

Reasoning Algebraically about Refinement on TSO Architectures

Brijesh Dongol¹, John Derrick¹, and Graeme Smith²

¹ Department of Computing, University of Sheffield, UK

² School of Information Technology and Electrical Engineering,
The University of Queensland, Australia

Abstract. The Total Store Order memory model is widely implemented by modern multicore architectures such as x86, where local buffers are used for optimisation, allowing limited forms of instruction reordering. The presence of buffers and hardware-controlled buffer flushes increases the level of non-determinism from the level specified by a program, complicating the already difficult task of concurrent programming. This paper presents a new notion of refinement for weak memory models, based on the observation that pending writes to a process' local variables may be treated as if the effect of the update has already occurred in shared memory. We develop an interval-based model with algebraic rules for various programming constructs. In this framework, several decomposition rules for our new notion of refinement are developed. We apply our approach to verify the spinlock algorithm from the literature.

1 Introduction

Logics for reasoning about concurrency in shared memory systems are based on the assumption that hardware is *sequentially consistent* [18], guaranteeing that instructions within each process are never executed out of order in memory. However, modern processors have abandoned sequential consistency in favour of weaker memory guarantees, using local buffers to offer greater scope for optimisation. There are several different weak memory models [1, 2, 23]; in this paper, we focus on the most restricted of these: the *Total Store Order* (TSO) memory model, which is implemented by architectures such as x86 (see Fig. 1). Under TSO, instead of committing writes immediately to main memory, the process executing the write stores it as a pending write in its local buffer. Pending writes are not visible to other processors until they are *flushed*, which commits the write to shared memory. A flush is either programmer controlled (via instructions such as `fence` or `lock`) or hardware controlled. Programmer-controlled flushes are ultimately expensive (and inefficient), hence, one would like to keep these to a minimum. On the other hand, reasoning about hardware-controlled flushes is difficult due to the increase in non-determinism of a program's behaviour.

Several approaches to program verification under TSO have been developed; we provide a brief survey. Researchers have considered direct methods, such as executable memory models [22], theorems for reduction [9], and identification of race conditions [20]. Others have linked programs under TSO executions to an abstract specification using linearizability [7, 16], however, these use abstract specifications different from the




Fig. 1. TSO hardware overview

natural abstractions one would expect; [7] requires buffers to be present in the abstract specification, while [16] uses a non-deterministic abstract specification.

An issue with many existing approaches is that program semantics is given at a low level of abstraction of individual read and writes, which means programs must be understood and analysed using a verbose representation. Our work is based on the desire to lift reasoning to higher levels of abstraction [15], which in turn improves scalability. To this end, we develop an interval-based semantics by adapting Interval Temporal Logic [19]. Such an approach has two distinct advantages: (a) it allows one to define *truly concurrent* executions [10, 11], providing a more accurate model of TSO-based hardware; and (b) it is amenable to algebraic reasoning [3, 13], which enables one to develop *algebraic laws* for syntactically manipulating formulas representing program behaviour. In this paper, we develop algebraic rules to verify *refinement* between a concrete program and its abstract representation. The development of algebraic laws is non-trivial. However, once available, they provide high-level reusable theories for verification. We do not claim to have a complete set of laws (this is a topic of future work), instead, we provide a set of rules that are required for proving the spinlock example we verify.

Within this interval-based logic, we develop a framework for reasoning on TSO, simplifying our existing semantics [15] and introducing enhancements specifically designed for reasoning about buffer-based programs. This includes a simplified permission framework (Section 3.1), a novel methodology for evaluating expressions in the presence of local buffers (Section 4.3) and a novel notion of *local buffer refinement* (Section 5.2). Local buffer refinement is based on the observation that: To show a command C refines another command A with respect to a process p , the pending writes to local variables of p may be treated as if they have already taken effect in A . Thus, local updates at the concrete level may be treated as if they occur in their program order (without waiting for their flush to occur). This benefits verification because the non-determinism from flushes of local variables is resolved earlier. We develop a number of algebraic transformation laws for both refinement and local buffer refinement.

