








View-Based Owicki–Gries Reasoning for Persistent x86-TSO^{*}

Eleni Vafeiadi Bila¹ , Brijesh Dongol¹  (✉), Ori Lahav² , Azalea Raad³ ,
and John Wickerson³ 

¹ University of Surrey, Guildford, UK b.dongol@surrey.ac.uk

² Tel Aviv University, Tel Aviv, Israel

³ Imperial College London, London, UK

Abstract. The rise of persistent memory is disrupting computing to its core. Our work aims to help programmers navigate this brave new world by providing a program logic for reasoning about x86 code that uses low-level operations such as memory accesses and fences, as well as persistency primitives such as flushes. Our logic, PIEROGI, benefits from a simple underlying operational semantics based on *views*, is able to handle *optimised* flush operations, and is mechanised in the Isabelle/HOL proof assistant. We detail the proof rules of PIEROGI and prove them sound. We also show how PIEROGI can be used to reason about a range of challenging single- and multi-threaded persistent programs.

Keywords: Persistent memory, x86-TSO, Owicki-Gries, Isabelle/HOL, verification

1 Introduction

In our era of big data, the long-established boundary between ‘memory’ and ‘storage’ is increasingly blurred. Persistent memory is a technology that sits in both camps, promising both the durability of disks and data access times similar to those of DRAM. Embracing this technology requires rethinking our decades-old programming paradigms. As data held in memory is no longer wiped after a system restart, there is an opportunity to write *persistent* programs – programs that can recover their progress and continue computing even after a crash.

However, writing persistent programs is extremely challenging, as it requires the programmer to keep track of which memory writes have become persistent,

^{*} Vafeiadi Bila is supported by VeTSS. Dongol is supported by EPSRC grants EP/V038915/1, EP/R032556/1, EP/R025134/2 and ARC Discovery Grant DP190102142. Lahav is supported by the Israel Science Foundation (grant 1566/18), by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 851811), and by the Alon Young Faculty Fellowship. Raad is supported by a UKRI Future Leaders Fellowship [grant number MR/V024299/1]. Wickerson is supported by an EPSRC Programme Grant (EP/R006865/1).

and which have not. This is further complicated in a multi-threaded setting by the intricate interplay between the rules of memory *persistence* (which determine the order in which writes become persistent) and those of memory *consistency* (which determine what data can be observed by which threads).

To address this difficulty, we provide a foundation for persistent programming. We develop a program logic, PIEROGI, for reasoning about x86 code that uses low-level operations such as memory accesses and fences, as well as persistence primitives such as flushes. We demonstrate the utility of PIEROGI by using it to reason about a range of challenging single- and multi-threaded persistent programs, including some that demonstrate the subtle interplay between optimised flush (**flush_{opt}**) and store fence (**sfence**) instructions. Using the Isabelle/HOL proof assistant, we have mechanised the PIEROGI rules and proved them sound with respect to an operational semantics for x86 persistence [9]. One benefit of our Isabelle/HOL formalisation is that PIEROGI is already partially automated: once the user has produced a proof outline (i.e. annotated each instruction with a postcondition), they can simply use Isabelle/HOL’s *sledgehammer*, which automatically decides which axioms and rules of the proof system need invoking to verify the whole program. Our mechanisation, which includes all the example programs discussed in this paper, is available as auxiliary material [4, 5].

State of the art To our knowledge, the only program logic for persistent programs is POG (Persistent Owicki–Gries) [31]. As with PIEROGI, POG enables reasoning about persistent x86 programs and is based on the Owicki–Gries method [30]. However, unlike PIEROGI, POG is not mechanised in a proof assistant, and does not support optimised flush (**flush_{opt}**) instructions. Optimised flush instructions are an important persistence primitive as they are considerably faster than ordinary flush instructions. Indeed, Intel’s experiments on their Skylake microarchitecture indicate that they can be *nine times* faster when applied to buffers that hold tens of kilobytes of data [19, p. 289], and hence programmers are impelled, “If **flush_{opt}** is available, use **flush_{opt}** over **flush**.” However, **flush_{opt}** is a tricky instruction for programmers and program logic designers alike: compared to **flush**, **flush_{opt}** can be reordered with more instructions under x86.

PIEROGI can reason efficiently about x86 persistence (including **flush_{opt}** instructions) thanks to two key recent advances: 1) Px86_{view} [9], the view-based operational semantics of x86 persistence; and 2) the C11 Owicki–Gries logic [11–13] to reason about view-based operational semantics, which we adapt to Px86_{view}.

Our contributions 1) We present a program logic, called PIEROGI, for reasoning about persistent x86 programs. 2) We mechanise (and partially automate) PIEROGI in Isabelle/HOL, and prove it sound relative to an established operational semantics for x86 persistence. 3) We demonstrate the utility of PIEROGI by using it to verify several idiomatic persistent x86 programs.

Outline We begin with an overview of memory consistency and persistence in x86 and provide an example-driven account of PIEROGI reasoning (§2). We describe the assertion language and proof rules of PIEROGI in §3, and verify a selection of programs using PIEROGI in §4. We present the view-based operational semantics of x86 persistence and prove the soundness of PIEROGI in §5.

Auxiliary material Additional examples as well as the proofs of theorems stated in the paper are given in the accompanying technical appendix [5]. Our Isabelle/HOL mechanisation is available as auxiliary material [4].

2 Overview and Motivation

Recent operational models for weak memory use *views* to capture relaxed behaviours of concurrent programs [9, 11, 21, 22], where the memory records the entire history of writes that have taken place thus far. This way, different threads can have different subsets of these writes (i.e. different *views*) visible to them. Below, we review $\text{Px86}_{\text{view}}$, a view-based operational semantics for x86 persistency (§2.1); we then describe PIEROGI (§2.2) using a series of running examples.

2.1 $\text{Px86}_{\text{view}}$ at a Glance

In the literature of concurrency semantics, *consistency* models describe the permitted behaviours of programs by constraining the volatile memory order, i.e. the order in which memory writes are made visible to other threads, while *persistency* models describe the permitted behaviours of programs upon recovering from a crash (e.g. a power failure) by defining the persistent memory order, i.e. the order in which writes are committed to persistent memory. To distinguish between the two, memory *stores* are differentiated from memory *persists*: the former denotes the process of making a write visible to other threads, whilst the latter denotes the process of committing writes to persistent memory (durably).

$\text{Px86}_{\text{view}}$ Consistency The consistency semantics of $\text{Px86}_{\text{view}}$ is that of the well-known TSO (total store ordering) [36] model, where later (in program order) reads can be reordered before earlier writes on different locations. This is illustrated in the *store buffering* (SB) example below (left):

$$\begin{array}{l} \mathbf{store} \ x \ 1; \\ a := \mathbf{load} \ y \end{array} \parallel \begin{array}{l} \mathbf{store} \ y \ 1; \\ b := \mathbf{load} \ x \end{array} \quad (\text{SB}) \qquad \begin{array}{l} \mathbf{store} \ x \ 42; \\ \mathbf{store} \ y \ 7 \end{array} \parallel \begin{array}{l} a := \mathbf{load} \ y; \\ b := \mathbf{load} \ x \end{array} \quad (\text{MP})$$

$$a = 0 \wedge b = 0 : \checkmark \qquad \qquad \qquad a = 7 \wedge b = 0 : \times$$

Specifically, assuming $x=y=0$ initially, since $a := \mathbf{load} \ y$ (resp. $b := \mathbf{load} \ x$) can be reordered before $\mathbf{store} \ x \ 1$ (resp. $\mathbf{store} \ y \ 1$), it is possible to observe the weak behaviour $a=0 \wedge b=0$. A well-known way of modelling such reorderings in TSO is through *store buffers*: when a thread τ executes a write $\mathbf{store} \ x \ v$, its effects are not immediately made visible to other threads; rather they are delayed in a thread-local (store) buffer only visible to τ , and propagated to the memory at a later time, whereby they become visible to other threads. For instance, when $\mathbf{store} \ x \ 1$ and $\mathbf{store} \ y \ 1$ are delayed in the respective thread buffers (and thus not visible to one another), then $a := \mathbf{load} \ y$ and $b := \mathbf{load} \ x$ may both read 0.

Cho et al. [9] capture this by associating each thread τ with a *coherence view* (also called a thread-observable view), describing the writes observable by τ . Distinct threads may have different coherence views. For instance, after executing $\mathbf{store} \ x \ 1$ and $\mathbf{store} \ y \ 1$, the coherence view of the left thread may include

store x 1 and *not* **store** y 1, while that of the right may include **store** y 1 and *not* **store** x 1. This way, $a := \mathbf{load}$ y (resp. $b := \mathbf{load}$ x) may read the initial value 0, as its coherence view does not include **store** y 1 (resp. **store** x 1).

After SC (sequential consistency) [27], TSO is one of the strongest consistency models and supports synchronisation patterns such as *message passing*, as shown in MP above, where $a = 7 \wedge b = 0$ cannot be observed. Specifically, (assuming $x=y=0$ initially) if the right thread reads 7 from y (written by the left thread), then the left thread passes a message to the right. Under TSO, message passing ensures that the instruction writing the message and all those ordered before it (e.g. **store** x 42; **store** y 7) are executed (ordered) before the instruction reading it (e.g. $a := \mathbf{load}$ y). As such, since $b := \mathbf{load}$ x is executed after $a := \mathbf{load}$ y , if $a=7$ (i.e. **store** x 42 is executed before $a := \mathbf{load}$ y), then $b=42$.

Px86_{view} Persistency Cho et al. [9] recently developed the Px86_{view} model, a view-based description of the Intel-x86 persistency semantics, which follows a *buffered, relaxed* persistency model. Under a buffered model, memory persists occur *asynchronously* [10]: they are buffered in a queue to be committed to persistent memory at a future time. This way, persists occur after their corresponding stores and as prescribed by the persistency semantics, while allowing the execution to proceed ahead of persists. As such, after recovering from a crash, only a *prefix* of the persistent memory order may have persisted. (The alternative is *unbuffered* persistency in which stores and persists happen simultaneously.)

Under relaxed persistency, the volatile and persistent memory orders may disagree: the order in which the writes are made visible to other threads may differ from the order in which they are persisted. (The alternative is *strict* persistency in which the volatile and persistent memory orders coincide.)

The relaxed and buffered persistency of Px86_{view} is shown in Fig. 1a. If a crash occurs during (or after) the execution of Fig. 1a, at crash time either write may have persisted and thus $x, y \in \{0, 1\}$ upon recovery. Note that the two writes cannot be reordered under Intel-x86 (TSO) consistency and thus at no point during the normal (non-crashing) execution of Fig. 1a is $x=0, y=1$ observable. Nevertheless, in case of a crash it is possible to observe $x=0, y=1$ after recovery. That is, due to the relaxed persistency of Px86_{view}, the store order (x before y) is separate from the persist order (y before x). More concretely, under Px86_{view} the writes may persist 1) in any order, when they are on distinct locations; or 2) in the volatile memory order, when they are on the same location.⁴

To afford more control over when pending writes are persisted, Intel-x86 provides explicit *persist* instructions such as **flush** x and **flush_{opt}** x that can be used to persist the pending writes on x .⁵ This is illustrated in Fig. 1b: executing **flush** x persists the earlier write on x (i.e. **store** x 1) to memory. As such, if

⁴ Given a *cache line* (a set of locations), writes on distinct cache lines may persist in any order, while writes on the same cache line persist in the volatile memory order. For brevity, we assume that each cache line contains a single location, thus forgoing the need for cache lines. However, it is straightforward to lift this assumption.

