complex than Intel-x86 with a bigger search space. We believe it is straightforward to adapt our approach to Intel-x86 persistency, especially given our unified semantic style for persistency.

$$p \ ::= s_1 \ || \ \ldots \ || \ s_n$$

$$s \in \text{St} \ ::= \text{skip} \ | \ s_1; s_2 \ | \ \text{if} \ (e) \ s_1 \ s_2 \ | \ \text{while} \ (e) \ s \ | \ r := e$$
$$| \ r := \text{load} \ [e] \ | \ \text{store} \ [e_1] \ e_2 \ | \ r := \text{rmw} \ rop \ [e]$$
$$| \ \text{fence}_f \ | \ \text{flush} \ e \ | \ \text{flushopt} \ e$$

$$rop \in \text{Rmw} \ ::= \text{fetch-op} \ op \ e \ | \ \text{cas} \ e_1 \ e_2$$

$$f \in \text{F} \ ::= \text{sfence} \ | \ \text{mfence}$$

$$e \in \text{Expr} \ ::= v \ | \ r \ | \ (e_1 \ op \ e_2) \qquad op \in \text{O} \ ::= + \ | - | \ \ldots$$

$$v \in \text{Val} \ = \mathbb{Z} \qquad r \in \text{Reg} \ = \mathbb{N} \qquad l \in \text{Loc} \ = \text{Val}$$

Figure 3. Intel-x86 concurrency and persistency language

3 Px86$_{\text{view}}$: A View-Based Model for Intel-x86 Persistency

We develop Px86$_{\text{view}}$, a view-based model for Intel-x86 persistency. We present a simple language for Intel-x86 concurrency and persistency (§3.1) used throughout this section. We develop x86$_{\text{view}}$, a new Intel-x86 concurrency model we use as a baseline (without persistency) and its two key ideas: store histories (§3.2) and views (§3.3); we describe how we support read-modify-writes (§3.4). We then extend x86$_{\text{view}}$ with persistency and develop Px86$_{\text{view}}$ (§3.5).

3.1 Language for Intel-x86 Persistency

To keep our presentation concrete, we use the language in Fig. 3 for Intel-x86 concurrency and persistency. A program p consists of concurrent statements run by distinct threads. A statement s is given by the standard 'while' language over register machines with concurrent memory instructions. The instruction $r := \text{load} \ [e]$ reads from the (NVM) location denoted by expression e and returns it in register r. The store $[e_1] \ e_2$ reads the value denoted by e_2 and stores it at the location denoted by e_1. Analogously, $r := \text{rmw} \ rop \ [e]$ evaluates the expressions rop, performs an RMW ('read-modify-write') operation (e.g., 'compare-and-swap') on the location denoted by e, and returns its old value in r. Finally, fence_f issues a memory 'fence' such as sfence or mfence; and flush e and flushopt e persist to NVM the pending stores on the cache line containing the location given by e.

3.2 The x86$_{\text{view}}$ Model

We present x86$_{\text{view}}$, our view-based operational model for Intel-x86 concurrency, in Fig. 4. Our design of x86$_{\text{view}}$ is inspired by Armv8$_{\text{view}}$, a view-based concurrency model of Armv8 [44]. As Intel-x86 concurrency is simpler than that of Armv8, we develop x86$_{\text{view}}$ by removing certain Armv8$_{\text{view}}$ features. Here, we highlight the interesting aspects of Intel-x86 and refer the reader to [9, Appendix A] for the full details.

States We represent a machine as a pair $\langle \vec{T}, M \rangle$, comprising a thread map $\vec{T}$ and a memory M. A thread map associates each thread with a statement and a thread state. A thread

$$\langle \vec{T}, M \rangle \in \text{Machine} \triangleq (\text{TId} \to \text{Thread}) \times \text{Memory}$$

$$tid \in \text{TId} \triangleq \mathbb{N} \quad T \in \text{Thread} \triangleq \text{St} \times \text{TState}$$

$$M \in \text{Memory} \triangleq \text{list Msg} \quad w \in \text{Msg} \triangleq \langle \text{loc} : \text{Loc}; \text{val} : \text{Val}; \text{tid} : \text{TId} \rangle$$

$$\langle l := v \rangle_{tid} \triangleq \langle \text{loc} = l; \text{val} = v; \text{tid} = tid \rangle \quad t \in \mathbb{T} \triangleq \mathbb{N} \quad \nu \in \mathbb{V} \triangleq \mathbb{T}$$

$$ts \in \text{TState} \triangleq \left\langle \begin{array}{l} \text{regs} : \text{Reg} \to \text{Val}; \\ \text{coh} : \text{Loc} \to \mathbb{V}; \quad \text{v}_{\text{rNew}} : \mathbb{V}; \end{array} \right\rangle$$

(INIT)

$$\frac{p = s_1 \mid\mid \dots \mid\mid s_n}{\text{init}(p, \langle \lambda tid. \langle s_{tid}, \langle \text{regs} = \lambda\_. \ 0; \text{coh} = \lambda\_. \ @0; \text{v}_{\text{rNew}} = @0 \rangle \rangle, [] \rangle)}$$

(MACHINE)

$$\frac{\vec{T}[tid], M \to_{tid} T', M'}{\langle \vec{T}, M \rangle \to \langle \vec{T}[tid \mapsto T'], M' \rangle}$$

(NOT-OVERWRITTEN)

$$\frac{\forall t \in (v_2, v_1]. \ M[t].\text{loc} \neq l}{v_1 \sqsubseteq_{M,l} v_2}$$

(STORE)

$$\frac{\begin{array}{c} l = [\![e_1]\!]_{ts.\text{regs}} \quad v = [\![e_2]\!]_{ts.\text{regs}} \quad ts.\text{coh}[l] \sqsubseteq t \quad \sqcup_l ts.\text{coh}[l] \sqsubseteq t \\ t = |M| + 1 \quad M' = M + [\langle l := v \rangle_{tid}@t] \quad ts' = ts[\text{coh}[l] \mapsto t] \end{array}}{(\text{store } [e_1] \ e_2, ts), M \to_{tid} (\text{skip}, ts'), M'}$$

(LOAD)

$$\frac{\begin{array}{c} l = [\![e]\!]_{ts.\text{regs}} \quad M[t] = \langle l := v \rangle \quad ts.\text{coh}[l] \sqsubseteq t \quad ts.\text{v}_{\text{rNew}} \sqsubseteq_{M,l} t \\ ts' = ts[\text{regs}[r] \mapsto v, \text{coh}[l] \mapsto t, \text{v}_{\text{rNew}} \mapsto_{\sqcup} t \neq ts.\text{coh}[l] \ ? \ t] \end{array}}{(r := \text{load } [e], ts), M \to_{tid} (\text{skip}, ts'), M}$$

(MFENCE)

$$\frac{ts' = ts[\text{v}_{\text{rNew}} \mapsto_{\sqcup} \sqcup_l ts.\text{coh}[l]]}{(\text{mfence}, ts), M \to_{tid} (\text{skip}, ts'), M}$$

Figure 4. States and transitions of $x86_{\text{view}}$ (excerpt)

state $ts \in \text{TState}$ consists of a register map, regs, assigning values to registers, and per-thread 'views' (described in §3.3). A memory is a list of messages; a message is a triple $\langle l := v \rangle_{tid}$ comprising a memory location (l), a value stored (v), and the id (tid) of the thread storing it. We write $\langle l := v \rangle_{tid}@t$ to denote that $\langle l := v \rangle_{tid}$ is issued at timestamp (index) t, starting from index @1. For simplicity, we assume a memory contains the initial message $\langle l := 0 \rangle@0$ for each l.

Transitions of $x86_{\text{view}}$ In the initial state for a program p **(INIT)**, thread statements are those in p; the register maps are $\lambda\_.0$; the views are @0; and the memory is empty ([]).

The transitions for control flow and assignment are standard (omitted). The **(MACHINE)** transition of $x86_{\text{view}}$ models thread interleaving as in sequential consistency (SC) [32].

Nevertheless, $x86_{\text{view}}$ allows relaxed (weaker than SC) behaviors since it records the entire history of stores in its memory as a list of messages, and allows threads to read stale values. Ignoring the colored premises (described later), when executing a store **(STORE)**, a thread determines the location l and the value v, and appends a new message $\langle l := v \rangle$ to the memory. Analogously, when executing $r := \text{load } [e]$

Table 1. Informal description of concurrency views

View	Past	Future
$\text{coh}[l]$	Upper bound of past reads and writes on l	Lower bound of future reads and writes on l
v_{rNew}	Upper bound of past updates; upper bound of external reads (from other threads)	Lower bound of future reads

(LOAD), a thread determines the location l, chooses a message $\langle l := v \rangle@t$ from the memory, and assigns v to r in the register map. Crucially, the chosen message need not be the latest one, thus allowing a stale value to be read. However, the chosen message should not have been overwritten **(NOT-OVERWRITTEN)** from the thread's point of view. We describe the remaining transitions shortly.

Store Buffering Recording stores as messages allows store buffering, a representative relaxed behavior of Intel-x86:

$$\begin{array}{ll} (a) \ \text{x} := 1 & (c) \ \text{y} := 1 \\ (b) \ r_1 := \text{y} \ //= 0 & (d) \ r_2 := \text{x} \ //= 0 \end{array} \quad \text{(SB)}$$

While the relaxed outcome $r_1 = r_2 = 0$ is prohibited under SC, it is allowed under Intel-x86 and may arise in $x86_{\text{view}}$ by: (a) writing $\langle \text{x} := 1 \rangle_{tid_1}@1$; (b) reading $\langle \text{y} := 0 \rangle@0$; (c) writing $\langle \text{y} := 1 \rangle_{tid_2}@2$; and most importantly, (d) reading the old value $\langle \text{x} := 0 \rangle@0$ that is overwritten by (a).

3.3 Concurrency Views

The model described thus far is too weak in that it allows behaviors prohibited under Intel-x86. We next describe how we strengthen the model to forbid such behaviors through *views*, as summarized in Table 1.