2 Background

2.1 Total Store Order Example

Total Store Order (TSO) memory allows a process to store a write in its local buffer and continue processing without waiting for this write to be committed to memory (i.e., while the write is pending). The values in the buffer are flushed in a FIFO order. To see

```

word x = 1;
void acquire() {
a1 while(1) {
a2   lock;
a3   if (x = 1) {
a4     x := 0;
a5     unlock;
a6     return }
a7   unlock;
a8   while(x = 0);
} }

void release() {
r1 x := 1; }

bool tryacquire() {
t1 lock;
t2 if (x = 1) {
t3   x := 0;
t4   unlock;
t5   return true }
t6 unlock;
t7 return false; }

```

Fig. 2. Spinlock algorithm

```

Noncritical section ;
acquire();
Critical section ;
release();

Fig. 3. Spinlock client (a)

Noncritical section;
if tryacquire() {
  Critical section ;
  release() ;
}

```

Fig. 4. Spinlock client (b)

the effect of this, consider the following classic example with processes p and q that modify shared variables x and y , which are initialised to 0. In this paper, we assume maximum parallelism and that each thread resides in exactly one core, therefore, the words *process* and *core* are used synonymously.

```

word x=0, y=0;
p { p1: x := 1 ;
    p2: r1 := y }
q { q1: y := 1 ;
    q2: r2 := x }

```

Under sequentially consistent memory, at the end of execution, at least one of $r1$ or $r2$ would have a value 1. However, in TSO memory, it is possible to end execution so that both $r1$ and $r2$ read the original values of x and y , i.e., both $r1$ and $r2$ are 0 at termination. One such execution is $\langle p1, p2, q1, q2, \text{flush}(p), \text{flush}(p), \text{flush}(q), \text{flush}(q) \rangle$, where $\text{flush}(p)$ denotes a (hardware-controlled) flush event for process p . The write to x at $p1$ is not seen by process q until p 's buffer is flushed, and symmetrically for the write to y at $q1$. Hence, it is possible for q to read a value 0 for x at $q2$ even though $q2$ is executed after $p1$.

In addition to the above behaviour, each TSO process reads pending writes from its own buffer if possible, and hence, may obtain values that are not yet globally visible to other processes, e.g., if $p2$ is replaced with $r1 := x$, process p would read $x = 1$ even if the write to x is pending. If there are multiple pending writes to the same location, then the write value corresponding to the last pending write is returned.

2.2 Case Study: Spinlock

Spinlock [4] is a locking mechanism designed to avoid operating system overhead associated with process scheduling and context switching. A typical implementation of spinlock is shown in Fig. 2, where a global variable x represents the lock and is set to 0 when the lock is held by a process, and 1 otherwise. The lock x is itself acquired using a secondary hardware lock (see Fig. 1), and this hardware lock is acquired/released using lock/unlock instructions. A process trying to acquire the lock x *spins*, i.e., waits in a loop and repeatedly checks the lock for availability.

Operation `acquire` only terminates if it successfully obtains the lock x . It will first lock the hardware so that no other process can access x . If, another process has already acquired x (i.e., $x = 0$) then it will release the hardware lock at `a7` and spin at `a8`, i.e., loop in the while-loop until x becomes free, before starting over from `a2`. Otherwise, it acquires the lock at `a4` by setting x to 0 , releases the hardware lock at `a5` and returns at `a6`. The operation `release` releases the lock by setting x to 1 . The operation `tryacquire` is similar to `acquire`, but unlike `acquire` it only makes one attempt to acquire the lock. If this is successful it returns `true`, otherwise it returns `false`. Under TSO, a process p executing an assignment (e.g., $x := 0$) places a pending write in p 's local buffer, which is not visible to other processes until the buffer is flushed.