⁵ Executing **flush** x or **flush_{opt}** x persists the pending writes on *all locations in the cache line of x* . However, as discussed, we assume cache lines contain single locations.

store x 1; store y 1 (a)	store x 1; flush x ; store y 1 (b)	store x 1; flush_{opt} x ; store y 1 (c)	store x 1; flush_{opt} x ; sfence ; store y 1 (d)	store x 1; \parallel $a := \text{load } y$; flush x ; \parallel if ($a=1$) store y 1 \parallel store z 1 (e)
$\zeta : x, y \in \{0, 1\}$	$\zeta : y=1 \Rightarrow x=1$	$\zeta : x, y \in \{0, 1\}$	$\zeta : y=1 \Rightarrow x=1$	$\zeta : z=1 \Rightarrow x=1$

Fig. 1: Example Px86_{view} programs and possible values after recovery from a crash (ζ). In all examples x, y, z are distinct locations in persistent memory such that $x=y=z=0$ initially, and a is a (thread-local) register.

the execution of Fig. 1b crashes and upon recovery $y=1$, then $x=1$. That is, if **store** y 1 has executed and persisted before the crash, then so must the earlier **store** x 1; **flush** x . Note that $y=1 \Rightarrow x=1$ describes a *crash invariant*, in that it holds upon crash recovery *regardless* of when (i.e. at which program point) the crash may have occurred. Observe that this crash invariant is guaranteed thanks to the ordering constraints on **flush** instructions. Specifically, **flush** instructions are ordered with respect to all writes; as such, **flush** x in Fig. 1b cannot be reordered with respect to either write, and thus upon recovery $y=1 \Rightarrow x=1$.

However, instruction reordering means that persist instructions may not execute at the intended program point and thus not guarantee the intended persist ordering. Specifically, **flush_{opt}** x is only ordered with respect to earlier writes on x , and may be reordered with respect to later writes, as well as earlier writes on different locations. This is illustrated in Fig. 1c: **flush_{opt}** x is not ordered with respect to **store** y 1 and may be reordered after it. Therefore, if a crash occurs after **store** y 1 has executed and persisted but before **flush_{opt}** x has executed, then it is possible to observe $y=1, x=0$ on recovery. That is, there is no guarantee that **store** x 1 persists before **store** y 1, *despite* the intervening **flush_{opt}** x .

In order to prevent such reorderings and to strengthen the ordering constraints between **flush_{opt}** and later instructions, one can use either *fence* instructions, namely **sfence** (store fence) and **mfence** (memory fence), or atomic *read-modify-write* (RMW) instructions such as compare-and-set (**CAS**) and fetch-and-add (**FAA**). More concretely, **sfence**, **mfence** and RMW instructions are ordered with respect to all (both earlier and later) **flush_{opt}**, **flush** and write instructions, and can be used to prevent reorderings such as that in Fig. 1c. This is illustrated in Fig. 1d. Unlike in Fig. 1c, the intervening **sfence** ensures that **flush_{opt}** in Fig. 1d is ordered with respect to **store** y 1 and cannot be reordered after it, ensuring that **store** x 1 persists before **store** y 1 (i.e. $y=1 \Rightarrow x=1$ upon recovery), as in Fig. 1b. Note that replacing **sfence** in Fig. 1d with **mfence** or an RMW yields the same result. Alternatively, one can think of **flush_{opt}** x executing *asynchronously*, in that its effect (persisting x) does not take place immediately upon execution, but rather at a later time. However, upon executing a barrier instruction (i.e. **mfence**, **sfence** or an RMW), execution is blocked until the effect of earlier **flush_{opt}** instructions take place; that is, executing such barrier instructions ensures that earlier **flush_{opt}** behave *synchronously* (like **flush**).

$$\begin{array}{l}
P : \{a = b = 0 \wedge \forall \tau \in \{1, 2\}. [x]_\tau = [y]_\tau = \{0\}\} \\
P_1 : \{7 \notin [y]_2 \wedge a = 0\} \\
\text{store } x \ 42; // \text{SP}_1, \text{Cons} \\
P_2 : \{[x]_1 = \{42\} \wedge 7 \notin [y]_2\} \\
\text{store } y \ 7; // \text{SP}_1, \text{Cons} \\
P_3 : \{\text{true}\} \\
Q : \{a = 7 \Rightarrow b = 42\} \\
Q_1 : \{[y]_2 \subseteq \{0, 7\} \wedge (7 \in [y]_2 \Rightarrow \langle y, 7 \rangle [x]_2 = \{42\})\} \\
a := \text{load } y; // \text{LP}_2 \\
Q_2 : \{a \in \{0, 7\} \wedge (a = 7 \Rightarrow [x]_2 = \{42\})\} \\
b := \text{load } x; // \text{LP}_1, \text{Cons} \\
Q_3 : \{a = 7 \Rightarrow b = 42\}
\end{array}$$

Fig. 2: A PIEROGI proof sketch of message passing (MP), where the // annotation at each step identifies the PIEROGI proof rule (in §3.4) applied, and the highlighted assertions capture the effects of the preceding instruction.

The example in Fig. 1e illustrates how message passing can impose persist orderings on the writes of *different* threads. (Note that the program in the left thread of Fig. 1e is that of Fig. 1b.) As in MP, if $a = 1$, then **store** $x \ 1$; **flush** x is executed before $a := \text{load } y$ (thanks to message passing). Consequently, since **store** $z \ 1$ is executed after $a := \text{load } y$ when $a = 1$, we know **store** $x \ 1$; **flush** x is executed before **store** $z \ 1$. Therefore, if upon recovery $z=1$ (i.e. **store** $z \ 1$ has persisted before the crash), then $x=1$ (**store** $x \ 1$; **flush** x must have also persisted before the crash). As before, replacing **flush** x in Fig. 1e with **flush**_{opt} x ; C yields the same result upon recovery when C is an **sfence**/**mfence** or an RMW.

2.2 PIEROGI: View-Based Owicki–Gries Reasoning for P_{x86}_{view}

Sequential Reasoning about Consistency using Views In Fig. 2 we present a PIEROGI proof sketch of MP. Recall that in order to account for possible write-read reorderings on Intel-x86 architectures, P_{x86}_{view} associates each thread τ with a coherence view, describing the writes visible to τ . To reason about such thread-observable views, PIEROGI supports assertions of the form $[x]_\tau = S$, stating that τ may read any value in the set S for location x . That is, the coherence view of τ for x consists of the writes whose values are those in S .

In the remainder of this article we enumerate the threads in our examples from left to right; e.g. the left and right threads in Fig. 2 are identified as 1 and 2, respectively. Moreover, we assume the registers of distinct threads have distinct names. The precondition P in Fig. 2 thus states that both threads may initially only read 0 for both x and y : $\forall \tau \in \{1, 2\}. [x]_\tau = [y]_\tau = \{0\}$.

In the case of thread 1, we can weaken P (using the standard rule of consequence of Hoare logic – see **Cons** in §3) to obtain P_1 . Upon executing **store** $x \ 42$ (1) we weaken the resulting assertion by dropping the $a = 0$ conjunct; and (2) we update the observable view of thread 1 on x to reflect the new value of x : $[x]_1 = \{42\}$; that is, after executing **store** $x \ 42$, the only value observable by thread 1 for x is 42. Similarly, after executing **store** $y \ 7$, we could assert $[y]_1 = \{7\}$; however, this is not necessary for establishing the final postcondition Q , and we thus simply weaken the postcondition to **true** (P_3).

<pre> {[y]^P = {0}} store x 1; // SP₁ {[x]₁ = {1} ∧ [y]^P = {0}} flush x; // FP₁ {[x]₁ = {1} ∧ [x]^P = {1} ∧ [y]^P = {0}} store y 1; // SP₁ {[x]₁ = {1} ∧ [x]^P = {1} ∧ [y]₁ = {1}} {!z : [y]^P = {1} ⇒ [x]^P = {1}} </pre>	<pre> {[y]^P = {0}} store x 1; // SP₁ {[x]₁ = {1} ∧ [y]^P = {0}} flush_{opt} x; // OP₁ {[x]₁ = {1} ∧ [x]₁^A = {1} ∧ [y]^P = {0}} sfence; // SFP₁ {[x]₁ = {1} ∧ [x]^P = {1} ∧ [y]^P = {0}} store y 1; // SP₁ {[x]₁ = {1} ∧ [x]^P = {1} ∧ [y]₁ = {1}} {!z : [y]^P = {1} ⇒ [x]^P = {1}} </pre>
--	---

Fig. 3: Proof sketches of Fig. 1b (left) and Fig. 1d (right)

Analogously, in the case of thread 2 we weaken P to obtain Q_1 : $[y]_2 = \{0\}$ implies $[y]_2 \subseteq \{0, 7\}$ and $7 \in [y]_2 \Rightarrow \langle y, 7 \rangle [x]_2 = \{42\}$. Note that $7 \in [y]_2 \Rightarrow \langle y, 7 \rangle [x]_2 = \{42\}$ yields a vacuously true implication as $[y]_2 = \{0\}$ and thus $7 \notin [y]_2$. The $\langle y, 7 \rangle [x]_2$ denotes a *conditional view assertion* [11] that describes how reading a value on one location (y) affects the thread-observable view on a different location (x). More concretely, $\langle y, 7 \rangle [x]_2 = \{42\}$ states that if thread 2 executes a load on y and reads value 7, it subsequently may only observe value 42 for x . This is indeed the essence of message passing in MP: once thread 2 reads 7 from y , it may only read 42 for x thereafter. As such, after executing the read instruction $a := \mathbf{load} \ y$ (1) we apply the LP₁ rule (in Fig. 7) which simply replaces $[y]_2$ with the local register a in which the value of y is read; and (2) we replace the conditional assertion $\langle y, 7 \rangle [x]_2 = \{42\}$ with the implication $a = 7 \Rightarrow [x]_2 = \{42\}$, stating that if the value read by thread 2 for y (in a) is 7, then its observable view for x is $\{42\}$. Similarly, upon executing $b := \mathbf{load} \ x$ we simply apply LP₁ to replace $[x]_2$ with the local register b in which the value of x is read. Lastly, the final postcondition Q is given by the conjunction of the thread-local postconditions ($P_3 \wedge Q_3$).

Concurrent Reasoning and Stability In our description of the PIEROGI proof sketch in Fig. 2 thus far we focused on *sequential* (per-thread) reasoning, ignoring how concurrent threads may affect the validity of assertions at each program point. Specifically, as in existing concurrent logics [11, 26, 30, 31], we must ensure that the assertions at each program point are *stable* under concurrent operations. For instance, to ensure that P_1 remains stable under the concurrent operation $a := \mathbf{load} \ y$, we require that executing $a := \mathbf{load} \ y$ on states satisfying the conjunction of P_1 and the precondition of $a := \mathbf{load} \ y$ (i.e. Q_1) not invalidate P_1 , in that the resulting states continue to satisfy P_1 ; that is, $\{P_1 \wedge Q_1\} a := \mathbf{load} \ y \{P_1\}$ holds. Similarly, we must ensure that P_1 is stable under $b := \mathbf{load} \ x$, i.e. $\{P_1 \wedge Q_2\} b := \mathbf{load} \ x \{P_1\}$ holds. Analogously, we must establish the stability of P_2 , P_3 , Q_1 , Q_2 and Q_3 under concurrent operations. In §3 we present syntactic rules that simplify the task of checking stability obligations. It is then straightforward to show that the assertions in Fig. 2 are stable.

Reasoning about flush Persistency To reason about the relaxed, buffered persistency of $\text{Px86}_{\text{view}}$, Cho et al. [9] introduce *persistency views*, determining the possible *persisted* values for each location; i.e. the values of those writes that may have persisted to memory. Note that the persistency view determines the possible values observable upon recovery from a crash. By contrast, the (per-thread) coherence views determine the observable values during normal (non-crashing) executions, and have no bearing on the post-crash values.

Analogously, we extend PIEROGI with assertions of the form $[x]^P = S$, stating that the persistent view for x includes writes whose values are given by S . To see this, consider the PIEROGI proof sketch of Fig. 1b in Fig. 3 (left). Initially, y holds 0 in persistent memory: $[y]^P = \{0\}$. (Note that the precondition could additionally include $[x]_1 = [y]_1 = \{0\} \wedge [x]^P = \{0\}$ to denote that initially the thread may only observe 0 for x and y and that x holds 0 in persistent memory; however, this is not needed for the proof and we thus forgo it.)

As before, after executing **store** x 1, the observable value for x is updated, as denoted by $[x]_1 = \{1\}$. Moreover, after executing **flush** x , the persisted value for x is 1, as denoted by $[x]^P = \{1\}$, by committing (persisting) the observable value for x (i.e., $[x]_1 = \{1\}$) to memory (see FP_1 in Fig. 7). Finally, after executing **store** y 1, the observable value for y is updated, as denoted by $[y]_1 = \{1\}$.