Coherence Intel-x86 orders loads and stores on the same location in a single thread as illustrated below:

$$\begin{array}{llll} (a) \ \text{x} := 20 & //= @2 & (d) \ r_2 := \text{y} & //= @4 \\ (b) \ \text{x} := 10 & //\neq @1 & (e) \ \text{y} := 30 & //\neq @3 \quad \cdots \quad \text{(COH)} \\ (c) \ r_1 := \text{x} & //\neq @1 & (f) \ r_3 := \text{y} & //\neq @3 \end{array}$$

The first thread issues $\langle \text{x} := 20 \rangle@2$, and then writes to x again and reads from it. Coherence orders (a) before (b) and (c) since they access the same location, thus forbidding them from accessing earlier timestamps, e.g., @1. Similarly, the second thread reads the message $\langle \text{y} := 40 \rangle@4$, and then writes to y and reads from it again. Coherence orders (d) before (e) and (f) as they access the same location, thus forbidding them from accessing earlier timestamps, e.g., @3.[7]

---

[7]Indeed, coherence between an access and a write is already enforced through **(STORE)** : stores always append messages to the end of memory. Nevertheless, we explicitly order them with views to achieve (i) uniformity with other coherence orders and (ii) correspondence with $\text{Armv8}_{\text{view}}$, where stores may add messages in places other than the end of memory.

To enforce coherence, we introduce *coherence views* that record past thread behaviors and simultaneously constrain future thread behaviors. Specifically, for each location l, a thread state ts records a coherence view in $ts.\text{coh}[l]$ as a timestamps (initialised to @0) representing an index in memory. The $ts.\text{coh}[l]$ represents the maximum (latest) timestamp observed for l by the thread; moreover, it forbids the thread from accessing messages of l with earlier timestamps than $ts.\text{coh}[l]$. Put formally, in (LOAD) and (STORE) we additionally require $ts.\text{coh}[l] \sqsubseteq t$ in the premise and update $ts'.\text{coh}[l] \mapsto t$ in the conclusion.

These changes indeed forbid the undesirable behavior in Coh: (a) updates $ts.\text{coh}[\text{x}]$ to @2, forbidding (b) and (c) from accessing @1. Similarly, (d) updates $ts.\text{coh}[\text{y}]$ to @4, forbidding (e) and (f) from accessing @3.

Message Passing In addition to coherence, Intel-x86 orders certain accesses on different locations via 'message passing':

$$
\begin{array}{ll}
(a) \text{ data} := 42 & (c) \ r_1 := \text{flag} \quad //= 1 \\
(b) \text{ flag} := 1 & (d) \ r_2 := \text{data} \quad //\neq 0
\end{array} \quad \text{(MP)}
$$

If the right thread reads 1 is from flag, then it should read 42 from data as (a) is ordered before (d) as follows:

- (a) **before** (b): A load or store is ordered before later stores. To enforce this, in (STORE) we additionally require $\bigsqcup_l ts.\text{coh}[l] \sqsubseteq t$ in the premise.[8]
- (b) **before** (c): A store is ordered before loads that read from it ("message passing"). This is already enforced as the store message read by the load is issued before it.
- (c) **before** (d): A load is ordered before a later load. To enforce this, we introduce the *new-read view*. Specifically, a thread state ts includes a 'new-read' view, $ts.\text{v}_{\text{rNew}}$, recording the maximum (latest) view previously read by the thread. Moreover, it forbids the thread's future loads (on any location) from reading messages that are overwritten by $ts.\text{v}_{\text{rNew}}$. Put formally, in (LOAD) we require $ts.\text{v}_{\text{rNew}} \sqsubseteq_{M,l} t$ (i.e., t is not overwritten by $ts.\text{v}_{\text{rNew}}$ in M as far as l is concerned; see (NOT-OVERWRITTEN) for details) in the premise and $ts'.\text{v}_{\text{rNew}} \mapsto_{\sqcup} t$ (shorthand for $ts'.\text{v}_{\text{rNew}} = ts.\text{v}_{\text{rNew}} \sqcup t$) in the conclusion.

These changes ensure 'message passing' in MP: (a) the left thread issues $\langle \text{data} := 42 \rangle$@1, updating $ts_1.\text{coh}[\text{data}]$ to @1; and (b) issues $\langle \text{flag} := 1 \rangle$@2, updating $ts_1.\text{coh}[\text{flag}]$ to @2; (c) the right thread reads $\langle \text{flag} := 1 \rangle$@2, updating $ts_2.\text{coh}[\text{flag}]$ and $ts_2.\text{v}_{\text{rNew}}$ to @2; (d) it then cannot read $\langle \text{data} := 0 \rangle$@0 as $ts_2.\text{v}_{\text{rNew}} = @2 \not\sqsubseteq_{M,\text{data}} @0$.

Store Buffering with Fences As shown in SB, Intel-x86 may reorder a store and a later load on different locations. If

---

[8]The astute reader may have noticed that this condition is stronger than the coherence requirement $ts.\text{coh}[l] \sqsubseteq t$ and thus makes it redundant. Nevertheless, we explicit include the two conditions to emphasize the two requirements, namely coherence and ordering.

(RMW-FAIL)

$$
\frac{
\begin{array}{c}
l = \llbracket e \rrbracket_{ts.\text{regs}} \quad M[t] = \langle l := v \rangle \quad \llbracket rop \rrbracket_{ts.\text{regs}}(v, \bot) \\
ts.\text{coh}[l] \sqsubseteq t \quad ts.\text{v}_{\text{rNew}} \sqsubseteq_{M,l} t \\
ts' = ts[\text{regs}[r] \mapsto v, \text{coh}[l] \mapsto t, \text{v}_{\text{rNew}} \mapsto_{\sqcup} t \neq ts.\text{coh}[l] ? t]
\end{array}
}{
(r := \text{rmw } rop \ [e], ts), M \rightarrow_{tid} (\text{skip}, ts'), M
}
$$

(RMW)

$$
\frac{
\begin{array}{c}
l = \llbracket e \rrbracket_{ts.\text{regs}} \quad M[t_1] = \langle l := v_1 \rangle \quad \llbracket rop \rrbracket_{ts.\text{regs}}(v_1, v_2) \\
t_2 = |M| + 1 \quad M' = M + [\langle l := v_2 \rangle_{tid}@t_2] \quad t_2 - 1 \sqsubseteq_{M,l} t_1 \\
ts.\text{coh}[l] \sqsubseteq t_1, t_2 \quad ts.\text{v}_{\text{rNew}} \sqsubseteq_{M,l} t_1 \quad \sqcup_l ts.\text{coh}[l] \sqsubseteq t_2 \\
ts' = ts[\text{regs}[r] \mapsto v_1, \text{coh}[l] \mapsto t_2, \text{v}_{\text{rNew}} \mapsto_{\sqcup} \sqcup_l ts.\text{coh}[l] \sqcup t_2]
\end{array}
}{
(r := \text{rmw } rop \ [e], ts), M \rightarrow_{tid} (\text{skip}, ts'), M'
}
$$

Figure 5. RMW transitions of x86$_{\text{view}}$

necessary, one can prevent this by inserting fences:

$$
\begin{array}{ll}
(a) \ \text{x} := 1 & (d) \ \text{y} := 1 \\
(b) \ \text{mfence} & (e) \ \text{mfence} \\
(c) \ r_1 := \text{y} \ //= 0 & (f) \ r_2 := \text{x} \ //\neq 0
\end{array} \quad \text{(SBFence)}
$$

To model this, in the conclusion of (MFENCE) we join $ts.\text{v}_{\text{rNew}}$ with $\bigsqcup_l ts.\text{coh}[l]$, thus forbidding store buffering. Without loss of generality, assume $M = [\langle \text{x} := 1 \rangle$@1$, \langle \text{y} := 1 \rangle$@2$]$. The right thread then (d) issues $\langle \text{y} := 1 \rangle$@2, updating $ts_2.\text{coh}[\text{y}]$ to @2; (e) executes mfence, updating $ts_2.\text{v}_{\text{rNew}}$ to @2; and (f) cannot read $\langle \text{x} := 0 \rangle$@0 as $ts_2.\text{v}_{\text{rNew}} = @2 \not\sqsubseteq_{M,\text{x}} @0$.

Forwarding By strengthening x86$_{\text{view}}$ we have precluded forbidden Intel-x86 behaviors. However, x86$_{\text{view}}$ is now too strong and must be weakened to allow store forwarding:

$$
\begin{array}{ll}
(a) \ \text{x} := 1 & (d) \ \text{y} := 1 \\
(b) \ r_1 := \text{x} \ //= 1 & (e) \ r_3 := \text{y} \ //= 1 \\
(c) \ r_2 := \text{y} \ //= 0 & (f) \ r_4 := \text{x} \ //= 0
\end{array} \quad \text{(SBFwd)}
$$

While (b) and (c) are ordered, (a) and (c) are not because (b) is forwarded from (a) in the same thread, thus allowing the reordering of (a) after (b) and (c). To model this, in (LOAD) the new-read view is joined with the read message's timestamp only if it is written by a different thread. This is denoted by the conditional notation $ts'.\text{v}_{\text{rNew}} \mapsto_{\sqcup} t \neq ts.\text{coh}[l] ? t$, stating that if $t \neq ts.\text{coh}[l]$, then $ts'.\text{v}_{\text{rNew}} \mapsto_{\sqcup} t$; otherwise, $ts'.\text{v}_{\text{rNew}}$ is left unchanged. These changes then admit the behavior in SBFwd. Without loss of generality, assume $M = [\langle \text{x} := 1 \rangle$@1$, \langle \text{y} := 1 \rangle$@2$]$. The right thread (d) writes $\langle \text{y} := 1 \rangle$@2, updating $ts_2.\text{coh}[\text{y}]$ to @2; (e) reads $\langle \text{y} := 1 \rangle$@2, without updating $ts_2.\text{v}_{\text{rNew}}$ thanks to forwarding; and (f) reads $\langle \text{x} := 0 \rangle$@0 as $ts_2.\text{v}_{\text{rNew}} = @0 \sqsubseteq_{M,\text{x}} @0$.