We refer to processes that use spinlock to provide mutual exclusion to a critical section of code as its *clients*. Here, as in [22], we assume that clients of the spinlock behave either as the program in Fig. 3 or Fig. 4. Thus, one can assume that a client only calls a `release` operation when it holds the lock.³ Note however, that the behaviours in Fig. 3 and Fig. 4 are not exhaustive. To admit other behaviours, one may formalise the additional client code, then apply our proof methods in this paper to verify this additional behaviour.

Clients can ensure mutual exclusion in the critical sections if in place of `acquire`, `release` and `tryacquire`, they use *abstract operations* `AAcq`, `ARel` and `ATry` below, respectively, which do not use buffers. We will refer to such clients as *abstract clients*.

```

word x = 1;
void AAcq() {
  await (x = 1) {
    x <== 0
  }
}
void ARel() {
  x <== 1
}
bool ATry() {
  if CAS(x, 1, 0) {
    return true
  } else return false }

```

Here, statement `await` denotes a blocking atomic test-and-set statement, e.g., `AAcq()` can only execute if $x = 1$ holds, and its execution atomically sets the value of x to 0 . Unlike the concrete program in Fig. 2, all reads and writes occur directly with main memory; we use assignments of the form $x <== 0$ (which directly updates the value of x in memory) to distinguish this difference. If $x = 0$, then `AAcq()` blocks and cannot execute further until its guard $x = 1$ is set to true by another process. Operation `ATry()` attempts to update x to 0 using a non-blocking atomic compare-and-swap operation CAS, and returns 1 if the operation is successful and 0 , otherwise.⁴

Our notion of correctness of the spinlock will be to show that every possible execution of a spinlock client is a possible execution of an abstract client. To this end, we prove *refinement* between the behaviour of the two executions (see Section 5.3). Proving refinement under TSO is difficult; one must not only verify concurrency effects, but additionally consider the effect of accessing the buffer during a program's execution. Furthermore, the level of atomicity at which these effects are visible is fine-grained,

³ Such restrictions on client behaviour must be made due to the simplicity of the spinlock algorithm in Fig. 2. Arbitrary client behaviour e.g., two consecutive calls to `release` without acquiring the lock will result in incorrect behaviour under TSO.

⁴ CAS(a, b, c) is equivalent to `atomic { if a = b then a := c ; return true else return false}`.

occurring at the level of individual reads and writes. This paper develops a high-level approach for proving refinement that avoids the need to consider low-level (fine-grained) effects whenever possible by developing an interval-based semantics for programs under TSO. This allows one to view the concurrent execution of two processes as the conjunction of their behaviours over an interval (as opposed to an interleaving of their traces), reducing the impact of non-determinism due to concurrency.

3 Interval-Based Reasoning

3.1 Permission Monitoring

Using an interleaved execution semantics, one can guarantee that a variable will not be simultaneously written, or read and written as part of the same transition. This is not true for shared-memory true concurrency, where one must model variable access by the different processes (e.g., two processes simultaneously modifying variable x in Fig. 2).

Our solution is to explicitly define read/write permissions. To this end, we assume that programs are executed by processes from a set $Proc$; each process represents a concurrent thread which modifies a set of variables from a set Var . The TSO architecture uses sophisticated coherence protocols to provide an illusion of shared memory. One may assume the following about read and write instructions:

- Two simultaneous writes (by different processes) to the same variable do not occur.
- A simultaneous read and write of the same variable does not occur.
- A process never has access permission to the local variable of another process.

As we shall see in Sections 4.2 and 4.4, permissions also provide a convenient mechanism for formalising the effect of a `lock-unlock` block.

In previous work [11], we have modelled permissions using a fractional encoding (inspired by [5]). Here, we simplify these general notions and define the *permission space* as $Perm \hat{=} Proc \rightarrow Var \rightarrow \mathbb{P}\{wr, rd\}$, where wr and rd denote write and read permission, respectively. Using ‘.’ for function application, given $\pi \in Perm$, we interpret $wr \in \pi.p.v$ (resp., $rd \in \pi.p.v$) as $p \in Proc$ has permission to write to (resp., read the value of) $v \in Var$.