Crash Invariants Recall that $\not\downarrow: y=1 \Rightarrow x=1$ in Fig. 1b denotes a *crash invariant* in that it describes the persistent memory upon recover from a crash at *any* program point. This is because we have no control over when a crash may occur. To capture such invariants, in PIEROGI we write *quadruples* of the form $\{P\} C \{Q\} \{\{\not\downarrow : I\}\}$, where $\{P\} C \{Q\}$ denotes a Hoare triple and I denotes the crash invariant. If C is a sequential program, I must follow from *every* assertion (including P and Q) in the proof. For instance, in the proof outline of Fig. 3 (left) all four assertions imply the invariant $[y]^P = \{1\} \Rightarrow [x]^P = \{1\}$. We discuss the meaning of crash invariants for concurrent programs below.

Reasoning about flush_{opt} Persistency Recall that unlike **flush**, **flush_{opt}** instructions (due to instruction reordering) may behave asynchronously and their effects may not take place immediately after execution. As such, unlike for **flush** x , after executing **flush_{opt}** x we cannot simply copy the observable view on x to the persistent view on x .

To capture the asynchronous nature of **flush_{opt}**, Cho et al. [9] introduce yet another set of views, namely the *thread-local asynchronous view*: the asynchronous view of thread τ on x describes the values (writes) that will be persisted at a later time (asynchronously) by τ upon executing a barrier instruction. That is, 1) when thread τ executes **flush_{opt}** x , its asynchronous view of x is advanced to at least its observable view of x ; and 2) when τ executes a barrier (**sfence**, **mfence** or RMW), then its persistent view for each location is advanced to at least its corresponding asynchronous view. We model this in PIEROGI by 1) setting $[x]_\tau^A$ to be a subset of $[x]_\tau$ when **flush_{opt}** x is executed; and 2) setting $[x]^P$ to be a subset of $[x]_\tau^A$ (for each location x) when a barrier is executed.

This is illustrated in the proof sketch of Fig. 1d in Fig. 3 (right). In particular, unlike the proof sketch of Fig. 1b in Fig. 3 (left), after executing **flush_{opt}** x we

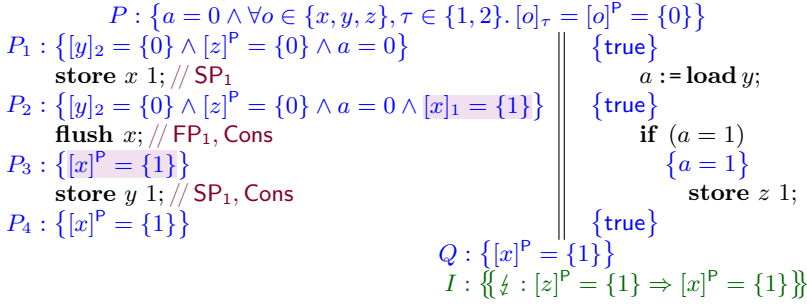


Fig. 4: A PIEROGI proof sketch of Fig. 1e

cannot simply copy the thread-observable view to the persistent view. Rather, we copy the thread-observable view $[x]_1$ to its asynchronous view and assert $[x]_1^A = \{1\}$; and upon executing the subsequent **sfence**, we copy the thread-asynchronous view to the persistent view and assert $[x]^P = \{1\}$.

Putting It All Together We next present a PIEROGI proof sketch of Fig. 1e in Fig. 4. The proof of the left thread is analogous to that in Fig. 3 (left); the proof of the right thread is straightforward and applies standard reasoning principles. The final postcondition Q is obtained by weakening the conjunction of per-thread postconditions.

Note that the crash invariant I follows from the assertions at each program point of thread 1 (i.e. $P_1 \vee P_2 \vee P_3 \vee P_4 \Rightarrow I$). That is, the crash invariant must follow from the assertions at *all* program points of *some* thread (e.g. thread 1 in Fig. 4). In the case of sequential programs (e.g. in Fig. 3), this amounts to all program points (of the only executing thread). Intuitively, we must ensure that the crash invariant holds at every program point regardless of how the underlying state changes. As the assertions are stable under concurrent operations, it is thus sufficient to ensure that there exists some thread whose assertions at each program point imply the crash invariant.

3 The PIEROGI Proof rules and Reasoning Principles

We proceed with a description of our verification framework. As with prior work [11], the view-based semantics for persistent TSO [9] allows us to use the standard Owicki–Gries rules [2, 30] for compound statements. The main adjustment is the introduction of a new specialised assertion language capable of expressing properties about the different “views” described intuitively in §2. As such, since view updates are highly non-deterministic, the standard “assignment axiom” of Hoare Logic (and by extension Owicki–Gries) is no longer applicable. Moreover, unlike SC, reads in a weak memory setting have a side-effect: their interaction with the memory location being read causes the view of the executing

$$\begin{array}{l}
v, u \in \text{VAL} \triangleq \mathbb{N} \quad x, y, \dots \in \text{LOC} \quad a, b, \dots \in \text{REG} \quad \tau \in \text{TID} \triangleq \mathbb{N} \quad i, j, k, \dots \in \text{LAB} \\
\hat{a}, \hat{b}, \dots \in \text{AUXVAR} \quad \hat{e} \in \text{AUXEXP} ::= v \mid \hat{a} \mid \hat{e} + \hat{e} \mid \dots \\
e \in \text{EXP} ::= v \mid a \mid e + e \mid \dots \quad B \in \text{BEXP} ::= \text{true} \mid B \wedge B \mid \dots \\
\alpha \in \text{AST} ::= \text{skip} \mid a := e \mid a := \text{load } x \mid \text{store } x \ e \\
\quad \mid a := \text{CAS } x \ e_1 \ e_2 \mid \text{sfence} \mid \text{mfence} \mid \text{flush } x \mid \text{flush}_{\text{opt}} \ x \\
ls \in \text{LST} ::= \alpha \ \mathbf{goto} \ j \mid \mathbf{if} \ B \ \mathbf{goto} \ j \ \mathbf{else} \ \mathbf{to} \ k \mid \langle \alpha \ \mathbf{goto} \ j, \hat{a} := \hat{e} \rangle \\
\Pi \in \text{PROG} \triangleq \text{TID} \times \text{LAB} \rightarrow \text{LST} \quad \vec{pc} \in \text{PC} \triangleq \text{TID} \rightarrow \text{LAB}
\end{array}$$

Fig. 5: The PIEROGI domains and programming language

thread to advance. Therefore, we resort to a set of proof rules that describe how views are modified and manipulated, as formalised by our view-based assertions.

3.1 The PIEROGI Programming Language

We present the programming language in Fig. 5. *Atomic statements* (in AST) comprise **skip**, assignment, memory reads and writes, barrier instructions and explicit persists. Specifically, $a := e$ evaluates expression e and returns it in (thread-local) register a ; $a := \text{load } x$ reads from memory location x and returns it in register a ; and **store** $x \ e$ writes the evaluated value of e into location x . The $a := \text{CAS } x \ e_1 \ e_2$ denotes ‘compare-and-set’ on location x , from the evaluated value of e_1 to the evaluated value of e_2 , and sets a to 1 if the CAS succeeds and to 0, otherwise. Finally, **mfence** denotes a memory fence, **sfence** denotes a store fence, and **flush** x and **flush**_{opt} x denote explicit persist instructions (see §2).

Formally, we model a program Π as a function mapping each pair (τ, i) of thread identifier and label to the *labelled statement* (in LST) to be executed. A labelled statement may be 1) a plain statement of the form $\alpha \ \mathbf{goto} \ j$, comprising an atomic statement α to be executed and the label j of the next statement; 2) a conditional statement of the form **if** $B \ \mathbf{goto} \ j \ \mathbf{else} \ \mathbf{to} \ k$ to accommodate branching, which proceeds to label j if B holds and to k , otherwise; or 3) a statement with an auxiliary update $\langle \alpha \ \mathbf{goto} \ j, \hat{a} := \hat{e} \rangle$, which behaves as $\alpha \ \mathbf{goto} \ j$, but in addition (in the same atomic step) updates the value of the auxiliary variable \hat{a} with the auxiliary expression \hat{e} . It is well known that Owicki–Gries proofs require auxiliary variables to record the history of executions to differentiate states that would otherwise not be distinguishable [30]. We show how auxiliary variables are used in PIEROGI in the flush buffering example (§4).

We track the control flow within each thread via the *program counter function*, \vec{pc} , recording the program counter of each thread. We assume a designated label, $\iota \in \text{LAB}$, representing the *initial label*; i.e. each thread begins execution with $\vec{pc}(\tau) = \iota$. Similarly, $\zeta \in \text{LAB}$ represents the *final label*. Moreover, if $\vec{pc}(\tau) = i$ at the current execution step, then: 1) when $\Pi(\tau, i) = \alpha \ \mathbf{goto} \ j$ or $\Pi(\tau, i) = \langle \alpha \ \mathbf{goto} \ j, \hat{a} := \hat{e} \rangle$, then $\vec{pc}(\tau) = j$ at the next step; 2) when $\Pi(\tau, i) = \mathbf{if} \ B \ \mathbf{goto} \ j \ \mathbf{else} \ \mathbf{to} \ k$ at the current step, then if B holds in the current state, then $\vec{pc}(\tau) = j$ at the next step; otherwise $\vec{pc}(\tau) = k$ at the next step.

Example 1. The program in Fig. 4, assuming that the left thread has id 1, is given as follows. The formalisation of the right thread is omitted, but is similar.

$$\Pi \triangleq \left\{ \begin{array}{l} (1, \iota) \mapsto \mathbf{store} \ x \ 1 \ \mathbf{goto} \ 2, (1, 2) \mapsto \mathbf{flush} \ x \ \mathbf{goto} \ 3, \\ (1, 3) \mapsto \mathbf{store} \ y \ 1 \ \mathbf{goto} \ \zeta, \dots \end{array} \right\}$$

3.2 View-Based Expressions

As with prior work on the RC11 model [21], we interpret PIEROGI expressions directly over a view-based state. We use expressions tailored for the view-based P_{x86}_{view} model [9], which allow us to express relationships between different system components, including the persistent memory.

Our expressions fall into one of four categories: 1) *current view* expressions, which describe the current views of different system components (e.g. the persistent view); 2) *conditional view* expressions [11], which describe a view on a location after reading a particular value on a *different* location; 3) *last view* expressions, which hold if a component is viewing the last write to a location; and 4) *write-count* expressions, which describe the number of writes to a location.

Our current view expressions comprise $[x]_\tau$, $[x]^P$ and $[x]^A$, as described below; as shown in §2, each of these expressions describes a *set* of possible values.

$[x]_\tau$ denotes the *coherence view* of thread τ : the set of values τ may read for x . $[x]^P$ denotes the *persistent memory view*: the set of values that x may hold in (persistent) memory.

$[x]^A$ denotes the *asynchronous memory view* of thread τ : the set of values that can be persisted after a barrier instruction (**sfence**/**mfence**/RMW) is executed by τ (see rule OP in Fig. 7). Asynchronous views are updated after executing a **flush**_{opt}; however, unlike persistent memory views, the values in asynchronous views are not guaranteed to be persisted until a subsequent barrier is executed by the same thread.

Conditional view expressions are of the form $\langle x, v \rangle [y]_\tau$, as described below. As discussed in §2, conditional expressions capture the crux of message passing.

$\langle x, v \rangle [y]_\tau$ returns a set of values that τ may read for y after it reads value v for x . In particular, if $\langle x, v \rangle [y]_\tau = S$ holds for some set S and τ executes $a := \mathbf{load} \ x$, then in the state immediately after the load, if $a = v$, then $[y]_\tau \subseteq S$ (see LP₂ in Fig. 7).

Last-view expressions (*cf.* [16]) are boolean-valued and hold if a particular component is synchronised (i.e. observes the latest value) on the given location. Such expressions provide determinism guarantees on **load** and **flush**. For instance if the view of τ is the last write on x , then a read from x by τ will load this last value. Last-view expressions comprise $\llbracket x \rrbracket_\tau$ and $\llbracket x \rrbracket_\tau^E$:

$\llbracket x \rrbracket_\tau$ holds iff τ is currently viewing the *last* write to x . Thus, for example, if $\llbracket x \rrbracket_\tau$ holds, then a **load** from x by τ reads the last write to x . Note that unlike architectural operational models [36], in the view model [9], writes are visible to all threads as soon as they occur.

$\llbracket x \rrbracket_\tau^F$ holds iff a **flush** of x by τ is guaranteed to flush the *last* write to x to persistent memory.