3.4 Supporting Read-Modify-Writes (RMW)

The RMW transitions (Fig. 5) are obtained by combining the transitions of loads, stores and mfences. A failed RMW (RMW-FAIL) degenerates to a load;[9] if an RMW fails, then

---

[9]The semantics of failed RMWs in Intel-x86 is not fully agreed upon in the literature. Our model assumes a failed RMW to degenerate to a load; an alternative model may additionally assume that failed RMWs execute a

$$ts \in \text{TState} \triangleq \Big\langle \ \dots; \quad \mathsf{v_{pReady}} : \mathbb{V}; \quad \mathsf{v_{pAsync}}, \mathsf{v_{pCommit}} : \text{Loc} \to \mathbb{V}; \ \Big\rangle$$

(FLUSH)

$$\frac{l = [\![e]\!]_{ts.\text{regs}} \quad v = \sqcup_{l'} ts.\text{coh}[l']}{ts' = ts[\mathsf{v_{pAsync}} \mapsto_\sqcup \lambda l'.\ cl(l,l') \ ? \ v, \mathsf{v_{pCommit}} \mapsto_\sqcup \lambda l'.\ cl(l,l') \ ? \ v]}{(\text{flush } e, ts), M \to_{tid} (\text{skip}, ts'), M}$$

(FLUSHOPT)

$$\frac{l = [\![e]\!]_{ts.\text{regs}} \quad v = \sqcup_{l'} \ cl(l,l') \ ? \ ts.\text{coh}[l']}{ts' = ts[\mathsf{v_{pAsync}} \mapsto_\sqcup \lambda l'.\ cl(l,l') \ ? \ (v \sqcup ts.\mathsf{v_{pReady}})]}{(\text{flushopt } e, ts), M \to_{tid} (\text{skip}, ts'), M}$$

(SFENCE)

$$\frac{ts' = ts[\mathsf{v_{pReady}} \mapsto_\sqcup \sqcup_l \ ts.\text{coh}[l], \mathsf{v_{pCommit}} \mapsto_\sqcup ts.\mathsf{v_{pAsync}}]}{(\text{sfence}, ts), M \to_{tid} (\text{skip}, ts'), M}$$

(LOAD)

$$\frac{\cdots \quad ts' = ts[\cdots, \mathsf{v_{pReady}} \mapsto_\sqcup \ t \neq ts.\text{coh}[l] \ ? \ t]}{(r := \text{load } [e], ts), M \to_{tid} (\text{skip}, ts'), M}$$

(CRASH)

$$\frac{\forall l. \ \exists t. \ M[t] = \langle l := SM[l]\rangle \ \wedge \ \forall (\_, ts) \in \vec{T}. \ ts.\mathsf{v_{pCommit}}[l] \sqsubseteq_{M,l} t}{\langle \vec{T}, M\rangle \to_{\text{crash}} SM}$$

Figure 6. States and transitions of $\text{Px86}_{\text{view}}$ where the highlighted rule denotes the *extension* of **LOAD** transition from Fig. 4 as shown; the premises of **MFENCE**, **RMW** and **RMW-FAIL** are analogously extended and omitted here.

$[\![rop]\!]_{ts.\text{regs}}(v_1, \bot)$ holds (e.g., $[\![\text{cas } 4 \ 5]\!]_{rmap}(3, \bot)$ but not $[\![\text{cas } 4 \ 5]\!]_{rmap}(4, \bot)$[10]). A successful RMW **(RMW)** atomically reads from and writes to a location; if an RMW succeeds, then $[\![rop]\!]_{ts.\text{regs}}(v_1, v_2)$ holds (e.g., $[\![\text{cas } 3 \ 5]\!]_{rmap}(3, 5)$ or $[\![\text{fetch-add } 1]\!]_{rmap}(4, 5)$). Moreover, atomicity requires that there be no intervening messages on the same location between those read and written by the RMW; i.e., $t_2 - 1 \sqsubseteq_{M,l} t_1$. Lastly, as with mfences, we join $ts.\mathsf{v_{rNew}}$ with $\bigsqcup_l ts.\text{coh}'[l]$.

As we show in §5, our x86_{view} model is *equivalent* to the authoritative axiomatic model reviewed by Intel engineers.

3.5 Persistency Views

We next develop $\text{Px86}_{\text{view}}$ by extending x86_{view} with persistency. As discussed in §2.3, the key idea is *persistency views*, determining persisted messages as summarized in Table 2.

Synchronous Flush As shown in Fig. 6, in order to model the behaviour of flush instructions synchronously, we extend a thread state ts with a persistency view, $ts.\mathsf{v_{pCommit}}$. For each location l, the $ts.\mathsf{v_{pCommit}}[l]$ denotes the maximum view (timestamp) of the messages on l that have persisted

memory fence. Nevertheless, we can straightforwardly adapt our model to support this by extending **(RMW-FAIL)** with the effects of **(MFENCE)**.

[10]Here we assume compare-and-swaps are strong: they do not fail spuriously. In our Coq formalization, we also support weak compare-and-swaps.

Table 2. Informal description of persistency views

View	Past	Future
$\mathsf{v_{pReady}}$	Upper bound of past external reads (from other threads)	Lower bound of messages to be asynchronously flushed by future flushopt
$\mathsf{v_{pAsync}}[l]$	Upper bound of past flush/flushopt on the same cache line as l	Lower bound of messages on l to be persisted by future fences/updates
$\mathsf{v_{pCommit}}[l]$	Upper bound of past (1) flush l'; and (2) flushopt l' followed by fences/updates, where l' is on the same cache line as l	Lower bound of persisted messages on l to survive a crash

to NVM. Executing a flush **(FLUSH)** determines the location l, and for each location l' on the same cache line as l, joins $ts.\mathsf{v_{pCommit}}[l']$ with the maximum coherence view v, thus persisting those messages of l' propagated to the thread (i.e., all earlier writes on l'). (The *asynchronous persistency view*, $ts.\mathsf{v_{pAsync}}[l']$, will be described shortly.) After a crash **(CRASH)**, the contents of NVM, SM ('sequential memory'), satisfy the following condition for each location l: $SM[l]$ holds the value of some message on l whose timestamp t is not overwritten by any thread's persistency view on l.

This indeed establishes the invariant $I \triangleq$ commit=1 $\Rightarrow$ data=42 for COMMIT2. After executing the left thread, $M = [\langle \text{data} := 42\rangle@1]$. The right thread (b) reads $\langle \text{data} := 42\rangle@1$, updating $ts_2.\text{coh}[\text{data}]$ and $ts_2.\mathsf{v_{rNew}}$ to @1; (c) persists the message $ts_2.\text{coh}[\text{data}] = @1$, updating $ts_2.\mathsf{v_{pCommit}}[\text{data}]$ to @1; and (d) writes $\langle \text{commit} := 1\rangle@2$. After a crash, if $\langle \text{commit} := 1\rangle@2$ has persisted, then (d) must have been executed; therefore $ts_2.\mathsf{v_{pCommit}}[\text{data}] = @1$ and data = 42.

Asynchronous Flush flushopt is a weaker variant of flush that may be reordered after certain instructions, and thus its execution may be delayed until a later fence/RMW. This may improve performance when persisting multiple locations:

(a) data1 := 42	(c) if (data2 != 0) {
(b) data2 := 7	(d)　　flushopt data1
	(e)　　flushopt data2　　(COMMITOPT)
	(f)　　sfence
	(g)　　commit := 1 }

Similarly to COMMIT2, the invariant $I' \triangleq$ commit=1 $\Rightarrow$ data1=42 $\wedge$ data2=7 always holds. The sfence (f) awaits the completion of both (d) and (e), reducing I/O latency.

To model flushopt instructions, we extend a thread state ts with (1) $ts.v_{\text{pReady}}$, denoting the view to be persisted asynchronously at a subsequent flushopt; and (2) $ts.v_{\text{pAsync}}[l]$, denoting the maximum view of messages on l that have been persisted asynchronously.

The additional transitions of $\text{Px86}_{\text{view}}$ are given in Fig. 6. Executing flushopt l (**FLUSHOPT**) or flush l (**FLUSH**) joins, for each location l' on the same cache line as l, $ts.v_{\text{pAsync}}[l']$ with $ts.v_{\text{pReady}}$ and the maximum coherence view, v, of the cache line. Executing a fence in (**MFENCE**) and (**SFENCE**), or a successful RMW in (**RMW**), joins $ts.v_{\text{pCommit}}$ with $ts.v_{\text{pAsync}}$ and $ts.v_{\text{pReady}}$ with $\bigsqcup_l ts.\text{coh}[l]$. Executing a load or a failed RMW in (**LOAD**) and (**RMW-FAIL**) joins $ts.v_{\text{pReady}}$ with the read message's timestamp unless forwarded.

This allows us to establish I' for COMMITOPT. Without loss of generality, let $M = [\langle \text{data1} := 42 \rangle @1, \langle \text{data2} := 7 \rangle @2]$. The right thread then (c) reads $\langle \text{data2} := 7 \rangle @2$, updating $ts_2.\text{coh}[\text{data2}]$, $ts_2.v_{\text{rNew}}$ and $ts_2.v_{\text{pReady}}$ to @2; (d, e) asynchronously persists $ts_2.v_{\text{pReady}} = @2$ to data1 and data2, updating $ts_2.v_{\text{pAsync}}[\text{data1}]$, $ts_2.v_{\text{pAsync}}[\text{data2}]$ to @2; (f) awaits the completion of (d) and (e), updating $ts_2.v_{\text{pCommit}}[\text{data1}]$ and $ts_2.v_{\text{pCommit}}[\text{data2}]$ to @2; (g) writes $\langle \text{commit} := 1 \rangle @3$. After a crash, if $\langle \text{commit} := 1 \rangle @3$ is persisted, then (g) must have been executed; $ts_2.v_{\text{pCommit}} = [\text{data1} \mapsto @2, \text{data2} \mapsto @2]$; and thus data1 = 42 and data2 = 7.