A system at any time is described by a state of type $State \hat{=} Var \rightarrow Val$, where Val is the set of values. The system over time is formalised by a *stream*, which is a total function of type $Stream \hat{=} \mathbb{Z} \rightarrow (State \times Perm)$. Therefore, for each time in \mathbb{Z} , a stream formalises the state of the system and the permissions for each process and variable. Properties of a system are given by *predicates*; a predicate of type T is a member of $\mathcal{PT} \hat{=} T \rightarrow \mathbb{B}$, e.g., $\mathcal{P}State$, $\mathcal{P}Stream$, and $\mathcal{P}Perm$ are state, stream, and permission predicates, respectively. We assume pointwise lifting of boolean operators over predicates in the normal manner.

If $\pi \in Perm$, then $\mathcal{W}.p.v.\pi \hat{=} (wr \in \pi.p.v)$, $\mathcal{R}.p.v.\pi \hat{=} (rd \in \pi.p.v)$ and $\mathcal{N}.p.v.\pi \hat{=} (\pi.p.v = \emptyset)$ denote permission predicates that hold iff process p has write, read or no access to v in the permission space π , respectively.

Example 1. Suppose $Var = \{u, v\}$, $Proc = \{p, q\}$ and

$$\pi \hat{=} \{p \mapsto \{u \mapsto \{rd, wr\}, v \mapsto \{rd\}\}, q \mapsto \{u \mapsto \emptyset, v \mapsto \{rd\}\}\}$$

Then, $\mathcal{W}.p.u.\pi$ (p has write permission to u), $(\mathcal{R}.p \wedge \mathcal{R}.q).v.\pi$ (both p and q have read permission to v) and $\mathcal{N}.q.u.\pi$ (q has no permission to u) in space π . Note that due to pointwise lifting $(\mathcal{R}.p \wedge \mathcal{R}.q).v.\pi = \mathcal{R}.p.v.\pi \wedge \mathcal{R}.q.v.\pi$. \square

The assumptions on reads and writes above are then taken into account by assuming that each *valid* permission space π satisfies the following, where p and q such that $p \neq q$ are processes, v is a variable, and u_p is a local variable of process p .

$$(\mathcal{W}.p.v.\pi \Rightarrow \mathcal{N}.q.v.\pi) \wedge (\mathcal{R}.p.v.\pi \Rightarrow \neg \mathcal{W}.q.v.\pi) \wedge \pi.p.u_p = \{\text{wr}, \text{rd}\} \quad (1)$$

Note that the third conjunct combined with the first ensures that q does not have read nor write permission to any local variable of p . This is lifted to the level of streams by defining a *valid* stream to be one in which each $s.t$ is valid for $t \in \mathbb{Z}$. For the rest of the paper, we assume each stream is valid.

To simplify the notation, for a state predicate b and permission predicate z , we assume $b.(\sigma, \pi) = b.\sigma$ and $z.(\sigma, \pi) = z.\pi$, where σ and π are a state and a permission state, respectively. We assume ‘ \uparrow ’ denotes a projection operator, e.g., $(x, y) \uparrow 1 = x$.

3.2 Interval Predicates

In this section, we provide the basics of interval predicates, which forms the logical basis of our program semantics. Our logic is an adaptation of Interval Temporal Logic [19]. An *interval* is a contiguous set of integers (denoted \mathbb{Z}), and hence the set of all intervals is $\text{Intv} \hat{=} \{\Delta \subseteq \mathbb{Z} \mid \forall t_1, t_2: \Delta \bullet \forall t: \mathbb{Z} \bullet t_1 \leq t \leq t_2 \Rightarrow t \in \Delta\}$.