Lastly, write-count expressions are of the form $|x, v|$, as described below. Such assertions are useful for inferring view expressions from known facts about the number of writes in the system with a particular value (see Fig. 11).

$|x, v|$ returns the number of writes to x with value v . If $|x, v|$ holds and τ writes to $y \neq x$, or writes a value $u \neq v$, then $|x, v|$ continues to hold afterwards.

3.3 Owicki–Gries Reasoning

We present the PIEROGI proof system, as an extension of Hoare Logic with Owicki–Gries reasoning to account for concurrency. The main differences are that 1) our program annotations contain view-based assertions that allow reasoning about weak and persistent memory behaviours; and 2) we define a crash invariant to describe the recoverable state of the program after a crash. We proceed by first defining proof outlines, then providing syntactic rules for proving their validity. Our proof rules are *syntactic*, and thus can be understood and used without having to understand the details of the underlying $\text{Px86}_{\text{view}}$ model.

We let $\text{ASSERTION}_{\text{PV}}$ be the set of *assertions* (i.e. predicates over $\text{Px86}_{\text{view}}$ states) that use view-based expressions (§3.2). A *crash invariant*, $I \in \text{INV} \subset \text{ASSERTION}_{\text{PV}}$, is defined over persistent views only, i.e. it only comprises the persistent view expressions of the form $[x]^P$. We model program annotations via an *annotation function*, $\text{ann} \in \text{ANN} = \text{TID} \times \text{LAB} \rightarrow \text{ASSERTION}_{\text{PV}}$, associating each program point (τ, i) with its associated assertion. A *proof outline* is a tuple (in, ann, I, fin) , where $in, fin \in \text{ASSERTION}_{\text{PV}}$ are the initial and final assertions.

Example 2. The annotation of the proof in Fig. 4 is given by ann , with the mappings of thread 1 as shown below; the mappings of thread 2 are similar.

$$ann \triangleq \{(1, \iota) \mapsto P_1, (1, 2) \mapsto P_2, (1, 3) \mapsto P_3, (1, \zeta) \mapsto P_4, \dots\}$$

Additionally, we have $in \triangleq a = 0 \wedge \forall o \in \{x, y, z\}, \tau \in \{1, 2\}. [o]_\tau = [o]^P = \{0\}$, $fin \triangleq [x]^P = \{1\}$ and $I \triangleq [z]^P = \{1\} \Rightarrow [x]^P = \{1\}$.

Definition 1 (Valid proof outline). A proof outline (in, ann, I, fin) is *valid* for a program Π iff the following hold:

Initialisation. For all $\tau \in \text{TID}$, $in \Rightarrow ann(\tau, \iota)$.

Finalisation. $(\bigwedge_{\tau \in \text{TID}} ann(\tau, \zeta)) \Rightarrow fin$.

Local correctness. For all $\tau \in \text{TID}$ and $i \in \text{LAB}$, either:

- $\Pi(\tau, i) = \alpha$ **goto** j and $\{ann(\tau, i)\} \alpha \{ann(\tau, j)\}$; or
- $\Pi(\tau, i) = \mathbf{if} B \mathbf{goto} j \mathbf{else to} k$ and both $ann(\tau, i) \wedge B \Rightarrow ann(\tau, j)$ and $ann(\tau, i) \wedge \neg B \Rightarrow ann(\tau, k)$ hold; or
- $\Pi(\tau, i) = \langle \alpha \mathbf{goto} j, \hat{a} := \hat{e} \rangle$ and $\{ann(\tau, i)\} \alpha \{ann(\tau, j)[\hat{e}/\hat{a}]\}$.

Stability. For all $\tau_1, \tau_2 \in \text{TID}$ such that $\tau_1 \neq \tau_2$ and $i_1, i_2 \in \text{LAB}$:

- if $\Pi(\tau_1, i_1) = \alpha$ **goto** j , then $\{ann(\tau_2, i_2) \wedge ann(\tau_1, i_1)\} \alpha \{ann(\tau_2, i_2)\}$;

- if $\Pi(\tau_1, i_1) = \langle \alpha \text{ goto } j, \hat{a} := \hat{e} \rangle$, then
 $\{ann(\tau_2, i_2) \wedge ann(\tau_1, i_1)\} \alpha \{ann(\tau_2, i_2)[\hat{e}/\hat{a}]\}$.

Persistence. There exists $\tau \in \text{TID}$ such that for all $i \in \text{LAB}$, $ann(\tau, i) \Rightarrow I$.

Intuitively, **Initialisation** (resp. **Finalisation**) ensures that the initial (resp. final) assertion of each thread holds at the beginning (resp. end); **Local correctness** establishes annotation validity for each thread; **Stability** ensures that each (local) thread annotation is *interference-free* under the execution of other threads [30]; and **Persistence** ensures that the crash invariant holds at every program point for some thread.

Example 3. Given the program in [Example 1](#) and its annotation in [Example 2](#), both **Initialisation** and **Finalisation** clearly hold. Moreover, **Persistence** holds for thread 1. For **Local correctness** of thread 1, we must prove (1)–(3) below; **Local correctness** of thread 2 is similar.

$$\{P_1\} \text{ store } x \ 1 \ \{P_2\} \tag{1}$$

$$\{P_2\} \text{ flush } x \ \{P_3\} \tag{2}$$

$$\{P_3\} \text{ store } y \ 1 \ \{P_4\} \tag{3}$$

For **Stability** of P (the precondition of **store** $x \ 1$ in thread 1) against thread 2 we must prove:

$$\{P_1\} a := \text{load } y \ \{P_1\} \tag{4}$$

$$\{P_1 \wedge a = 1\} \text{ store } z \ 1 \ \{P_1\} \tag{5}$$

Stability of other assertions (i.e., P_2 – P_4) is similar. We prove (1)–(5) in [§3.4](#).

3.4 PIEROGI Proof rules

One of the main benefits of PIEROGI is the ability to perform proofs at a high level of abstraction. In this section, we provide the set of proof rules that we use. The annotation within a proof outline is, in essence, an invariant mapping each program location to an assertion that holds at the program location. Thus, we prove local correctness by checking that each atomic step of a thread establishes the assertions in that thread. Similarly, we check stability by checking each assertion in one thread against each atomic step of the other threads. To enable proof abstraction, we introduce a set of proof rules that describe the interaction between the assertions from [§3.2](#) and the atomic program steps. We will use the standard decomposition rules from Hoare Logic to reduce proof outlines and enable our rules over atomic steps to be applied.

Standard Decomposition Rules The standard decomposition rules we use are given in [Fig. 6](#), which allow one to weaken preconditions and strengthen postconditions, and decompose conjunctions and disjunctions.

Rules for Atomic Statements and View-Based Assertions Weak and persistent memory models (e.g. Px86) are inherently non-deterministic. Moreover in contrast to sequential consistent, in view-based operational semantics

$$\text{Cons} \frac{P' \Rightarrow P \quad Q \Rightarrow Q' \quad \frac{\{P\} \text{ II } \{Q\}}{\{P'\} \text{ II } \{Q'\}}}{\{P'\} \text{ II } \{Q'\}} \quad \text{Conj} \frac{\frac{\{P_1\} \text{ II } \{Q_1\}}{\{P_2\} \text{ II } \{Q_2\}}}{\{P_1 \wedge P_2\} \text{ II } \{Q_1 \wedge Q_2\}} \quad \text{Disj} \frac{\frac{\{P_1\} \text{ II } \{Q_1\}}{\{P_2\} \text{ II } \{Q_2\}}}{\{P_1 \vee P_2\} \text{ II } \{Q_1 \vee Q_2\}}$$

Fig. 6: Standard decomposition rules of PIEROGI

Precondition	Statement	Postcondition	Const.	Ref.
$\{[x]_\tau = S\}$	$a := \mathbf{load} \ x$	$\{a \in S \wedge [x]_\tau \subseteq S\}$		LP ₁
$\{u \in [x]_\tau \Rightarrow \langle x, u \rangle [y]_\tau = S\}$		$\{a = u \Rightarrow [y]_\tau \subseteq S\}$		LP ₂
$\{ x, u = 1 \wedge \llbracket x \rrbracket_{\tau'} \wedge [x]_{\tau'} = \{u\}\}$		$\{a = u \Rightarrow [x]_\tau = \{u\}\}$		LP ₃
$\{true\}$	$\mathbf{store} \ x \ v$	$\{[x]_\tau = \{v\}\}$	$\tau \neq \tau'$	SP ₁
$\{[x]_{\tau'} = S\}$		$\{[x]_{\tau'} = S \cup \{v\}\}$		SP ₂
$\{[x]_{\tau'}^A = S\}$		$\{[x]_{\tau'}^A = S \cup \{v\}\}$		SP ₃
$\{[x]^P = S\}$		$\{[x]^P = S \cup \{v\}\}$		SP ₄
$\{[y]_\tau = S \wedge v \notin [x]_{\tau'}\}$		$\{\langle x, v \rangle [y]_{\tau'} \subseteq S\}$		SP ₅
$\{true\}$		$\{\llbracket x \rrbracket_\tau \wedge \llbracket x \rrbracket_{\tau'}^F\}$		SP ₆
$\{ x, v = n\}$		$\{ x, v = n + 1\}$		SP ₇
$\{[x]_\tau = S\}$	$\mathbf{flush} \ x$	$\{[x]^P \subseteq S \wedge [x]_\tau^A \subseteq S\}$		FP ₁
$\{[x]^P = S\}$		$\{[x]^P \subseteq S\}$		FP ₂
$\{\llbracket x \rrbracket_{\tau'} \wedge [x]_{\tau'} = \{u\} \wedge \llbracket x \rrbracket_\tau^F\}$		$\{[x]^P = \{u\}\}$		FP ₃
$\{[x]_\tau = S \vee [x]_\tau^A = S\}$	$\mathbf{flush}_{\text{opt}} \ x$	$\{[x]_\tau^A \subseteq S\}$		OP
$\{[x]_\tau^A = S \vee [x]^P = S\}$	\mathbf{sfence}	$\{[x]^P \subseteq S\}$		SFP

Fig. 7: Selected proof rules for atomic statements executed by thread τ

(such as $\text{Px86}_{\text{view}}$) instructions such as $a := \mathbf{load} \ x$ have may a side-effect since they may update the view of the thread performing the **load** (cf. [11]). Therefore, unlike Hoare Logic, which contains a single rule for assignment, we have a set of rules for atomic statements, describing their interaction with view-based assertions. Each of the rules in this section has been proved sound with respect to the view-based semantics encoded in Isabelle/HOL.

A selection of these rules for the atomic statements is given in Fig. 7, where the statement is assumed to be executed by thread τ . The first column contains the pre/post condition triple, the second any additional constraints and the third, labels that we use to refer to the rules in our descriptions below. Unless explicitly mentioned as a constraint, we do not assume that threads, locations and values are distinct; e.g. rule LP₃ (referring to τ and τ') holds regardless of whether $\tau = \tau'$ or not.

The rules in Fig. 7 provide high-level insights into the low-level semantics of $\text{Px86}_{\text{view}}$ without having to understand the operational details. The LP_{*i*} rules are for statement $a := \mathbf{load} \ x$. Rule LP₁ states that if τ 's view of x is the set of values S , then in the post state a is an element of S and moreover τ 's view of x is a subset of S (since τ 's view may have shifted). By LP₂, provided the conditional view of τ on y (with condition $x = u$) is S , if the load returns value u , then the view of τ is shifted so that $[y]_\tau \subseteq S$. We only have $[y]_\tau \subseteq S$ in the postcondition because there may be multiple writes to x with value u ; reading x

read may shift the view to the latter write, thus *reducing* the set of values that τ can read for y . LP_3 describes conditions for a deterministic load by thread τ . The precondition assumes that there is only one write to x with value u , that *some* thread τ' sees the last write to x with value u . Then, if τ reads u , its view of x is also constrained to just the set containing u .