The resulting model, $\text{Px86}_{\text{view}}$, is proven equivalent to the authoritative axiomatic model reviewed by Intel engineers [47] (modulo the fix discussed in §2.1 – see §5).

4 Fixing and Simplifying the Px86 Model

We present $\text{Px86}_{\text{axiom}}$, a new axiomatic model for Intel-x86 persistency that simplifies Px86 [47] and fixes its flaws discussed in §2.1. We present a short background on axiomatic models (§4.1); describe the baseline axiomatic model for Intel-x86 concurrency (§4.2); extend it to persistency and present $\text{Px86}_{\text{axiom}}$ (§4.3); and compare $\text{Px86}_{\text{axiom}}$ with Px86, proving their equivalence modulo our fixes in $\text{Px86}_{\text{axiom}}$ (§4.4).

4.1 Background on Axiomatic Models

Executions and Events In the literature of axiomatic (a.k.a. declarative) memory models, the traces of shared memory accesses of a program are represented as a set of *executions*, where each execution G is a graph comprising: (i) a set of events (graph nodes); and (ii) a number of relations on events (graph edges). We typically use a, b and e to range over events. Each event captures the execution of a primitive command (e.g., a load) and is a triple of the form $e = (n, \text{tid}, l)$, where $n \in \mathbb{N}$ is the (unique) *event identifier*; $\text{tid} \in \text{TId}$ identifies the executing thread; and $l \in \text{Lab}$ is the event *label*. Event labels are defined by the underlying memory model; for Intel-x86 a label l may be (1) (R, x, v) for reading (loading) value v from location x; (2) (W, x, v) for writing (storing) value v to location x; (3) (U, x, v, v') for a successful update (RMW) modifying x

to v' when its value matches v; (4) MF for executing an mfence. The functions loc, rval and wval respectively project the location, the read value and the written value of a label, where applicable. For instance, $\text{loc}(l) = x$ and $\text{wval}(l) = v$ for $l = (\text{W}, x, v)$. The functions thrd and lab respectively project the thread identifier and the label of an event.

Notation Given a relation r on a set A, we write $r^?$ and r^+ for the reflexive and transitive closures of r, respectively. We write r^{-1} for the inverse of r; $[A]$ for the identity relation on A, i.e., $\{(a, a) \mid a \in A\}$; and A_x for $\{a \in A \mid \text{loc}(a) = x\}$. We write ri for the internal subset of r (on events of the same thread), i.e., $ri \triangleq \{(a, b) \in r \mid \text{thrd}(a) = \text{thrd}(b)\}$; and re for the external subset of r (on events of different threads). Finally, we write $r_1; r_2$ for the relational composition of r_1 and r_2, i.e., $\{(a, b) \mid \exists c. (a, c) \in r_1 \wedge (c, b) \in r_2\}$.

Definition 4.1 (Executions). An *execution*, G, is a tuple of the form $(E, \text{po}, \text{rf}, \text{co})$, where:

- E is a set of events, including a set of *initialisation* events, $I \subseteq E$, comprising a single write event with label $(\text{W}, x, 0)$ for each $x \in \text{Loc}$. The set of *read* events in E is: $R \triangleq \{e \in E \mid \exists x, v. \text{lab}(e) = (\text{R}, x, v)\}$; the sets of *writes* ($W$), *RMW* ($U$) and *memory fence* ($MF$) events are analogous.
- $\text{po} \subseteq E \times E$ denotes the *'program-order'* relation, defined as a disjoint union of strict total orders, each ordering the events of one thread, together with $I \times (E \setminus I)$ that orders initialisation events before all others.
- $\text{rf} \subseteq (W \cup U) \times (R \cup U)$ denotes the *'reads-from'* relation on events of the same location with matching values; i.e., $(a, b) \in \text{rf} \Rightarrow \text{loc}(a) = \text{loc}(b) \wedge \text{wval}(a) = \text{rval}(b)$. Moreover, rf is total and functional on its range, i.e., every read/update is related to exactly one write/update. A read/update may be rf-related to an initialisation write.
- $\text{co} \subseteq E \times E$ is the *'coherence-order'*, defined as the disjoint union of relations $\{\text{co}_x\}_{x \in \text{Loc}}$, such that each co_x is a strict total order on $W_x \cup U_x$ and $I_x \times ((W_x \cup U_x) \setminus I) \subseteq \text{co}_x$.

In the context of an execution graph $(E, \text{po}, \text{rf}, \text{co})$, we define the *'from-reads'* relation as $\text{fr} \triangleq \text{rf}^{-1}; \text{co}$. Note that in this initial stage, executions are unrestricted: there are few constraints on rf and co. Such restrictions are determined by the set of model-specific *consistent* executions. We next define execution consistency for several models.

4.2 The $\text{x86}_{\text{axiom}}$ Model [3]

As the baseline axiomatic model for Intel-x86, we use that of Alglave et al. [3], presented in Fig. 7, which we refer to as $\text{x86}_{\text{axiom}}$.[11] We choose $\text{x86}_{\text{axiom}}$ as the baseline as it is stylistically similar with $\text{Armv8}_{\text{axiom}}$ [44], thus allowing a more uniform treatment of Intel-x86 and Armv8 persistency.

---

[11]For clarity, we rename the relations and axioms in [3] to highlight its similarity with the axiomatic model for Armv8 concurrency [44].

Kyeongmin Cho, Sung-Hwan Lee, Azalea Raad, and Jeehoon Kang

$$obs = co \cup rfe \cup fre$$
$$dob = ([W \cup U \cup R]; po; [W \cup U \cup R]) \setminus (W \times R)$$
$$bob = [W \cup U \cup R]; po; [MF]; po; [W \cup U \cup R]$$
$$ob = obs \cup dob \cup bob$$

$$(rf; po^?) \text{ irreflexive} \tag{co-rw}$$
$$(fr; po) \text{ irreflexive} \tag{co-wr}$$
$$ob \text{ acyclic} \tag{external}$$

Figure 7. The x86$_{axiom}$ model [3]

$$(\text{axioms of x86}_{axiom} \text{ (Fig. 7)})$$
$$fob = [W \cup U \cup R]; po; [FL]$$
$$\cup ([U \cup R] \cup ([W]; po; [MF \cup SF])); po; [FO]$$
$$\cup [W]; (po; [FL])^?; (po \cap CL); [FO]$$
$$ob = obs \cup dob \cup bob \cup fob \cup pf \cup fp \tag{redefined}$$
$$pf \subseteq (obs \cup dob \cup bob \cup fob \cup fp)^+ \tag{pf-min}$$
$$P = dom(pf; ([FL] \cup ([FO]; po; [MF \cup SF \cup U])))$$
$$\forall l. \exists w. SM(l) = \text{wval}(w) \ \wedge \ (P \times \{w\}) \cap Loc \subseteq co^? \tag{persist}$$

Figure 8. The Px86$_{axiom}$ model

The co-rw ('coherence-read-write') axiom requires that loads not read from later stores; co-wr ('coherence-write-read') ensures that loads do not read values overwritten by earlier stores; and external, ensures that externally visible events can be linearized with respect to the *'ordered-before'* relation (ob). The existence of such a globally-agreed order of events makes x86$_{axiom}$ multi-copy-atomic.

The ob relation enforces the order (a, b) if: (1) a is a store overwritten by b (co); (2) a is a store read by b in a different thread (rfe); (3) a reads a value overwritten by b in a different thread (fre); (4) a, b are accesses by the same thread and $(a, b) \notin W \times R$ (dob); or (5) a, b are accesses by the same thread and are separated by a fence (bob).

Coherence between two writes (resp. two reads) is derived from the axioms. Specifically, $co \cup ([W \cup U]; po; [W \cup U]) \subseteq ob$ and acyclicity of ob ensure irreflexivity of $co; po$. Similarly, $(fre; rfe) \cup (fri; rfi) \cup ([U \cup R]; po; [U \cup R]) \subseteq ob$ and acyclicity of ob ensure irreflexivity of $fr; rf; po$.

4.3 The Px86$_{axiom}$ Model

We extend x86$_{axiom}$ with persistency semantics and develop the Px86$_{axiom}$ model as presented in Fig. 8. We first define:

- FL and FO: the set of synchronous flush (flush) and asynchronous flush (flushopt) events, respectively;
- SF: the set of sfence events;
- $pf \subseteq (W \cup U) \times (FL \cup FO)$: the *'persists-from'* relation, relating each flush to the co-latest store for each location

persisted by the flush. This is analogous to the rf relation; however, while rf relates a load to a *single* store, pf may relate a flush to multiple stores (one for each location) on the same cache line.

- $fp \triangleq pf^{-1}; co$: the *'from-persists'* relation (analogous to fr), relating a flush to co-later stores (*cf.* $fr \triangleq rf^{-1}; co$).

The ob relation is extended with fob ('flush-ordered-before'), ordering earlier events and a later flush as per the Intel manual [20]. Furthermore, pf and fp are included in ob for the same reason rf and fr are; i.e., because Intel-x86 is multi-copy atomic. P denotes the set of writes that must be persisted, i.e., those writes that are persisted by a synchronous (FL) or an asynchronous flush (FO) followed by a fence ($MF \cup SF \cup U$).[12] The persist axiom states that in case of a crash, the persisted value (in NVM) of each location l in $SM[l]$ should not be coherence-before the writes in P. For simplicity, the pf-min axiom ensures that P is minimal, i.e., a flush persists only those writes that are strictly ordered before it. However, this minimality axiom is optional (Lemma 4.2).