We let $\text{lub}.\Delta$ and $\text{glb}.\Delta$ denote the *least upper* and *greatest lower* bounds of an interval Δ , respectively. Furthermore, we define $\text{inf}.\Delta \hat{=} (\text{lub}.\Delta = \infty)$, $\text{fin}.\Delta \hat{=} \neg \text{inf}.\Delta$, and $\varepsilon.\Delta \hat{=} (\Delta = \emptyset)$. We define an ordering $\Delta_1 < \Delta_2 \hat{=} \forall t_1, t_2: \Delta_2 \bullet t_1 < t_2$. To facilitate reasoning about specific parts of a stream, we use *interval predicates*, which have type $\text{IntvPred} \hat{=} \text{Intv} \rightarrow \mathcal{P}\text{Stream}$ [11, 13].

Example 2. Given *Var*, *Proc* and π as defined in Example 1, we define

$$\begin{aligned} \sigma_1 &\hat{=} \{u \mapsto 500, v \mapsto 42\} & s &\hat{=} \lambda t \bullet \text{if } t \geq 10 \text{ then } (\sigma_1, \pi) \text{ else } (\sigma_2, \pi) \\ \sigma_2 &\hat{=} \{u \mapsto 0, v \mapsto 1\} & g &\hat{=} \lambda \Delta \bullet \lambda s \bullet \forall t: \Delta \bullet ((s.t) \uparrow 1).u \geq 300 \\ b &\hat{=} \lambda \sigma \bullet \sigma.u < \sigma.v \end{aligned}$$

Then, σ_1, σ_2 are states, b a state predicate, s is a stream and g is an interval predicate. Each of $(\neg b).(\sigma_1, \pi)$, $b.(\sigma_2, \pi)$, $\neg g.[-3, 3].s$ and $g.[10, 100].s$ hold.⁵ \square

We define *universal implication* $g_1 \Rightarrow g_2 \hat{=} \forall \Delta: \text{Intv}, s: \text{Stream} \bullet g_1.\Delta.s \Rightarrow g_2.\Delta.s$ for interval predicates g_1 and g_2 , and write $g_1 \equiv g_2$ iff both $g_1 \Rightarrow g_2$ and $g_2 \Rightarrow g_1$ hold.

4 Concurrent Programming with Intervals

4.1 Operators to Model Programming Constructs

In this section, we introduce interval predicate operators used to formalise common programming constructs: sequential composition, branching, loops, and parallel composition. To model sequential composition, we define the *chop* operator [19, 13]. Unlike

⁵ Here, $[-3, 3]$ is the closed interval from -3 to 3 and $[10, 100)$ is the right-open interval from 10 to 100 .

Interval Temporal Logic, which requires adjoining intervals to overlap at a single point, adjoining intervals in our logic are disjoint.

$$(g_1 ; g_2).\Delta.s \hat{=} (\exists \Delta_1, \Delta_2 \bullet (\Delta_1 \cup \Delta_2 = \Delta) \wedge \Delta_1 < \Delta_2 \wedge g_1.\Delta_1.s \wedge g_2.\Delta_2.s) \vee (\text{inf} \wedge g_1).\Delta.s$$

Thus, $(g_1 ; g_2).\Delta.s$ holds iff either interval Δ may be split into two adjoining parts Δ_1 and Δ_2 so that g_1 holds for Δ_1 and g_2 holds for Δ_2 in s , or the least upper bound of Δ is ∞ and g_1 holds for Δ in s . Inclusion of the second disjunct $(\text{inf} \wedge g_1).\Delta.s$ enables g_1 to model an infinite (divergent or non-terminating) program. We assume that ‘;’ binds tighter than all other binary operators, e.g., $g_1 ; g_2 \vee h = (g_1 ; g_2) \vee h$.