The store rules, SP_i , reflect that fact that a new write modifies the views of the other threads as well as the persistent memory and asynchronous views. The first four rules describe the interaction of a **store** by thread τ with current view assertions. By SP_1 , the **store** ensures that the current view of τ is solely the value v written by τ . This is because in $\text{Px86}_{\text{view}}$, new writes are introduced by the executing thread, τ , with a maximal timestamp (see **STORE** rule in Fig. 12), and τ 's view is updated to this new write. SP_2 , SP_3 and SP_4 are similar, and assuming that the view (of another thread, persistent memory and asynchronous view, respectively) in the pre-state is S , shows that the view in the post state is $S \cup \{v\}$. Rule SP_5 allows one to *introduce* a conditional observation assertion $\langle x, v \rangle [y]_{\tau'}$ where $\tau' \neq \tau$. The pre-state of SP_5 assumes that τ 's view of y is the set S , and that τ' cannot view value v for y . Rule SP_6 introduces last-view assertions for τ after τ performs a write to x , and finally SP_7 states that the number of writes to x with value v increases by 1 after executing **store** $x v$.

Rules FP_i describe the effect of **flush** x on the state. FP_1 states that, provided that the current view of τ for x is the set of values S , after executing **flush** x , we are guaranteed that both the persistent view and asynchronous view of τ for x are subsets of S . We obtain a subset in the post state since the $\text{Px86}_{\text{view}}$ semantics potentially moves the persistent and asynchronous views forward. Similarly, by FP_2 if the current persistent view of x is S , then after executing **flush** x the persistent view will be a subset of S . Finally, FP_3 provides a mechanism for establishing a deterministic persistent view u for x . The precondition assumes that *some* thread's view of x is the last write with value u and that τ 's view is such that the flush is guaranteed to flush to this last write to x .

Rule OP describes how the asynchronous view of τ in the postcondition of **flush**_{opt} x is related to the current view of τ and the asynchronous view in the precondition. Finally, rule SFP describes the relationship between the persistent view in the postcondition and the asynchronous view and persistent view in the precondition for an **sfence** instruction.

Our Isabelle/HOL development contains further rules for the other instructions, including **mfence** and **cas**, which we omit here for space reasons. In addition, we prove the stability of several assertions (see Fig. 8 for a selection). An assertion P is *stable* over a statement α executed by τ iff $\{P\} \alpha \{P\}$ holds.

Well-formedness The final major aspect of our framework is a well-formedness condition that describes the set of reachable states in the $\text{Px86}_{\text{view}}$ semantics. The condition is expressed as an invariant of the semantics: it holds initially, and is stable under every possible transition of $\text{Px86}_{\text{view}}$. In fact, the rules in Figs. 7 and 8 are proved with respect to this well-formedness condition.

The majority of the well-formedness constraints are straightforward, e.g. describing the relationship between the views of different components. The most

Statement	Stable Assert.	Const.	Ref.	Statement	Stable Assert.	Const.	Ref.
$a := \mathbf{load} \ x$	$\{[y]_{\tau'} = S\}$	$\tau \neq \tau'$	LS ₁	$\mathbf{store} \ x \ v$	$\{[y]_{\tau'} = S\}$	$x \neq y$	WS ₁
	$\{[y]^P = S\}$		LS ₂		$\{[y]^P = S\}$	$x \neq y$	WS ₂
	$\{[y]_{\tau'}^A = S\}$		LS ₃		$\{[y]_{\tau'}^A = S\}$	$x \neq y$	WS ₃
	$\{a = k\}$		LS ₄		$\{a = k\}$		WS ₄
	$\{\llbracket y \rrbracket_{\tau'}\}$		LS ₅		$\{\llbracket y \rrbracket_{\tau'}\}$	$x \neq y$	WS ₅
$\mathbf{flush} \ x$	$\{[y]_{\tau'} = S\}$	$x \neq y$	FS ₁	$\mathbf{flush}_{\text{opt}} \ x$	$\{[y]_{\tau'} = S\}$		OS ₁
	$\{[y]^P = S\}$		FS ₂		$\{[y]^P = S\}$		OS ₂
	$\{\llbracket y \rrbracket_{\tau'}\}$		FS ₃		$\{\llbracket y \rrbracket_{\tau'}^F\}$	$x \neq y \vee v \neq v'$	OS ₃
	$\{\llbracket y \rrbracket_{\tau'}^F\}$		FS ₄		$\{[y, v'] = n\}$		
	$\{[y, v] = n\}$		FS ₅				
\mathbf{sfence}	$\{[x]_{\tau'} = S\}$		SFS ₁				
	$\{[x, v] = n\}$		SFS ₂				

Fig. 8: Selection of stable assertions for atomic statements executed by thread τ

important component of the well-formedness condition is a non-emptiness condition on views, which states that $[x]_{\tau} \neq \emptyset \wedge [x]^P \neq \emptyset \wedge [x]_{\tau}^A \neq \emptyset$. For instance, a consequence of this condition is that, in combination with LP₁, we have:

$$\{[y]_{\tau} = \{v\}\} \ a := \mathbf{load} \ x \ \{[y]_{\tau} = \{v\}\} \quad (6)$$

Worked Example We now return to the proof obligations from [Example 3](#) and demonstrate how they can be discharged using the proof rules described above. For Local correctness, condition (1) holds by Conj (from [Fig. 6](#)) together with stability rules WS₁, WS₂ and WS₄ (from [Fig. 8](#)) which establish the first three conjunctions in the postcondition, and SP₁ from [Fig. 7](#), which establishes the final conjunction. Condition (2) holds by FP₁ in [Fig. 7](#) together with Cons (from [Fig. 6](#)). Finally, condition (3) holds by WS₂ (from [Fig. 8](#)).

Both the Stability conditions (4) and (5) from [Example 3](#) hold by the stability rules in [Fig. 8](#) together with Cons and Conj (from [Fig. 6](#)). In particular, for (4), we use rules LS₁, LS₂ and LS₄, and for (5), we use WS₁, WS₂ and WS₄.

4 Examples

In this section we present a selection of programs that we have verified in Isabelle/HOL. These examples highlight specific aspects of Px86, in particular, the interaction between $\mathbf{flush}_{\text{opt}}$ and \mathbf{sfence} , as well as aspects of our view-based assertion language that simplifies verification.

Optimised Message Passing We start by considering a variant of [Fig. 1e](#), which contains two optimisations. First, we notice that flushing of the write to x in thread 1 can be moved to thread 2 since the write to z is guarded by whether or not thread 2 reads the flag y . Second, it is possible to replace the \mathbf{flush} by a more optimised $\mathbf{flush}_{\text{opt}}$ followed by an \mathbf{sfence} . We confirm correctness of these optimisations via the proof outline in [Fig. 9](#). The optimised message passing in [Fig. 9](#) ensures the same persistent invariant as [Fig. 1e](#). However, the way in

$$\begin{array}{l}
\{ \forall o \in \{x, y, z\}, \tau \in \{1, 2\}. [o]_\tau = [o]^P = [o]_\tau^\Delta = \{0\} \} \\
\left\{ \begin{array}{l}
\{ (1 \in [y]_2 \Rightarrow \langle y, 1 \rangle [x]_2 = \{1\}) \wedge [y]_2 \subseteq \{0, 1\} \wedge [z]^P = \{0\} \} \\
a := \mathbf{load} \ y; \\
\{ (a = 1 \Rightarrow [x]_2 = \{1\}) \wedge [z]^P = \{0\} \} \\
\mathbf{if} \ (a \neq 0) \\
\quad \{ [x]_2 = \{1\} \wedge [z]^P = \{0\} \} \\
\quad \mathbf{flush}_{\text{opt}} \ x; \\
\quad \{ [x]_2^\Delta = \{1\} \wedge [z]^P = \{0\} \} \\
\quad \mathbf{sfence}; \\
\quad \{ [x]^P = \{1\} \} \\
\quad \mathbf{store} \ z \ 1; \\
\quad \{ [z]^P = \{0\} \vee [x]^P = \{1\} \} \\
\quad \{ [z]^P = \{0\} \vee [x]^P = \{1\} \} \\
\quad \{ \hat{z} : [z]^P = \{1\} \Rightarrow [x]^P = \{1\} \}
\end{array} \right\} \\
\left\{ \begin{array}{l}
\{ [y]_2 = \{0\} \} \\
\mathbf{store} \ x \ 1; \\
\left\{ \begin{array}{l}
\{ [y]_2 = \{0\} \wedge \\
[x]_1 = \{1\} \}
\end{array} \right\} \\
\mathbf{store} \ y \ 1; \\
\{ \mathbf{true} \}
\end{array} \right\}
\end{array}$$

Fig. 9: Proof outline for optimised message passing

which this is established differs. In particular, in Fig. 1e, the persistent invariant holds due to thread 1, whereas in Fig. 9 it holds due to thread 2.

With respect to the persistent invariant, the most important sequence of steps takes place in thread 2 if it reads 1 for y . Note that by the conditional view assertion in the precondition of $a := \mathbf{load} \ y$, thread 2 is guaranteed to read 1 for x after reading 1 for y . Thus, if the test of \mathbf{if} statement succeeds, then thread 2 must see 1 for x . This view is translated into an asynchronous view after the $\mathbf{flush}_{\text{opt}}$ is executed, and then to the persistent view after executing \mathbf{sfence} . Note that until this occurs, we can guarantee that $[z]^P = \{0\}$, which trivially guarantees the persistent invariant.

Flush Buffering Our next example is a variation of store buffering (**SB**) and is used to highlight how writes by different threads on different locations interact with flushes. Here, thread 1 writes to x and flushes y , while thread 2 writes to y then flushes x .⁶ The writes to w and z are used to witness whether the flushes in both threads have occurred. The persistent invariant states that, if both w and z hold 1 in persistent memory, then either x or y has the new value (i.e. 1) in persistent memory. If both threads perform their \mathbf{flush} operations, then at least one must flush value 1 since a \mathbf{flush} cannot be reordered with a \mathbf{store} .

Although simple to state, the proof is non-trivial since it requires careful analysis of the order in which the stores to x and y occur. In the semantics of Cho et al. [9], the \mathbf{flush} corresponding to the *second store* instruction executed synchronises with writes to *all* locations. Thus, for example, if thread 1's store to x is executed after thread 2's store to y , then the subsequent \mathbf{flush} in thread 1 is guaranteed to flush the new write to y .

The above intuition requires reasoning about the order in which operations occur. To facilitate this, we use auxiliary variables \hat{a} and \hat{b} to record the order in which the writes to x and y occur; $\hat{a} = 1$ iff the write to x occurs before the

⁶ Note that the \mathbf{flush} operations here are analogous to the \mathbf{load} instructions in **SB**.

$$\left\{ \begin{array}{l}
(\forall \tau \in \{1, 2\}, o \in \{x, y, z\}, [o]_\tau = [o]^P = \{0\}) \wedge a = 0 \\
\{[y]^P = \{0\} \wedge [z]^P = \{0\} \wedge (|x, 2| \in \{0, 1\})\} \\
\text{store } x \ 1; \\
\left\{ \begin{array}{l}
([x]_2 = 1 \vee \left([x]_2 = \{1, 2\} \wedge |x, 2| = 1 \wedge \right)) \wedge \\
\left(\begin{array}{l}
\llbracket x \rrbracket_1 \wedge [x]_1 = 2 \\
[y]^P = \{0\} \wedge [z]^P = \{0\}
\end{array} \right) \wedge
\end{array} \right\} \\
a := \text{load } x; \\
\{(a = 2 \Rightarrow [x]_2 = \{2\}) \wedge [y]^P = \{0\} \wedge [z]^P = \{0\}\} \\
\text{flush}_{\text{opt}} x; \\
\{(a = 2 \Rightarrow [x]_2^A = \{2\}) \wedge [y]^P = \{0\} \wedge [z]^P = \{0\}\} \\
\text{if } (a = 2) \\
\quad \{[x]_2^A = \{2\} \wedge [y]^P = \{0\} \wedge [z]^P = \{0\}\} \\
\quad \text{store } y \ 1; \\
\quad \{([x]_2^A = \{2\} \vee [y]^P = \{0\}) \wedge [z]^P = \{0\}\} \\
\text{sfence}; \\
\quad \{[x]^P = \{2\} \vee [y]^P = \{0\}\} \\
\text{store } z \ 1; \\
\quad \{[x]^P = \{2\} \vee [y]^P = \{0\} \vee [z]^P = \{0\}\} \\
\{[x]^P = \{2\} \vee [y]^P = \{0\} \vee [z]^P = \{0\}\} \\
\{\{ \not\exists : [y]^P = \{1\} \wedge [z]^P = \{1\} \Rightarrow [x]^P = \{2\} \}\}
\end{array} \right\}$$

Fig. 11: Proof outline for epoch persistency

invariant for thread 2 requires reasoning about the view of thread 2 for address x (i.e. $[x]_2$) after the execution of the instruction $a := \text{load } x$. Notice here that $a := \text{load } x$ is ordered with respect to the later $\text{flush}_{\text{opt}} x$ instruction. Consequently, any impact of the execution of the **load** on $[x]_2$, will also affect $[x]_2^A$. Taking into account the ordering of the writes at the address x , we can conclude that if thread 2 reads the value 2, it reads the value of the last write at x . This is expressed with the assertion $\llbracket x \rrbracket_1$ in the precondition of $a := \text{load } x$, which states that the threads 1's view of x is the last write to x . By rule LP₃, if a thread τ 's view of an address x contains only the last write at this address, and the last value written at this address appears only once at the memory, then if a thread τ read this value at x , its view of x (i.e. $[x]_\tau$) is guaranteed to contain only the last written value at x . Consequently, after reading value 2, thread 2's view of x contains only the value 2 (i.e. $[x]_2 = \{2\}$). Execution of $\text{flush}_{\text{opt}} x$ ensures $[x]_2^A$ (by rule OP). As a result, in the case that the **if** statement succeeds, after the execution of the **sfence** it is guaranteed that the value 2 is persisted at x (i.e. $[x]^P = \{2\}$). In the case that the **if** statement fails, $[y]^P = \{0\}$ must hold, thus the persistent invariant holds trivially.