Lemma 4.2. *A behavior is allowed under Px86$_{axiom}$ with axiom pf-min iff it is allowed under Px86$_{axiom}$ without pf-min.*

Proof. The proof is given in [9, Appendix C]. ∎

4.4 Comparing Px86$_{axiom}$ to Px86 in [47]

Fix Our Px86$_{axiom}$ model indeed fixes the Px86 shortcomings described in §2.1. In particular, as discussed in §2.1, we first strengthen Px86 to SPx86 by additionally requiring that flush instructions behave synchronously – see [9, Figs. 18 and 19] for the definitions of Px86 and SPx86.[13] In Theorem 4.3 below we then prove that Px86$_{axiom}$ and SPx86 are equivalent.

Theorem 4.3. *A behavior is allowed under SPx86 iff it is allowed under Px86$_{axiom}$.*

Proof. The proof is given in [9, Appendix D]. ∎

The Px86 and SPx86 models are based on the axiomatic Intel-x86 model known as TSO [41, 49], henceforth referred to as x86$_{man}$ (given in [9, Fig. 18]). As such, in order to prove Theorem 4.3 we first show that x86$_{man}$ and x86$_{axiom}$ are equivalent. In particular, existing equivalence results between x86$_{man}$ and x86$_{axiom}$ cover loads and stores only and not RMWs and fences [2]. We extend this result for the first time to cover RMWs and fences in Theorem 4.4 below.

Theorem 4.4. *A behavior is allowed under x86$_{man}$ iff it is allowed under x86$_{axiom}$.*

Proof. The proof is given in [9, Appendix D.2]. ∎

---

[12]One may expect an asynchronous flush to complete also when the thread terminates. But this is defined neither in the Intel manual [20] nor in its libraries [19]. We thus assume an asynchronous flush *not* to be completed when a thread terminates. However, we can easily change this by appending *TERM* to $MF \cup SF \cup U$, where *TERM* denotes thread termination. Analogously, we can adapt Px86$_{view}$ in §3 to account for terminated threads.

[13]For clarity, we adapted Px86 from [47] to match our style.

Simplification Our Px86$_{\text{axiom}}$ model is simpler than Px86 in [47] in the following aspects:

- While tso ('total store order'), nvo ('non-volatile order'), and P ('persisted stores') components of Px86 are *existentially quantified*, thus increasing non-determinism, the analogous ob and P in Px86$_{\text{axiom}}$ are *constructed*.

- While the conditions for intra-thread, inter-thread, and CPU-NVM communications are intertwined in Px86, they are separated and constrained by distinct axioms in Px86$_{\text{axiom}}$: intra-thread ones by CO-RW and CO-WR, inter-thread ones by EXTERNAL and CPU-NVM ones by PERSIST. To achieve this, Px86$_{\text{axiom}}$ orders fewer flush events than the Intel reference manual [20] does; e.g., unlike the manual, Px86$_{\text{axiom}}$ does not order FL before R.

- Px86$_{\text{axiom}}$ may optionally require the minimality of pf, which is beneficial for e.g., reducing the search space significantly for stateless model checking. By contrast, Px86 does not require a similar minimality in tso.

As we show in §5, the constructive and succinct nature of Px86$_{\text{axiom}}$ and its stylistic similarity to the axiomatic Armv8 model [44] make it easier to prove its equivalence to Px86$_{\text{view}}$.

5 Equivalence of Px86$_{\text{view}}$ and Px86$_{\text{axiom}}$

To evaluate the fidelity of Px86$_{\text{view}}$, we show that it is equivalent to Px86$_{\text{axiom}}$. To do this, we first prove the equivalence of x86$_{\text{view}}$ and x86$_{\text{axiom}}$ by adapting the equivalence proof of the view-based and axiomatic models for Armv8 concurrency [44], and then generalize it to Intel-x86 persistency. All theorems in this section are mechanized in Coq [10].

Equivalence of x86$_{\text{view}}$ and x86$_{\text{axiom}}$ In order to reuse the existing equivalence proof of the view-based and axiomatic models for Armv8 concurrency [44] maximally, we appeal to a new model, x86$_{\text{prom}}$, the *promising* view-based model for Intel-x86 concurrency, as the bridge between x86$_{\text{view}}$ and x86$_{\text{axiom}}$. Compared to x86$_{\text{view}}$, x86$_{\text{prom}}$ additionally allows 'promises', modeling speculative writes (see §6.2). Specifically, we employ the following proof strategy:

(1) We prove that x86$_{\text{view}}$ and x86$_{\text{prom}}$ are equivalent and that promises do not enable additional behaviors as their effect is cancelled out by concurrency views (Lemma 5.1).

(2) We prove that x86$_{\text{prom}}$ and x86$_{\text{axiom}}$ are equivalent by adapting the analogous equivalence proof for Armv8 concurrency [44] as x86$_{\text{prom}}$ and x86$_{\text{axiom}}$ respectively have the same style as the (view-based) Armv8$_{\text{view}}$ and (axiomatic) Armv8$_{\text{axiom}}$ models of Armv8 concurrency.

Combining the two steps we then establish the desired equivalence in Theorem 5.2.

Lemma 5.1. *A behavior is allowed under x86$_{\text{prom}}$ iff it is allowed under x86$_{\text{view}}$.*

$$
\begin{array}{llr}
 & \ldots \text{ (the language for Intel-x86 in Fig. 3)} & \\
s \in \text{St} & ::= \cdots & \textit{statement} \\
 & \mid r := \text{load}_{xcl,rk} \, [e] & \textit{load} \\
 & \mid r_{\text{succ}} := \text{store}_{xcl,wk} \, [e_1] \, e_2 & \textit{store} \\
 & \mid \text{isb} \mid \text{dmb}.f \mid \text{dsb}.f & \textit{fence} \\
 & \mid \text{flushopt } e & \textit{flush} \\
f \in \text{F} & ::= \text{ld} \mid \text{st} \mid \text{sy} & \textit{order} \\
\end{array}
$$

$$
\begin{array}{llr}
xcl \in \mathbb{B} & ::= \text{false} \mid \text{true} & \textit{exclusivity} \\
rk \in \text{RK} & ::= \text{pln} \mid \text{wacq} \mid \text{acq} & \textit{read kind} \\
wk \in \text{WK} & ::= \text{pln} \mid \text{wrel} \mid \text{rel} & \textit{write kind} \\
\end{array}
$$

Figure 9. The Armv8 concurrency/persistency language

Theorem 5.2. *A behavior is allowed under x86$_{\text{view}}$ iff it is allowed under x86$_{\text{axiom}}$.*

Equivalence of Px86$_{\text{view}}$ and Px86$_{\text{axiom}}$ We next extend Theorem 5.2 to Intel-x86 persistency (Theorem 5.3). To do this, we relate each view of an x86$_{\text{view}}$ execution to a set of events in the corresponding x86$_{\text{axiom}}$ execution; similarly for the persistency views in Px86$_{\text{view}}$. For example, the v_{pCommit} view of a thread state is related to the set P of persisted writes in the corresponding Px86$_{\text{axiom}}$ execution. This then allows us to prove the equivalence of Px86$_{\text{view}}$ and Px86$_{\text{axiom}}$.

Theorem 5.3. *A behavior is allowed under Px86$_{\text{axiom}}$ iff it is allowed under Px86$_{\text{view}}$.*

6 View-Based and Axiomatic Models for Armv8 Persistency

In §3–5 we presented view-based and axiomatic models for Intel-x86 persistency and proved their equivalence. We next do the same for Armv8. As Intel-x86 and Armv8 persistency are highly similar, we focus on their differences (§6.1; see [9, Appendix B] for the full details). We then present the view-based Armv8 persistency model (§6.2), fix and simplify the axiomatic model for Armv8 persistency due to Raad et al. [48] as discussed in §2.2 (§6.3), and finally prove the equivalence of our view-based and axiomatic models (§6.4).

6.1 Armv8 versus Intel-x86 Persistency

We present the Armv8 language in Fig. 9, which is similar to that for Intel-x86 (Fig. 3), modulo the following:

Ordering: Armv8 ordering constraints are weaker and more elaborate than those of Intel-x86. Specifically, Armv8 loads and stores are annotated with *access ordering* constraints (rk or wk in Fig. 9). Moreover, Armv8 fences are more diverse: isb orders loads and later dependent accesses; dmb.f orders accesses according to the ordering constraint f (see Fig. 9); and dsb.sy additionally awaits the completion of pending flush instructions.

Exclusivity: Unlike Intel-x86, Armv8 supports exclusive *load-link* and *store-conditional* instructions [25] that (if

successful) prohibit intervening stores between the load and store. Exclusive instructions are more primitive than RMWs: RMWs can be implemented via exclusive instructions but not vice versa.[14] As such, loads and stores are annotated with *exclusivity* tags (*xcl* in Fig. 9).

Flush: All Armv8 flushes are asynchronous (flushopt).

As we describe shortly, these differences are largely orthogonal to modeling persistency, except for the relaxed ordering of writes. Specifically, Armv8 allows (unlike Intel-x86) speculative execution of writes, interacting with NVM in an interesting way. To see this, we review the relaxed 'load buffering' behavior of Armv8 due to speculative writes:

$$
\begin{array}{l|l}
(a)\ r_1 := \mathsf{y}\ //= 1 & (c)\ r_2 := \mathsf{x}\ //= 1 \\
(b)\ \mathsf{x} := 1 & (d)\ \mathsf{y} := 1
\end{array}
\qquad \text{(LB)}
$$

As Armv8 does not order a read and a subsequent write, (a) and (b) may be reordered; similarly for (c) and (d). As such, Armv8 allows an execution where (b), (d), (a), and (c) are executed in order, thus allowing the $r_1 = r_2 = 1$ behavior.

6.2 PArmv8_view: View-Based Armv8 Persistency

As with Px86_view, the view-based Armv8 persistency model, PArmv8_view, follows the same interleaving model over the history of stores. However, PArmv8_view differs from Px86_view in that (1) its views are more elaborate; and (2) it introduces *promises* to model speculative writes.