Example 3. For b , g and s as defined in Example 2, if $h \hat{=} \lambda \Delta \bullet \lambda s \bullet \exists t: \Delta \bullet b.(s.t)$, then $(h ; g).[0, 100).s$ holds because both $h.[0, 10).s$ and $g.[10, 100).s$ hold. Note that there may be more than one possible way to split up an interval when applying the definition of chop. \square

Non-deterministic choice is modelled by (lifted) disjunction, and hence, for example, the behaviour of `if b then S1 else S2` is given by $\text{test}.b ; \text{beh}.S1 \vee \text{test}.(\neg b) ; \text{beh}.S2$, where $\text{test}.b$ and $\text{beh}.S1$ are interval predicate formalisations of evaluating b and executing $S1$, respectively. The precise value of $\text{test}.b$ depends on the atomicity assumptions of the program under consideration, and hence, the interpretation of $\text{test}.b$ is non-trivial (see [17, 13, 11]). The value of $\text{test}.b$ is modelled by command $[b]$ (see Section 4.2), whereas at the concrete level its value is formalised by a different command $\llbracket b \rrbracket$, which takes the effect of the buffer into account (see Section 4.3).

Iteration g^* and g^ω are the least and greatest fixed points of $\lambda z \bullet gz \vee \varepsilon$, respectively [14], where g^* allows empty and finite iterations and g^ω allows empty, finite and infinite iterations of g . We also define strictly finite and possibly infinite *positive* iterations.

$$\begin{aligned} g^* &\hat{=} \mu z \bullet ((g ; z) \vee \varepsilon) & g^+ &\hat{=} g ; g^* \\ g^\omega &\hat{=} \nu z \bullet ((g ; z) \vee \varepsilon) & g^{\omega+} &\hat{=} g ; g^\omega \end{aligned}$$

A thorough algebraic treatment of loops using iteration is given in [3]. For example, program code `while b do S` is modelled by $(\text{test}.b ; \text{beh}.S)^\omega ; \text{test}.(\neg b)$. In this paper, we use the following rule.

Law 1 (Leapfrog [3]) *For interval predicates g and h , both $g ; (h ; g)^\omega \equiv (g ; h)^\omega ; g$ and $g ; (h ; g)^* \equiv (g ; h)^* ; g$ hold.*

We are interested in modelling true concurrency and therefore simply treat the parallel composition of two or more processes using (lifted) logical conjunction. For example, the behaviour of $g_1 ; g_2$ in parallel with $h_1 ; h_2$ over an interval Δ in stream s is given by $(g_1 ; g_2 \wedge h_1 ; h_2).\Delta.s$. Using pointwise lifting, this is equivalent to $(g_1 ; g_2).\Delta.s \wedge (h_1 ; h_2).\Delta.s$, which holds iff (a) Δ can be split into adjoining intervals Δ_1 and Δ_2 such that $g_1.\Delta_1.s \wedge g_2.\Delta_2.s$ holds; and (b) Δ can also be split into adjoining intervals Δ_3 and Δ_4 such that $h_1.\Delta_3.s \wedge h_2.\Delta_4.s$. Note that there is no immediate correlation between the lengths of Δ_1 and Δ_3 , i.e. g_1 could terminate earlier than h_1 , and vice versa.

Modelling tests. Interval predicates provide a flexible approach to non-deterministic state predicate evaluation [17], where expression evaluation is assumed to take time (as opposed to being instantaneous). In this paper, guards and assignments are restricted

to contain at most one shared variable.⁶ Given that c is either a state or permission predicate, and Δ and s are an interval and stream, respectively, we define:

$$(\Box c).\Delta.s \hat{=} \neg\varepsilon.\Delta \wedge \forall t: \Delta \bullet c.(s.t)$$

Thus, $(\Box c).\Delta.s$ holds iff Δ is non-empty and c holds for each state of s within Δ . For example, $\neg\varepsilon \wedge g \equiv \Box(u \geq 300)$, where g is the interval predicate defined in Example 2.