5 PIEROGI Soundness

In this section we present the $\text{Px86}_{\text{view}}$ model from [9] (§5.1), formally interpret our assertions as predicates on states of that model (§5.2), and establish the soundness of the proposed reasoning technique (§5.3).

$$\begin{array}{c}
\text{(ASSIGN)} \\
\frac{\alpha = a := e \quad v = T.\text{regs}(e) \quad T' = T[\text{regs}(a) \mapsto v]}{\langle T, M \rangle \xrightarrow{\alpha} \langle T', M \rangle} \\
\\
\text{(STORE)} \\
\frac{\alpha = \mathbf{store} \ x \ e \quad v = T.\text{regs}(e) \quad M' = M \uparrow\uparrow [\langle x := v \rangle] \quad T' = T[\text{coh}(x) \mapsto |M|]}{\langle T, M \rangle \xrightarrow{\alpha} \langle T', M' \rangle} \\
\\
\text{(LOAD-INTERNAL)} \\
\frac{\alpha = a := \mathbf{load} \ x \quad M[t] = \langle x := v \rangle \quad T.\text{coh}(x) = t \quad T' = T[\text{regs}(a) \mapsto v]}{\langle T, M \rangle \xrightarrow{\alpha} \langle T', M \rangle} \\
\\
\text{(LOAD-EXTERNAL)} \\
\frac{\alpha = a := \mathbf{load} \ x \quad M[t] = \langle x := v \rangle \quad T.\text{coh}(x) < t \quad x \notin M(t..T.\text{VrNew}] \quad T' = T \left[\begin{array}{l} \text{regs}(a) \mapsto v, \\ \text{coh}(x) \mapsto t, \\ \text{VrNew} \mapsto_{\sqcup} t, \\ \text{VpReady} \mapsto_{\sqcup} t \end{array} \right]}{\langle T, M \rangle \xrightarrow{\alpha} \langle T', M \rangle} \\
\\
\text{(SFENCE)} \\
\frac{\alpha = \mathbf{sfence} \quad T' = T \left[\begin{array}{l} \text{VpReady} \mapsto_{\sqcup} T.\text{maxcoh}, \\ \text{VpCommit} \mapsto_{\sqcup} T.\text{VpAsync} \end{array} \right]}{\langle T, M \rangle \xrightarrow{\alpha} \langle T', M \rangle} \\
\\
\text{(FLUSH)} \\
\frac{\alpha = \mathbf{flush} \ x \quad T' = T \left[\begin{array}{l} \text{VpAsync}(x) \mapsto_{\sqcup} T.\text{maxcoh}, \\ \text{VpCommit}(x) \mapsto_{\sqcup} T.\text{maxcoh} \end{array} \right]}{\langle T, M \rangle \xrightarrow{\alpha} \langle T', M \rangle} \\
\\
\text{(FLUSHOPT)} \\
\frac{\alpha = \mathbf{flush}_{\text{opt}} \ x \quad T' = T[\text{VpAsync}(x) \mapsto_{\sqcup} T.\text{coh}(x) \sqcup T.\text{VpReady}]}{\langle T, M \rangle \xrightarrow{\alpha} \langle T', M \rangle} \\
\hline
\text{(PROGRAM-NORMAL)} \\
\frac{\vec{p}\vec{c}(\tau) = i \quad \Pi(\tau, i) = \alpha \ \mathbf{goto} \ j \quad \langle \vec{T}(\tau), M \rangle \xrightarrow{\alpha} \langle T', M' \rangle \quad \vec{p}\vec{c}' = \vec{p}\vec{c}[\tau \mapsto j] \quad \vec{T}' = \vec{T}[\tau \mapsto T']}{\langle \vec{p}\vec{c}, \vec{T}, M, G \rangle \Rightarrow_{\Pi} \langle \vec{p}\vec{c}', \vec{T}', M', G \rangle} \\
\\
\text{(PROGRAM-IF)} \\
\frac{\vec{p}\vec{c}(\tau) = i \quad \Pi(\tau, i) = \mathbf{if} \ B \ \mathbf{goto} \ j \ \mathbf{else} \ \mathbf{to} \ k \quad \vec{p}\vec{c}' = \vec{p}\vec{c} \left[\tau \mapsto \begin{cases} j & \vec{T}(\tau).\text{regs}(B) = \mathbf{true} \\ k & \vec{T}(\tau).\text{regs}(B) = \mathbf{false} \end{cases} \right]}{\langle \vec{p}\vec{c}, \vec{T}, M, G \rangle \Rightarrow_{\Pi} \langle \vec{p}\vec{c}', \vec{T}, M, G \rangle} \\
\\
\text{(PROGRAM-GHOST)} \\
\frac{\vec{p}\vec{c}(\tau) = i \quad \Pi(\tau, i) = \langle \alpha \ \mathbf{goto} \ j, \hat{a} := \hat{e} \rangle \quad \langle \vec{T}(\tau), M \rangle \xrightarrow{\alpha} \langle T', M' \rangle \quad \vec{p}\vec{c}' = \vec{p}\vec{c}[\tau \mapsto j] \quad \vec{T}' = \vec{T}[\tau \mapsto T'] \quad G' = G[\hat{a} \mapsto G(\hat{e})]}{\langle \vec{p}\vec{c}, \vec{T}, M, G \rangle \Rightarrow_{\Pi} \langle \vec{p}\vec{c}', \vec{T}', M', G' \rangle}
\end{array}$$

Fig. 12: Transitions of $\text{Px86}_{\text{view}}$ for a program Π

5.1 The $\text{Px86}_{\text{view}}$ Model

Like previous view-based models, $\text{Px86}_{\text{view}}$ employs a non-standard memory capturing all previously executed writes, alongside so-called “thread views” that track several position(s) of each thread in that history and enforce limitations on the ability of the thread to read from and write to the memory. In addition, the thread views contain the necessary information for determining the possible contents of the non-volatile memory upon a system crash. Formally, $\text{Px86}_{\text{view}}$ ’s memory and thread states are defined as follows.

Definition 2 (Px86_{view}’s memory). A *memory* $M \in \text{MEMORY}$ is a list of *messages*, where each message has the form $\langle x := v \rangle$ for some $x \in \text{LOC}$ and $v \in \text{VAL}$. We use $w.\text{loc}$ and $w.\text{val}$ to refer to the two components of a message

w . We use standard list notations for memories (e.g. $M_1 ++ M_2$ for appending memories, $[w]$ for a singleton memory, and $|M|$ for the length of M). We refer to indices (starting from 0) in a memory M as *timestamps*, and denote the t 'th element of M as $M[t]$. We use \sqcup for obtaining the maximum among timestamps (i.e. $t_1 \sqcup t_2 = \max(t_1, t_2)$), and extend this notation pointwise to functions. We write $x \notin M(t_2..t_1]$ for the condition $\forall t_2 < t \leq t_1. M[t].\text{loc} \neq x$.

Definition 3 (Px86_{view}'s thread states). A *thread state* $T \in \text{THREAD}$ is a record consisting of the following fields: $\text{coh} : \text{LOC} \rightarrow \mathbb{N}$, $\text{v}_{\text{rNew}} : \mathbb{N}$, $\text{v}_{\text{pReady}} : \mathbb{N}$, $\text{v}_{\text{pAsync}} : \text{LOC} \rightarrow \mathbb{N}$, and $\text{v}_{\text{pCommit}} : \text{LOC} \rightarrow \mathbb{N}$. We use standard function/record update notation (e.g. $T' = T[\text{coh}(x) \mapsto t]$ denotes the thread state obtained from T by modifying the x entry in the coh component of T to t). In addition, \mapsto_{\sqcup} is used to incorporate certain timestamps in fields (e.g. $T[\text{v}_{\text{rNew}} \mapsto_{\sqcup} t]$ denotes the thread state obtained from T by modifying the v_{rNew} component of T to $T.\text{v}_{\text{rNew}} \sqcup t$). We denote by $T.\text{maxcoh}$ the maximum among the coherence view timestamps ($T.\text{maxcoh} = \bigsqcup_x T.\text{coh}(x)$).

The two components, together with program counters and the “ghost memory”, are combined in Px86_{view}'s machine states as defined next.

Definition 4 (Px86_{view}'s machine states). A *machine state* is a tuple $\sigma = \langle \vec{p}\vec{c}, \vec{T}, M, G \rangle$ where $\vec{p}\vec{c} : \text{TID} \rightarrow \text{LAB}$ is a mapping assigning the next program label to be executed by each thread, $\vec{T} : \text{TID} \rightarrow \text{THREAD}$ is a mapping assigning the current thread state to each thread, $M \in \text{MEMORY}$ is the current memory, and $G : \text{AUXVAR} \rightarrow \text{VAL}$ is storing the current values of the auxiliary variables. Below we assume that G is extended to expressions $\hat{e} \in \text{AUXEXP}$ in a standard way. We denote the components of a machine state σ by $\sigma.\vec{p}\vec{c}$, $\sigma.\vec{T}$, $\sigma.M$, and $\sigma.G$. In addition, we denote by $\sigma.\text{maxpCommit}(x)$ the maximum among the persistency view timestamps for location x ($\sigma.\text{maxpCommit} = \bigsqcup_{\tau} \sigma.\vec{T}(\tau).\text{v}_{\text{pCommit}}(x)$).

The transitions of Px86_{view} are presented in Fig. 12. These closely follow the model in [9] with minor presentational simplifications. Note, however, that, for simplicity and following [23], we conservatively assume that writes persist atomically at the location granularity (representing, e.g. machine words) rather than at the granularity of the width of a cache line. We refer the interested reader to [9] for a detailed discussion of the transitions rules in Fig. 12.

The above operational definitions naturally induce a notion of a execution (or a “run”) of Px86_{view} on a certain program Π starting from some initial state of the form $\langle \lambda\tau. \iota, \vec{T}, M, G \rangle$. A system crash might occur at any point during the execution. Again, following the model of [9], the non-volatile memory (NVM) is not modeled as a concrete part of the state. Instead, the possible contents of the NVM can be inferred from the machine state (specifically from the memory and the $\text{v}_{\text{pCommit}}$ views of the different threads), as defined next. This definition is presented as “crash transition” in [9].

Definition 5. A non-volatile memory $\text{NVM} : \text{LOC} \rightarrow \text{VAL}$ is *possible in a state* σ if for every $x \in \text{LOC}$, there exists some t such that $\sigma.M[t] = \langle x := \text{NVM}(x) \rangle$ and $x \notin \sigma.M(t..\sigma.\text{maxpCommit}(x))$.

5.2 The Semantics of PIEROGI Assertions

We present the formal definitions of the expressions introduced in §3.2 in terms of $\text{Px86}_{\text{view}}$'s machine states.

Current and conditional views When formalising the *current* and *conditional view* expressions, we start with auxiliary functions that return the sets of observable timestamps visible to the components in question, then extract the values in memory corresponding these timestamps. To facilitate this, we define

$$\text{Vals}(M, TS) \triangleq \{M[t].\text{loc} \mid t \in TS\}$$

where $M \in \text{MEMORY}$ and TS is a set of timestamps.