Views To model the ordering constraints and exclusivity of Armv8, the PArmv8_view thread state in [9, Fig. 15] has additional view components compared to x86_view in Fig. 4. These additional components are those of Armv8_view [44]; i.e., the PArmv8_view thread state is that of Armv8_view extended with persistency views ($\mathsf{v}_{\mathsf{pReady}}$, $\mathsf{v}_{\mathsf{pAsync}}$ and $\mathsf{v}_{\mathsf{pCommit}}$ in §3.5).

Promises The additional views, however, are not sufficient to model LB: without further instrumentation, the model remains interleaving, where either (a) or (c) is executed first, reading the initial value 0.

To model speculative writes, Armv8_view [44] introduces the notion of a *promise*: a message that may be speculatively added to the memory (or *promised*) without executing a store, provided that the promised message is later substantiated (or *fulfilled*) by executing a corresponding store. Put formally, a thread state ts contains the set $ts.\mathsf{prom}$ of the message ids that are promised by the thread but not yet fulfilled.

Using promises, we can model the LB behavior as follows, where tid_1 and tid_2 denote the left and right threads, respectively: (b-prom) tid_1 promises $\langle \mathsf{x} := 1 \rangle_{tid_1}@1$ with $ts_1.\mathsf{prom} = \{@1\}$; ($c$) tid_2 reads $\langle \mathsf{x} := 1 \rangle_{tid_1}@1$, updating $ts_2.\mathsf{coh}[\mathsf{x}]$ and

$ts_2.\mathsf{v}_{\mathsf{rOld}}$ ('old-read view'[15]) to @1; (d) tid_2 writes $\langle \mathsf{y} := 1 \rangle_{tid_2}@2$, updating $ts_2.\mathsf{coh}[\mathsf{y}]$ and $ts_2.\mathsf{v}_{\mathsf{wOld}}$ ('old-write view') to @2; (a) tid_1 reads $\langle \mathsf{y} := 1 \rangle_{tid_2}@2$, updating $ts_1.\mathsf{coh}[\mathsf{y}]$ and $ts_1.\mathsf{v}_{\mathsf{rOld}}$ to @2; and (b-fulfill) tid_1 fulfills $\langle \mathsf{x} := 1 \rangle_{tid_1}@1$, yielding $ts_1.\mathsf{prom}=\emptyset$ and $ts_1.\mathsf{v}_{\mathsf{wOld}}=@1$. Effectively, the write (b) is speculatively executed before the read (a) is executed.

To ensure that all speculations are substantiated, we require that a thread state's prom set be empty at the end of an execution; otherwise, the execution is deemed invalid.

Promises and Persistency The promises in PArmv8_view similarly model speculative writes. Indeed, promises are largely orthogonal to persistency, except in the case of a crash. Specifically, in case of a crash in the presence of unfulfilled promises, we must determine the NVM contents.

On the one hand, one may argue that unfulfilled promises should persist (remain in NVM) as they have been made visible to other threads. To see this, consider Commit2 and suppose that the left thread promises $\langle \mathsf{data} := 42 \rangle@1$ which is yet unfulfilled, the right thread reads $\langle \mathsf{data} := 42 \rangle@1$ and writes $\langle \mathsf{commit} := 1 \rangle@2$, and then a crash occurs. If upon recovery $\langle \mathsf{commit} := 1 \rangle@2$ has persisted, then $\langle \mathsf{data} := 42 \rangle@1$ (which is an unfulfilled promise) should have also persisted.

On the other hand, one may argue that unfulfilled promises should not persist as they are not substantiated by a store. For example, suppose that the left thread in Commit2 promises to write $\langle \mathsf{data} := 23 \rangle@1$ without fulfilling it, and then it crashes. The promised write then should not persist as it is unsubstantiated; i.e., otherwise 23 appears *out-of-thin-air*.

To resolve this dilemma, we allow an execution to crash only if it has no unfulfilled promises. This then admits only the desired behaviors in Commit2: the execution cannot crash if either $\langle \mathsf{data} := 42 \rangle@1$ or $\langle \mathsf{data} := 23 \rangle@1$ is promised and not yet fulfilled. At first glance, this may seem restrictive as micro-architecturally an execution may crash even in the presence of uncommitted speculative writes. However, when this is the case, executing the remaining instructions to commit speculative writes does not constrain the NVM contents. Moreover, we formally justify our design by proving that PArmv8_view and PArmv8_axiom are equivalent (see §6.4).

6.3 PArmv8_axiom: Fixing and Simplifying PArmv8

We use the model of Pulte et al. [44] as the baseline axiomatic model for Armv8 concurrency, presented as Armv8_axiom in [44, Appendix D].[16] The Armv8_axiom model is equivalent to the authoritative axiomatic model in [43] which is reviewed by Arm engineers. Note that Armv8_axiom has the same style as x86_axiom in Fig. 7, except that: (1) all coherence constraints are captured by a single axiom (INTERNAL) since (CO-WW)

---

[14]While Armv8.1 also supports RMWs, they are currently missing in Armv8_view and Armv8_axiom [44]. Accordingly, we do not extend them to support RMWs as this is orthogonal to our objectives here.

[15]While reads update $\mathsf{v}_{\mathsf{rNew}}$ in x86_view, they update $\mathsf{v}_{\mathsf{rOld}}$ in Armv8_view. We refer the reader to [9, Appendix B] for more details.

[16]We refactor the relations in [44] to replace dmb with dmb ∪ dsb. The latter is a straightforward extension as dsb is strictly stronger than dmb [4].

$$(\text{axioms of Armv8}_{\text{axiom}} \; [44] \; [9, \text{Fig. 20}])$$

$$\text{fob} = [W \cup R]; \text{po}; [\text{dmb.sy} \cup \text{dsb.sy}]; \text{po}; [FO]$$

$$\cup \; [W \cup R]; (\text{po} \cap \text{CL}); [FO]$$

$$\text{ob} = \text{obs} \cup \text{dob} \cup \text{aob} \cup \text{bob} \cup \text{fob} \cup \text{pf} \cup \text{fp} \qquad (\textit{redefined})$$

$$\text{pf} \subseteq (\text{obs} \cup \text{dob} \cup \text{aob} \cup \text{bob} \cup \text{fob} \cup \text{fp})^{+} \qquad (\textsc{pf-min})$$

$$P = dom(\text{pf}; [FO]; \text{po}; [\text{dsb.sy}])$$

$$\forall l. \; \exists w. \; SM(l) = \text{wval}(w) \; \wedge \; (P \times \{w\}) \cap \text{Loc} \subseteq \text{co}^{?} \qquad (\textsc{persist})$$

Figure 10. The PArmv8$_{\text{axiom}}$ model

and (CO-RR) no longer follow from the other axioms;[17] (2) the ob component of Armv8$_{\text{axiom}}$ is more elaborate, modeling the weak ordering constraints of Armv8; and (3) Armv8$_{\text{axiom}}$ has an additional axiom (ATOMIC) that ensures the exclusivity of load-link/store-conditional instructions.

We next define an axiomatic model for Armv8 persistency, PArmv8$_{\text{axiom}}$ in Fig. 10, by extending Armv8$_{\text{axiom}}$ with persistency in the same style as Px86$_{\text{axiom}}$. The key differences from Px86$_{\text{axiom}}$ are that: (1) flush instructions impose different ordering constraints; (2) PArmv8$_{\text{axiom}}$ has no strong flush instructions; and (3) optimized flush instructions are guaranteed to commit only upon executing dsb.sy fences.

The PF-MIN axiom is optional as in Px86$_{\text{axiom}}$ (Lemma 4.2).

Fix Our PArmv8$_{\text{axiom}}$ model fixes the PArmv8 problem discussed in §2.2. Put formally, we prove the equivalence of PArmv8$_{\text{axiom}}$ and SPArmv8 which denotes strengthening PArmv8 by extending ob with pf and fp.

Theorem 6.1. *A behavior is allowed under SPArmv8 iff it is allowed under PArmv8$_{\text{axiom}}$.*

Proof. The proof is given in [9, Appendix E].

6.4   Equivalence of PArmv8$_{\text{view}}$ and PArmv8$_{\text{axiom}}$

Finally, we prove that PArmv8$_{\text{axiom}}$ and PArmv8$_{\text{view}}$ are equivalent by generalizing the analogous concurrency result in [44, Theorem 6.1] (showing that Armv8$_{\text{axiom}}$ and Armv8$_{\text{view}}$ are equivalent) and extending it with persistency.

Theorem 6.2. *A behavior is allowed under PArmv8$_{\text{axiom}}$ iff it is allowed under PArmv8$_{\text{view}}$.*

Proof. The proof is mechanized in [10].

7   Model Checking Persistency Patterns

We develop a stateless model checker for PArmv8$_{\text{view}}$ by generalizing and extending the Armv8$_{\text{view}}$ model checking framework in [44] to support persistency and account for crashes (§7.1). We use our model checker to verify representative persistent synchronization examples, including the ATOMICPERSISTS example [45] that emulates a persistent

---

[17]We could replace INTERNAL with irreflexivity of po; (co ∪ rf ∪ fr ∪ fr; rf) for uniformity with x86$_{\text{axiom}}$. We forwent this to use Armv8$_{\text{axiom}}$ [44] as is.

transaction [9, Appendix F]. Our model checking tool and verified examples are open source and publicly available [10].

7.1   Model Checking Tool

Model Checking Tool for Armv8$_{\text{view}}$ We first briefly review the baseline model checking tool for Armv8$_{\text{view}}$ [44], which is a part of RMEM [5]. The tool consists of two parts: the executable model for sequential semantics of Armv8 ISA written in Sail [5]; and the executable memory model for concurrency written in Lem [39]. The former is adopted from [43], and the latter is split into two modes: the "promise-mode" which approximately enumerates the reachable final memories; and the "non-promise-mode" that checks if each potentially reachable final memory is actually reachable by thread executions to the end without promises. The two-mode execution is sound for the Armv8$_{\text{view}}$ model: a reachable state in Armv8$_{\text{view}}$ is also reachable by first promising to write all messages and then fulfilling the promises by executing the threads [44, Theorem 7.1].