Thread view To define the meaning of the thread view expression, $[x]_\tau$, we use:

$$\begin{aligned} \text{TS}_\tau^{\text{OF}}(\sigma, x, t) &\triangleq \{t' \mid \sigma.M[t'].\text{loc} = x \wedge \sigma.\vec{T}(\tau).\text{coh}(x) \leq t' \wedge x \notin \sigma.M(t'..t)\} \\ \text{TS}_\tau(\sigma, x) &\triangleq \text{TS}_\tau^{\text{OF}}(\sigma, x, \sigma.\vec{T}(\tau).\text{vrNew}) \end{aligned}$$

$\text{TS}_\tau^{\text{OF}}(\sigma, x, t)$ returns the set of *timestamps* that are *observable from* timestamp t for thread τ to read for location x in state σ ; and $\text{TS}_\tau(\sigma, x)$ returns the set of *timestamps* that are *observable* for τ to read x in σ . Note that after instantiating t to $\sigma.\vec{T}(\tau).\text{vrNew}$ in $\text{TS}_\tau^{\text{OF}}(\sigma, x, t)$, we obtain the premises of the load rules in Fig. 12. Then, $[x]_\tau \triangleq \lambda\sigma. \text{Vals}(\sigma.M, \text{TS}_\tau(\sigma, x))$, i.e. is the set of values in $\sigma.M$ corresponding to the timestamps in $\text{TS}_\tau(\sigma, x)$.

Persistent memory view For the persistent memory view expression, $[x]^{\text{P}}$, we use:

$$\text{TS}^{\text{P}}(\sigma, x) = \{t \mid \sigma.M[t].\text{loc} = x \wedge x \notin \sigma.M(t..\sigma.\text{maxpCommit}(x))\}$$

which returns the set of *timestamps* that are observable to the *persistent memory* for x in σ . Then, $[x]^{\text{P}} \triangleq \lambda\sigma. \text{Vals}(\sigma.M, \text{TS}^{\text{P}}(\sigma, x))$. Note that the second conjunct within the definition of $\text{TS}^{\text{P}}(\sigma, x)$ is precisely the condition that links $\text{Px86}_{\text{view}}$ states to NVM states (Definition 5). Given this definition, we have:

Proposition 1. *A non-volatile memory NVM : LOC \rightarrow VAL is possible in a state σ iff $\text{NVM}(x) \in [x]^{\text{P}}(\sigma)$ for every $x \in \text{LOC}$.*

Asynchronous memory view To define the meaning of the asynchronous memory view, $[x]_\tau^{\text{A}}$, we use:

$$\text{TS}_\tau^{\text{A}}(\sigma, x) \triangleq \{t \mid \sigma.M[t].\text{loc} = x \wedge x \notin \sigma.M(t..\sigma.\vec{T}(\tau).\text{vpASync}(x))\}$$

which returns the timestamps of the asynchronous view of thread τ in location x and state σ . Then, as before, $[x]_\tau^{\text{A}} \triangleq \lambda\sigma. \text{Vals}(\sigma.M, \text{TS}_\tau^{\text{A}}(\sigma, x))$.

Conditional view The functions used to define conditional memory view, $\langle x, v \rangle [y]_\tau$, are slightly more sophisticated than those above. We define:

$$\begin{aligned} \text{TS}_\tau^{\text{OV}}(\sigma, x, v) &\triangleq \left\{ t' \mid \begin{array}{l} \exists t \in \text{TS}_\tau(\sigma, x). \sigma.M[t].\text{val} = v \wedge \\ t' = \text{if } t = \sigma.\vec{T}(\tau).\text{coh}(x) \text{ then } \sigma.\vec{T}(\tau).\text{vrNew} \\ \text{else } t \sqcup \sigma.\vec{T}(\tau).\text{vrNew} \end{array} \right\} \\ \text{TS}_\tau^{\text{CO}}(\sigma, x, v, y) &\triangleq \bigcup \{ \text{TS}_\tau^{\text{OF}}(\sigma, y, t) \mid t \in \text{TS}_\tau^{\text{OV}}(\sigma, x, v) \} \end{aligned}$$

where $\text{TS}_\tau^{\text{OV}}(\sigma, x, v)$ returns the set of timestamps that τ can observe for x with value v . Assuming t is a timestamp that τ can observe for x , and the value for x at t is v , the corresponding timestamp t' that $\text{TS}_\tau^{\text{OV}}(\sigma, x, v)$ returns is $\sigma.\vec{T}(\tau).\text{vrNew}$ if τ 's coherence view for x is t , and the maximum of t and $\sigma.\vec{T}(\tau).\text{vrNew}$, otherwise. Given this, $\text{TS}_\tau^{\text{CO}}(\sigma, x, v, y)$ returns the timestamps that τ can observe for y , from any timestamp $t \in \text{TS}_\tau^{\text{OV}}(\sigma, x, v)$. Finally, the set of conditional values is defined by $\langle x, v \rangle [y]_\tau \triangleq \lambda \sigma. \text{Vals}(\sigma.M, \text{TS}_\tau^{\text{CO}}(\sigma, x, v, y))$.

Last view assertions We use the following auxiliary definition:

$$\text{Last}(M, x) \triangleq \bigsqcup \{t \mid M[t].\text{loc} = x\}$$

which returns the timestamp of the last write to x in M . Then, the last view assertions are given by:

- $\llbracket x \rrbracket_\tau \triangleq \{\sigma \mid \text{TS}_\tau(\sigma, x) = \{\text{Last}(\sigma.M, x)\}\}$, i.e. τ 's view of x in σ is the last write to x in σ .
- $\llbracket x \rrbracket_\tau^{\text{F}} \triangleq \{\sigma \mid \text{Last}(\sigma.M, x) \leq \sigma.\vec{T}(\tau).\text{maxcoh} \sqcup \sigma.\text{maxpCommit}(x)\}$, i.e. the maximum of τ 's maximum coherence view and the maximum commit view of x (over all threads) is beyond the last write to x in σ . This means that executing a **flush** x operation in τ will cause the last write of x to be flushed (see FLUSH rule in Fig. 12).

Value count Finally, the value count expression is defined as follows:

$$|x, v| \triangleq \lambda \sigma. |\{t \mid \sigma.M[t] = \langle x := v \rangle\}|$$

5.3 Soundness of PIEROGI

Given the above building blocks, the soundness of the proposed reasoning technique is stated as follows.

Theorem 1 (Soundness of PIEROGI). *Suppose that a program Π has a valid proof outline $\langle \text{in}, \text{ann}, I, \text{fin} \rangle$. Let σ be a state of $\text{Px86}_{\text{view}}$ that is reachable in an execution of Π from some state σ_{init} of the form $\langle \lambda \tau. \iota, \vec{T}_{\text{init}}, M_{\text{init}}, G_{\text{init}} \rangle$ such that $\sigma_{\text{init}} \in \text{in}$. Then, the following hold:*

- 1) For every $\tau \in \text{TID}$, we have that $\sigma \in \text{ann}(\tau, \sigma.\vec{pc}(\tau))$.
- 2) If $\sigma.\vec{pc}(\tau) = \zeta$ for every $\tau \in \text{TID}$, then $\sigma \in \text{fin}$.
- 3) Every non-volatile memory NVM that is possible in σ satisfies the crash invariant I .

Finally, it is straightforward to show the soundness of a standard ‘‘auxiliary variable transformation’’ [30] which removes all auxiliary variables from a program Π (translating each command $\langle \alpha \text{ goto } j, \hat{a} := \hat{e} \rangle$ into $\alpha \text{ goto } j$) provided that the crash invariant and the final assertion do not contain occurrences of the auxiliary variables. Indeed, it is easy to see that the auxiliary memory G in the operational semantics in Fig. 12 serves only as an instrumentation, and does not restrict the possible runs. (Formally, if Π' is obtained from Π by removing all auxiliary variables and $\langle \vec{pc}, \vec{T}, M, G' \rangle$ is reachable in $\Rightarrow_{\Pi'}$ from some initial state, then $\langle \vec{pc}, \vec{T}, M, G \rangle$ is reachable in \Rightarrow_{Π} from the same state for some G .)

6 Mechanisation

Perhaps the greatest strength of our development is an integrated Isabelle/HOL mechanisation providing a fully fledged semi-automated verification tool for $\text{Px86}_{\text{view}}$ programs. This mechanisation builds on the existing work on Owicki–Gries for RC11 by Dalvandi et al [11, 12] applying it to the $\text{Px86}_{\text{view}}$ semantics. We start by encoding the operational semantics of Cho et al. [9], followed by the view-based assertions described in §3.2. Then, we prove correctness of all of the proof rules for the atomic statements, including those described in §3.4. These rules can be challenging to prove since they require unfolding of the assertions and examination of the low-level operational semantics and their effect on the views of different system components.

Once proved, the rules provided are highly reusable, and are key to making verification feasible. Specifically, when showing the validity of a proof outline (Definition 1), Isabelle/HOL generates the necessary proof obligations (after minor interactions) and *automatically* finds the set of high-level proof rules needed to discharge each proof obligation via the built-in sledgehammer tool [6]. This enables a high degree of experimentation and debugging of proof outlines, including the ability to reduce assertion complexity once a proof outline is validated.

The base development (semantics, view-based assertions, and soundness of proof rules) comprise ~ 7000 lines of Isabelle/HOL code. With this base development in place, each example comprises 200–400 lines of code (including the encoding of the program, the annotations, and the proofs of validity). The entire development took approximately 3 months of full-time work.

7 Related Work

The soundness of PIEROGI is proven relative to the $\text{Px86}_{\text{view}}$ of Cho et al. [9]; there are however other equivalent models in the literature [1, 23, 32, 34], as well as other persistency models [33, 35]. While the original persistent x86 semantics has asynchronous explicit persist instructions [34], the underlying model assumed here is due to Cho et al. [9] with synchronous persist instructions. Nevertheless, Khyzha and Lahav [23] formally proved that the two alternatives are equivalent when reasoning about states after crashes (e.g. using our “crash invariants”).

As mentioned in §1, the only existing program logic for persistent programs is POG [31], which (as with PIEROGI) is a descendent of Owicki–Gries [30]. PIEROGI goes beyond POG by handling examples that involve $\mathbf{flush}_{\text{opt}}$ instructions, which cannot be directly verified using POG. Raad et al. [31] provide a transformation technique to replace certain patterns of $\mathbf{flush}_{\text{opt}}$ and \mathbf{sfence} with \mathbf{flush} . Specifically, given a program Π that includes $\mathbf{flush}_{\text{opt}}$ instructions, provided that Π meets certain conditions, this transformation mechanism rewrites Π into an equivalent program Π' that uses \mathbf{flush} instructions instead, allowing one to use POG. However, there are three limitations to this strategy: 1) the rewriting is an external mechanism that requires stepping outside the POG logic; 2) the rewriting is potentially expensive and must be done for every program

that includes **flush**_{opt}; and 3) the transformation technique is incomplete in that not all programs meet the stipulated conditions (e.g. Epoch Persistency 2), and thus cannot be verified using this technique. PIEROGI has no such limitations, as we showed in the examples in Section 4. Moreover, POG has no corresponding mechanisation, and developing a mechanisation that also efficiently handles the program transformation for **flush**_{opt} instructions would be non-trivial.

The Owicki–Gries method was first applied to non-SC memory consistency by Lahav et al. [26]. One way that their approach, which targets the release/acquire memory model, is different from ours is that they aim to use standard SC-like assertions; in order to retain soundness under a weak memory model, they had to strengthen the standard stability conditions on proof outlines. Dalvandi et al. [11, 13] took a different approach when designing their Owicki–Gries logic for the release/acquire fragment of C11: by employing a more expressive, view-based assertion language, they were able to stick with the standard stability requirement. In our work, we follow Dalvandi et al.’s approach. However, our assertions are fine-tuned to cope with the other types of view present in P_{x86}_{view}, such as those corresponding to the persistent and the asynchronous views. It is interesting that some of the principles of view-based reasoning apply to different memory models, and future work could look at unifying reasoning across models.

Dalvandi et al. [13] have developed a deeper integration of their view-based logic using the Owicki–Gries encoding of Nipkow and Prensa Nieto [28] in Isabelle/HOL. Such an integration would be straightforward for PIEROGI too, allowing verification to take place without translating programs into a transition system. This would be much more difficult for POG since Owicki–Gries rules themselves are different from the standard encoding in Isabelle/HOL, in addition to the transformation required for **flush**_{opt} instructions discussed above.

The idea of extending Hoare triples with crash conditions first appeared in the work of Chen et al. [8]. However, that work supports neither concurrency nor explicit flushing instructions. Related ideas are found in the works of Ntzik et al. [29] and Chajed et al. [7]. However, in contrast to PIEROGI, both of these works 1) assume sequentially consistent memory, as opposed to a weak memory model such as TSO; 2) assume strict persistency (where store and persist orders coincide); and 3) assume there is a synchronous **flush** operation, which is easier to reason about than the asynchronous **flush**_{opt} operation.

Besides program logics, there have been other recent efforts to help programmers reason about persistent programs. For instance, Abdulla et al. [1] have proven that state-reachability for persistent x86 is decidable, thus opening the door to automatic verification of persistent programs, and Gorjiara et al. [18] and Kokologiannakis et al. [25] have developed model checkers for finding bugs in persistent programs. Recent works have considered durable atomic objects such as concurrent data structures [17] and transactional memory [3] and their verification [3, 14, 15], which have been designed to satisfy conditions such as durable linearizability [20, 24] and durable opacity [3]. These proofs assume persistency under SC; our work provides foundations for extending these proofs to persistent x86-TSO.

References

1. Abdulla, P.A., Atig, M.F., Bouajjani, A., Kumar, K.N., Saivasan, P.: Deciding reachability under persistent x86-TSO. *Proc. ACM Program. Lang.* **5**(POPL), 1–32 (2021). <https://doi.org/10.1145/3434337>
2. Apt, K.R., de Boer, F.S., Olderog, E.: *Verification of Sequential and Concurrent Programs*. Texts in Computer Science, Springer (2009). <https://doi.org/10.1007/978-1-84882-745-5>
3. Bila, E., Doherty, S., Dongol, B., Derrick, J., Schellhorn, G., Wehrheim, H.: Defining and verifying durable opacity: Correctness for persistent software transactional memory. In: Gotsman, A., Sokolova, A. (eds.) *FORTE*. Lecture Notes in Computer Science, vol. 12136, pp. 39–58. Springer (2020). https://doi.org/10.1007/978-3-030-50086-3_3
4. Bila, E.V., Dongol, B., Lahav, O., Raad, A., Wickerson, J.: Isabelle/HOL files for "View-Based Owicki-Gries Reasoning for Persistent x86-TSO" (Jan 2022). <https://doi.org/10.6084/m9.figshare.18469103>
5. Bila, E.V., Dongol, B., Lahav, O., Raad, A., Wickerson, J.: View-based Owicki-Gries reasoning for persistent x86-TSO (extended version) (2022), <https://arxiv.org/abs/2201.05860>
6. Böhme, S., Nipkow, T.: Sledgehammer: Judgement day. In: Giesl, J., Hähnle, R. (eds.) *Automated Reasoning, 5th International Joint Conference, IJCAR 2010*, Edinburgh, UK, July 16-19, 2010. Proceedings. LNCS, vol. 6173, pp. 107–121. Springer (2010). https://doi.org/10.1007/978-3-642-14203-1_9
7. Chajed, T., Tassarotti, J., Kaashoek, M.F., Zeldovich, N.: Verifying concurrent, crash-safe systems with perennial. In: Brecht, T., Williamson, C. (eds.) *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP 2019*, Huntsville, ON, Canada, October 27-30, 2019. pp. 243–258. ACM (2019). <https://doi.org/10.1145/3341301.3359632>
8. Chen, H., Ziegler, D., Chajed, T., Chlipala, A., Kaashoek, M.F., Zeldovich, N.: Using crash hoare logic for certifying the FSCQ file system. In: Miller, E.L., Hand, S. (eds.) *Proceedings of the 25th Symposium on Operating Systems Principles, SOSP 2015*, Monterey, CA, USA, October 4-7, 2015. pp. 18–37. ACM (2015). <https://doi.org/10.1145/2815400.2815402>
9. Cho, K., Lee, S.H., Raad, A., Kang, J.: Revamping hardware persistency models: view-based and axiomatic persistency models for Intel-x86 and Armv8. In: Freund, S.N., Yahav, E. (eds.) *PLDI '21: 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, Virtual Event, Canada, June 20-25, 2021. pp. 16–31. ACM (2021). <https://doi.org/10.1145/3453483.3454027>
10. Condit, J., Nightingale, E.B., Frost, C., Ipek, E., Lee, B., Burger, D., Coetzee, D.: Better I/O through byte-addressable, persistent memory. In: *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*. pp. 133–146. SOSP '09, ACM, New York, NY, USA (2009). <https://doi.org/10.1145/1629575.1629589>
11. Dalvandi, S., Doherty, S., Dongol, B., Wehrheim, H.: Owicki-Gries reasoning for C11 RAR. In: Hirschfeld, R., Pape, T. (eds.) *34th European Conference on Object-Oriented Programming, ECOOP 2020*, November 15-17, 2020, Berlin, Germany (Virtual Conference). *LIPICs*, vol. 166, pp. 11:1–11:26. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). <https://doi.org/10.4230/LIPICs.ECOOP.2020.11>
12. Dalvandi, S., Doherty, S., Dongol, B., Wehrheim, H.: Owicki-Gries reasoning for C11 RAR (artifact). *Dagstuhl Artifacts Ser.* **6**(2), 15:1–15:2 (2020). <https://doi.org/10.4230/DARTS.6.2.15>

13. Dalvandi, S., Dongol, B., Doherty, S., Wehrheim, H.: Integrating Owicki-Gries for C11-style memory models into Isabelle/HOL. *J. Autom. Reason.* **66**(1), 141–171 (2022). <https://doi.org/10.1007/s10817-021-09610-2>
14. Derrick, J., Doherty, S., Dongol, B., Schellhorn, G., Wehrheim, H.: Verifying correctness of persistent concurrent data structures. In: ter Beek, M.H., McIver, A., Oliveira, J.N. (eds.) *Formal Methods - The Next 30 Years - Third World Congress, FM 2019, Porto, Portugal, October 7-11, 2019, Proceedings*. Lecture Notes in Computer Science, vol. 11800, pp. 179–195. Springer (2019). https://doi.org/10.1007/978-3-030-30942-8_12
15. Derrick, J., Doherty, S., Dongol, B., Schellhorn, G., Wehrheim, H.: Verifying correctness of persistent concurrent data structures: a sound and complete method. *Formal Aspects Comput.* **33**(4-5), 547–573 (2021). <https://doi.org/10.1007/s00165-021-00541-8>
16. Doherty, S., Dongol, B., Wehrheim, H., Derrick, J.: Verifying C11 programs operationally. In: Hollingsworth, J.K., Keidar, I. (eds.) *Proceedings of the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2019, Washington, DC, USA, February 16-20, 2019*. pp. 355–365. ACM (2019). <https://doi.org/10.1145/3293883.3295702>
17. Friedman, M., Herlihy, M., Marathe, V.J., Petrank, E.: A persistent lock-free queue for non-volatile memory. In: Krall, A., Gross, T.R. (eds.) *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2018, Vienna, Austria, February 24-28, 2018*. pp. 28–40. ACM (2018). <https://doi.org/10.1145/3178487.3178490>
18. Gorjiara, H., Xu, G.H., Demsky, B.: Jaaru: efficiently model checking persistent memory programs. In: Sherwood, T., Berger, E.D., Kozyrakis, C. (eds.) *ASPLOS '21: 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Virtual Event, USA, April 19-23, 2021*. pp. 415–428. ACM (2021). <https://doi.org/10.1145/3445814.3446735>
19. Intel Corporation: Intel 64 and IA-32 Architectures Optimization Reference Manual (2021), <https://software.intel.com/content/dam/develop/external/us/en/documents-tps/64-ia-32-architectures-optimization-manual.pdf>
20. Izraelevitz, J., Mendes, H., Scott, M.L.: Linearizability of persistent memory objects under a full-system-crash failure model. In: Gavoille, C., Ilcinkas, D. (eds.) *Distributed Computing - 30th International Symposium, DISC 2016, Paris, France, September 27-29, 2016*. Proceedings. Lecture Notes in Computer Science, vol. 9888, pp. 313–327. Springer (2016). https://doi.org/10.1007/978-3-662-53426-7_23
21. Kaiser, J., Dang, H.H., Dreyer, D., Lahav, O., Vafeiadis, V.: Strong logic for weak memory: Reasoning about release-acquire consistency in Iris. In: *ECOOP (2017)*
22. Kang, J., Hur, C., Lahav, O., Vafeiadis, V., Dreyer, D.: A promising semantics for relaxed-memory concurrency. In: Castagna, G., Gordon, A.D. (eds.) *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*. pp. 175–189. ACM (2017). <https://doi.org/10.1145/3009837.3009850>
23. Khyzha, A., Lahav, O.: Taming x86-TSO persistency. *Proc. ACM Program. Lang.* **5**(POPL), 1–29 (2021). <https://doi.org/10.1145/3434328>
24. Khyzha, A., Lahav, O.: Abstraction for crash-resilient objects. In: *Programming Languages and Systems*. Springer International Publishing, Cham (2022)
25. Kokologiannakis, M., Kaysin, I., Raad, A., Vafeiadis, V.: Persevere: Persistency semantics for verification under ext4. *Proc. ACM Program. Lang.* **5**(POPL) (jan 2021). <https://doi.org/10.1145/3434324>

26. Lahav, O., Vafeiadis, V.: Owicki-Gries reasoning for weak memory models. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) *Automata, Languages, and Programming*. pp. 311–323. Springer, Berlin, Heidelberg (2015)
27. Lamport, L.: How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers* **28**(9), 690–691 (Sep 1979). <https://doi.org/10.1109/TC.1979.1675439>
28. Nipkow, T., Prensa Nieto, L.: Owicki/Gries in Isabelle/HOL. In: Finance, J. (ed.) *FASE. Lecture Notes in Computer Science*, vol. 1577, pp. 188–203. Springer (1999). https://doi.org/10.1007/978-3-540-49020-3_13
29. Ntzik, G., da Rocha Pinto, P., Gardner, P.: Fault-tolerant resource reasoning. In: Feng, X., Park, S. (eds.) *APLAS. Lecture Notes in Computer Science*, vol. 9458, pp. 169–188. Springer (2015). https://doi.org/10.1007/978-3-319-26529-2_10
30. Owicki, S.S., Gries, D.: An axiomatic proof technique for parallel programs I. *Acta Informatica* **6**, 319–340 (1976). <https://doi.org/10.1007/BF00268134>
31. Raad, A., Lahav, O., Vafeiadis, V.: Persistent Owicki-Gries reasoning: a program logic for reasoning about persistent programs on Intel-x86. *Proc. ACM Program. Lang.* **4**(OOPSLA), 151:1–151:28 (2020). <https://doi.org/10.1145/3428219>
32. Raad, A., Maranget, L., Vafeiadis, V.: Extending Intel-X86 consistency and persistency: Formalising the semantics of Intel-X86 memory types and non-temporal stores. *Proc. ACM Program. Lang.* **6**(POPL) (jan 2022). <https://doi.org/10.1145/3498683>
33. Raad, A., Vafeiadis, V.: Persistence semantics for weak memory: Integrating epoch persistency with the TSO memory model. *Proc. ACM Program. Lang.* **2**(OOPSLA) (oct 2018). <https://doi.org/10.1145/3276507>
34. Raad, A., Wickerson, J., Neiger, G., Vafeiadis, V.: Persistency semantics of the Intel-x86 architecture. *Proc. ACM Program. Lang.* **4**(POPL), 11:1–11:31 (2020). <https://doi.org/10.1145/3371079>
35. Raad, A., Wickerson, J., Vafeiadis, V.: Weak persistency semantics from the ground up: Formalising the persistency semantics of ARMv8 and transactional models. *Proc. ACM Program. Lang.* **3**(OOPSLA) (oct 2019). <https://doi.org/10.1145/3360561>
36. Sewell, P., Sarkar, S., Owens, S., Nardelli, F.Z., Myreen, M.O.: x86-TSO: A rigorous and usable programmer’s model for x86 multiprocessors. *Commun. ACM* **53**(7), 89–97 (Jul 2010). <https://doi.org/10.1145/1785414.1785443>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