Extension for PArmv8$_{\text{view}}$ We extend the model checking tool for Armv8$_{\text{view}}$ as follows: (1) we add persistency instructions to the executable model for sequential semantics in Sail; (2) we add persistency views to the executable memory model for Armv8$_{\text{view}}$ in Lem; (3) we enumerate not only final but also intermediate reachable memories in the promise–mode; and (4) we allow each thread's execution to stop amidst the non-promise-mode; and (5) we enumerate all post-crash states from the reachable states of intermediate memories and persistency views.

The performance of the resulting model checking algorithm for PArmv8$_{\text{view}}$ is similar to that for Armv8$_{\text{view}}$ because (1), (2), (4), (5) introduce only a constant-factor overhead; and the number of intermediate memories in (3) is usually dominated by that of final memories.

8   Related and Future Work

Related Work on Hardware Persistency Models Existing literature includes several works on formalising and testing hardware persistency models [11, 26, 28, 36, 42, 47, 48]. As discussed in detail in §2–6, the works of [47, 48] are closest to ours. Pelley et al. [42] propose several persistency models including epoch persistency; however, these models have not been adopted by mainstream architectures as of yet. Condit et al. [11], Joshi et al. [26] describe epoch persistency under x86-TSO [49]. Liu et al. [36] develop the PMTest testing framework for finding persistency bugs in software running over hardware models. Izraelevitz et al. [23] give a formal semantics of epoch persistency under release consistency [16]. As discussed in §1, the PTSO model of Raad and Vafeiadis [46] formalises epoch persistency under x86$_{\text{man}}$ (TSO) as a *proposal* for Intel-x86. However, PTSO is rather different from the existing Intel-x86 persistency model in Intel [20] in that it does not support the fine-grained Intel primitives for

selectively persisting cache lines (`flush` and `flushopt`), and instead proposes coarse-grained instructions (for persisting *all* locations at once) that do not exist in Intel-x86.

Khyzha and Lahav [28] recently developed the PTSO$_{syn}$ model for Intel-x86 that satisfies the three properties of (A)–(C) discussed in §1. In particular, PTSO$_{syn}$ supports persistent synchronization patterns even in the presence of I/O (B) as it also models `flush` instructions synchronously like Px86$_{view}$ (§3.5). However, this problem of asynchronous modeling of `flush` regarding I/O is not discussed in the paper.

Intel recently introduced Optane Persistent Memory 200 Series [22] that feature Enhanced Asynchronous DRAM Refresh (eADR), which treats processor caches as persistent (rather than volatile) by automatically flushing cache data to NVM in case of a crash. When eADR is available, a store is guaranteed to persist when made visible to other threads (e.g., after executing an `mfence`/`sfence`, but not `clflush`/`clflushopt`). Nevertheless, we believe that our contributions still stand for the following reasons. First, to ensure backwards compatibility, programs must support persistency in the absence of eADR. That is, a *correct* NVM program must defensively check whether eADR is enabled, and if not insert appropriate `clflush` or `clflushopt` instructions per our models. Second, eADR may increase runtime cost. For example, to flush cache data to NVM when a crash occurs, eADR must drain more power with higher voltage level or larger capacity, the impact of which on power consumption has not been thoroughly analyzed as of yet. The increased power consumption may affect embedded systems worse, and to our knowledge, Arm currently has no plans for supporting an eADR-like feature in Armv8.

Related Work on Software Persistency Models The literature on software persistency is more limited [8, 17, 30]. Kolli et al. [30] propose *acquire-release persistency*, an analogue to release-acquire consistency in C/C++. Gogte et al. [17] propose *synchronisation-free regions* (regions delimited by synchronisation operations or system calls). Although both approaches enjoy good performance, their semantic models are rather fine-grained, paving the way towards more coarse-grained transactional models [6, 19, 31, 48, 51, 52].

Related Work on Verification There are several works on implementing and verifying algorithms that operate on NVM. Friedman et al. [15] developed persistent queue implementations using Intel-x86 persist instructions (e.g., `flush`). Similarly, Zuriel et al. [54] developed persistent set implementations using Intel-x86 persist instructions. Derrick et al. [14] provided a formal correctness proof of the implementation in [54]. All three of [14, 15, 54] assume that the underlying concurrency model is sequential consistency [33], rather than x86$_{man}$ (TSO). Recently, Raad et al. [45] developed a persistent program logic for verifying programs under the Px86 model. Finally, Kokologiannakis et al. [29] recently formalised the consistency and persistency semantics of the

Linux ext4 filesystem, and developed a model-checking algorithm and tool for verifying the consistency and persistency behaviors of ext4 applications such as text editors

Future Work We plan to build on this work in several ways. First, we will empirically validate the proposed models w.r.t. NVM hardware using custom SoC (ASIC or FPGA) that captures the traffic between CPU and NVM, as proposed also in [47]. Second, we will explore language-level persistency by researching persistency extensions of high-level languages such as C/C++. This will liberate programmers from understanding hardware-specific persistency guarantees and make persistent programming more accessible. Third, we will first specify existing persistent libraries such as PMDK [19] and then use our model checker (§7) to verify their implementations against our specifications. Lastly, in the spirit of persistency semantics defining the order in which writes are propagated to NVM in DIMM slots, we will study the semantics in the presence of accelerators (e.g., CXL [13] and CCIX [12]), defining the order in which writes are propagated to accelerators in PCIe slots or other peripheral interconnects.

Acknowledgments

References

[1] 2020. The Coq Proof Assistant. https://coq.inria.fr/

[2] Jade Alglave. 2012. A Formal Hierarchy of Weak Memory Models. *Form. Methods Syst. Des.* 41, 2 (2012).

[3] Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory. *ACM Trans. Program. Lang. Syst.* 36, 2 (2014).

[4] Arm. 2020. Arm architecture reference manual Armv8, for Armv8-A architecture profile (DDI 0487F.b). https://static.docs.arm.com/ddi0487/fb/DDI0487F_b_armv8_arm.pdf

[5] Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Alastair Reid, Kathryn E. Gray, Robert M. Norton, Prashanth Mundkur, Mark Wassell, Jon French, Christopher Pulte, Shaked Flur, Ian Stark, Neel Krishnaswami, and Peter Sewell. 2019. ISA Semantics for ARMv8-a, RISC-v, and CHERI-MIPS. *Proc. ACM Program. Lang.* 3, POPL (2019). https://doi.org/10.1145/3290384

[6] Hillel Avni, Eliezer Levy, and Avi Mendelson. 2015. Hardware Transactions in Nonvolatile Memory. In *Proceedings of the 29th International Symposium on Distributed Computing - Volume 9363* (Tokyo, Japan) *(DISC 2015)*. Springer-Verlag, Berlin, Heidelberg, 617–630. https://doi.org/10.1007/978-3-662-48653-5_41

[7] H. Alan Beadle, Wentao Cai, Haosen Wen, and Michael L. Scott. 2020. Nonblocking Persistent Software Transactional Memory. In *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (San Diego, California) *(PPoPP '20)*. Association for Computing Machinery, New York, NY, USA, 429–430. https://doi.org/10.1145/3332466.3374506

[8] Dhruva R. Chakrabarti, Hans-J. Boehm, and Kumud Bhandari. 2014. Atlas: Leveraging Locks for Non-volatile Memory Consistency. *SIGPLAN Not.* 49, 10 (Oct. 2014), 433–452. https://doi.org/10.1145/2714064.2660224

[9] Kyeongmin Cho, Sung-Hwan Lee, Azalea Raad, and Jeehoon Kang. 2021. Appendix for Revamping Hardware Persistency Models: View-Based and Axiomatic Persistency Models for Intel-x86 and Armv8. https://cp.kaist.ac.kr/pmem

[10] Kyeongmin Cho, Sung-Hwan Lee, Azalea Raad, and Jeehoon Kang. 2021. Artifact for Revamping Hardware Persistency Models: View-Based and Axiomatic Persistency Models for Intel-x86 and Armv8. https://doi.org/10.1145/3410292

[11] Jeremy Condit, Edmund B. Nightingale, Christopher Frost, Engin Ipek, Benjamin Lee, Doug Burger, and Derrick Coetzee. 2009. Better I/O Through Byte-addressable, Persistent Memory. In *Proceedings of the ACM SIGOPS 22Nd Symposium on Operating Systems Principles* (Big Sky, Montana, USA) *(SOSP '09)*. ACM, New York, NY, USA, 133–146. https://doi.org/10.1145/1629575.1629589

[12] CCIX Consortium. [n.d.]. Cache Coherent Interconnect for Accelerators. https://www.ccixconsortium.com/

[13] CXL Consortium. [n.d.]. Compute Express Link. https://www.computeexpresslink.org/

[14] John Derrick, Simon Doherty, Brijesh Dongol, Gerhard Schellhorn, and Heike Wehrheim. 2019. Verifying Correctness of Persistent Concurrent Data Structures. In *Formal Methods – The Next 30 Years*, Maurice H. ter Beek, Annabelle McIver, and José N. Oliveira (Eds.). Springer International Publishing, Cham, 179–195.

[15] Michal Friedman, Maurice Herlihy, Virendra Marathe, and Erez Petrank. 2018. A Persistent Lock-Free Queue for Non-Volatile Memory. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Vienna, Austria) *(PPoPP '18)*. Association for Computing Machinery, New York, NY, USA, 28–40. https://doi.org/10.1145/3178487.3178490

[16] Kourosh Gharachorloo, Daniel Lenoski, James Laudon, Phillip Gibbons, Anoop Gupta, and John Hennessy. 1990. Memory Consistency and Event Ordering in Scalable Shared-memory Multiprocessors. *SIGARCH Comput. Archit. News* 18, 2SI (May 1990), 15–26. https://doi.org/10.1145/325096.325102

[17] Vaibhav Gogte, Stephan Diestelhorst, William Wang, Satish Narayanasamy, Peter M. Chen, and Thomas F. Wenisch. 2018. Persistency for Synchronization-free Regions. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Philadelphia, PA, USA) *(PLDI 2018)*. ACM, New York, NY, USA, 46–61. https://doi.org/10.1145/3192366.3192367

[18] Taeho Hwang, Jaemin Jung, and Youjip Won. 2015. HEAPO: Heap-Based Persistent Object Store. *ACM Trans. Storage* 11, 1, Article 3 (Dec. 2015), 21 pages. https://doi.org/10.1145/2629619

[19] Intel. 2015. Persistent Memory Programming. https://pmem.io/

[20] Intel. 2019. Intel 64 and IA-32 Architectures Software Developer's Manual (Combined Volumes). https://software.intel.com/sites/default/files/managed/39/c5/325462-sdm-vol-1-2abcd-3abcd.pdf Order Number: 325462-069US.

[21] Intel. 2019. Intel® Optane™ Persistent Memory. https://www.intel.com/content/www/us/en/architecture-and-technology/optane-dc-persistent-memory.html

[22] Intel. 2020. Intel® Optane™ Persistent Memory 200 Series Product Specifications. https://ark.intel.com/content/www/us/en/ark/products/series/203877/intel-optane-persistent-memory-200-series.html

[23] Joseph Izraelevitz, Hammurabi Mendes, and Michael L Scott. 2016. Linearizability of persistent memory objects under a full-system-crash failure model. In *International Symposium on Distributed Computing.* Springer, 313–327.

[24] Joseph Izraelevitz, Jian Yang, Lu Zhang, Juno Kim, Xiao Liu, Amirsaman Memaripour, Yun Joon Soh, Zixuan Wang, Yi Xu, Subramanya R.

Dulloor, Jishen Zhao, and Steven Swanson. 2019. Basic Performance Measurements of the Intel Optane DC Persistent Memory Module. arXiv:1903.05714 [cs.DC]

[25] Erik Jensen, G. W. Hagensen, and J. Broughton. 1987. A new approach to exclusive data access in shared memory multiprocessors.

[26] Arpit Joshi, Vijay Nagarajan, Marcelo Cintra, and Stratis Viglas. 2015. Efficient Persist Barriers for Multicores. In *Proceedings of the 48th International Symposium on Microarchitecture* (Waikiki, Hawaii) *(MICRO-48)*. ACM, New York, NY, USA, 660–671. https://doi.org/10.1145/2830772.2830805

[27] Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A promising semantics for relaxed-memory concurrency. In *POPL 2017*.

[28] Artem Khyzha and Ori Lahav. 2021. Taming X86-TSO Persistency. *Proc. ACM Program. Lang.* 5, POPL (2021). https://doi.org/10.1145/3434328

[29] Michalis Kokologiannakis, Ilya Kaysin, Azalea Raad, and Viktor Vafeiadis. 2021. PerSeVerE: Persistency Semantics for Verification under Ext4. *Proc. ACM Program. Lang.* 5, POPL, Article 43 (Jan. 2021), 29 pages. https://doi.org/10.1145/3434324

[30] Aasheesh Kolli, Vaibhav Gogte, Ali Saidi, Stephan Diestelhorst, Peter M. Chen, Satish Narayanasamy, and Thomas F. Wenisch. 2017. Language-level Persistency. In *Proceedings of the 44th Annual International Symposium on Computer Architecture* (Toronto, ON, Canada) *(ISCA '17)*. ACM, New York, NY, USA, 481–493. https://doi.org/10.1145/3079856.3080229

[31] Aasheesh Kolli, Steven Pelley, Ali Saidi, Peter M. Chen, and Thomas F. Wenisch. 2016. High-Performance Transactions for Persistent Memories. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems* (Atlanta, Georgia, USA) *(ASPLOS '16)*. Association for Computing Machinery, New York, NY, USA, 399–411. https://doi.org/10.1145/2872362.2872381

[32] Leslie Lamport. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM* 21, 7 (July 1978).

[33] Leslie Lamport. 1979. How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. *IEEE Trans. Computers* 28, 9 (Sept. 1979), 690–691. https://doi.org/10.1109/TC.1979.1675439

[34] Sung-Hwan Lee, Minki Cho, Anton Podkopaev, Soham Chakraborty, Chung-Kil Hur, Ori Lahav, and Viktor Vafeiadis. 2020. Promising 2.0: Global Optimizations in Relaxed Memory Concurrency *(PLDI 2020)*. https://doi.org/10.1145/3385412.3386010

[35] Mengxing Liu, Mingxing Zhang, Kang Chen, Xuehai Qian, Yongwei Wu, Weimin Zheng, and Jinglei Ren. 2017. DudeTM: Building Durable Transactions with Decoupling for Persistent Memory. In *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems* (Xi'an, China) *(ASPLOS '17)*. Association for Computing Machinery, New York, NY, USA, 329–343. https://doi.org/10.1145/3037697.3037714

[36] Sihang Liu, Yizhou Wei, Jishen Zhao, Aasheesh Kolli, and Samira Khan. 2019. PMTest: A Fast and Flexible Testing Framework for Persistent Memory Programs. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '19)*.

[37] Youyou Lu, Jiwu Shu, Youmin Chen, and Tao Li. 2017. Octopus: an RDMA-enabled Distributed Persistent Memory File System. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*. USENIX Association, Santa Clara, CA, 773–785. https://www.usenix.org/conference/atc17/technical-sessions/presentation/lu

[38] Youyou Lu, Jiwu Shu, and Long Sun. 2016. Blurred Persistence: Efficient Transactions in Persistent Memory. *ACM Trans. Storage* 12, 1, Article 3 (Jan. 2016), 29 pages. https://doi.org/10.1145/2851504

[39] Dominic P. Mulligan, Scott Owens, Kathryn E. Gray, Tom Ridge, and Peter Sewell. 2014. Lem: Reusable Engineering of Real-World Semantics *(ICFP '14)*. https://doi.org/10.1145/2628136.2628143

[40] Ismail Oukid, Daniel Booss, Adrien Lespinasse, Wolfgang Lehner, Thomas Willhalm, and Grégoire Gomes. 2017. Memory Management Techniques for Large-Scale Persistent-Main-Memory Systems. *Proc. VLDB Endow.* 10, 11 (Aug. 2017), 1166–1177. https://doi.org/10.14778/3137628.3137629

[41] Scott Owens, Susmit Sarkar, and Peter Sewell. 2009. A Better X86 Memory Model: X86-TSO. In *TPHOL*.

[42] Steven Pelley, Peter M. Chen, and Thomas F. Wenisch. 2014. Memory Persistency. In *Proceeding of the 41st Annual International Symposium on Computer Architecuture* (Minneapolis, Minnesota, USA) *(ISCA '14)*. IEEE Press, 265–276.

[43] Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. 2017. Simplifying ARM Concurrency: Multicopy-Atomic Axiomatic and Operational Models for ARMv8. 2, POPL (Dec. 2017).

[44] Christopher Pulte, Jean Pichon-Pharabod, Jeehoon Kang, Sung-Hwan Lee, and Chung-Kil Hur. 2019. Promising-ARM/RISC-V: A Simpler and Faster Operational Concurrency Model. In *PLDI 2019*.

[45] Azalea Raad, Ori Lahav, and Viktor Vafeiadis. 2020. Persistent Owicki-Gries Reasoning. *Proc. ACM Program. Lang.* 3, OOPSLA (2020).

[46] Azalea Raad and Viktor Vafeiadis. 2018. Persistence Semantics for Weak Memory: Integrating Epoch Persistency with the TSO Memory Model. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 137 (Oct. 2018), 27 pages. https://doi.org/10.1145/3276507

[47] Azalea Raad, John Wickerson, Gil Neiger, and Viktor Vafeiadis. 2020. Persistency Semantics of the Intel-X86 Architecture. *Proc. ACM Program. Lang.* 4, POPL (2020).

[48] Azalea Raad, John Wickerson, and Viktor Vafeiadis. 2019. Weak Persistency Semantics from the Ground up: Formalising the Persistency Semantics of ARMv8 and Transactional Models. *Proc. ACM Program.*

Lang. 3, OOPSLA (2019).

[49] Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. 2010. X86-TSO: A Rigorous and Usable Programmer's Model for X86 Multiprocessors. *Commun. ACM* 53, 7 (2010).

[50] Yizhou Shan, Shin-Yeh Tsai, and Yiying Zhang. 2017. Distributed Shared Persistent Memory. In *Proceedings of the 2017 Symposium on Cloud Computing* (Santa Clara, California) *(SoCC '17)*. Association for Computing Machinery, New York, NY, USA, 323–337. https://doi.org/10.1145/3127479.3128610

[51] Hongping Shu, Hongyu Chen, Hao Liu, Youyou Lu, Qingda Hu, and Jiwu Shu. 2018. Empirical Study of Transactional Management for Persistent Memory. 61–66. https://doi.org/10.1109/NVMSA.2018.00015

[52] Arash Tavakkol, Aasheesh Kolli, Stanko Novakovic, Kaveh Razavi, Juan Gómez-Luna, Hasan Hassan, Claude Barthels, Yaohua Wang, Mohammad Sadrosadati, Saugata Ghose, Ankit Singla, Pratap Subrahmanyam, and Onur Mutlu. 2018. Enabling Efficient RDMA-based Synchronous Mirroring of Persistent Memory Transactions. *CoRR* abs/1810.09360 (2018). arXiv:1810.09360 http://arxiv.org/abs/1810.09360

[53] Haris Volos, Andres Jaan Tack, and Michael M. Swift. 2011. Mnemosyne: Lightweight Persistent Memory. In *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems* (Newport Beach, California, USA) *(ASPLOS XVI)*. Association for Computing Machinery, New York, NY, USA, 91–104. https://doi.org/10.1145/1950365.1950379

[54] Yoav Zuriel, Michal Friedman, Gali Sheffi, Nachshon Cohen, and Erez Petrank. 2019. Efficient Lock-Free Durable Sets. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 128 (Oct. 2019), 26 pages. https://doi.org/10.1145/3360554