Reasoning about pre/post assertions. One may define several additional interval predicate operators [13]. For the purposes of this paper, we find it useful to reason about properties that hold in the immediately preceding interval. We therefore define

$$(\Theta g).\Delta.s \hat{=} \neg\varepsilon.\Delta \wedge \text{glb}.\Delta \neq -\infty \wedge g.(\text{prev}.\Delta).s$$

where $\text{prev}.\Delta \hat{=} \{t: \mathbb{Z} \mid \forall u: \Delta \bullet t < u\}$ is the interval of all times before Δ . If c is a state or permission predicate, we use notation $\vec{c} \hat{=} \text{true}$; $\Box c$, where $\vec{c}.\Delta$ states that c holds at the end of Δ whenever $\text{inf}.\Delta \neq \infty$. Additionally, we define the following notation to reason about assertions that immediately precede, or are a result of a computation.

$$\{c\}g \hat{=} \Theta \vec{c} \wedge g \qquad g\{c\} \hat{=} g \wedge \vec{c}$$

Such a definition of a pre-assertion is necessary because we assume adjoining intervals do not overlap (unlike [19]). We have the following useful properties, which can be proved in a straightforward manner.

$$\{c\}(g_1 \vee g_2) \equiv (\{c\}g_1) \vee (\{c\}g_2) \tag{2}$$

$$g_1\{c\}; g_2 \equiv g_1; \{c\}g_2 \qquad \text{provided } g_1 \vee g_2 \Rightarrow \neg\varepsilon \tag{3}$$

4.2 Abstract Commands

Using the interval-based semantic basis from the previous sections, we formalise *commands*, which describe the behaviours of the system processes. Formally, a command is of type $\text{Cmd} \hat{=} \mathbb{P}_1 \text{Proc} \rightarrow \text{IntvPred}$, mapping non-empty sets of processes to an interval predicate representing their behaviour. We use $C.p$ as shorthand for $C.\{p\}$, where C is a command and p is a process.

The semantics of sequential composition, iteration, non-deterministic choice and parallel composition of commands are defined pointwise lifting of the interval predicate operators, and hence, are given in the same syntax, e.g., $(C_1; C_2).p = C_1.p; C_2.p$. What remains is to define the commands to model, say, guard evaluation and assignment.

We first present some basic commands that may be used to model the abstract (sequentially consistent) specification. In particular, we define *idling* (denoted id), *abstract guard evaluation* (denoted $[b]$), *memory update* (denoted $\bar{v} \Leftarrow \bar{e}$) and *locked access* (denoted $\bar{v} \bullet C$), where \bar{v} and \bar{e} denote vectors of variables and expressions, respectively. We define $\text{Deny}.\bar{v}.p \hat{=} \Box(\forall q: \text{Proc} \setminus \{p\}, v: \bar{v} \bullet (\neg\mathcal{W} \wedge \neg\mathcal{R}).q.v)$, which states that the variables in \bar{v} are not accessed by processes other than p . We assume that $\text{vars}.b$ denotes the free variables in b .

$$\begin{aligned} \text{id}.p &\hat{=} \forall v: \text{Var} \bullet \Box \neg\mathcal{W}.p.v & (\bar{v} \Leftarrow \bar{e}).p &\hat{=} \exists \bar{k} \bullet \{\bar{e} = \bar{k}\} \\ \text{id}.p &\hat{=} \neg\varepsilon \Rightarrow \text{id}.p & & \Box(\bar{v} = \bar{k}) \wedge \Box(\forall v: \bar{v} \bullet \mathcal{W}.p.v) \\ [b].p &\hat{=} \Box b \wedge \text{id}.p \wedge & \bar{v} \bullet C.p &\hat{=} (\bar{v} \neq \emptyset \Rightarrow \text{Deny}.\bar{v}.p) \wedge C.p \wedge \\ &\forall v: \text{vars}.b \bullet \Box \mathcal{R}.p.v & & (\forall v: \bar{v} \bullet \Box(\mathcal{W} \wedge \mathcal{R}).p.v) \end{aligned}$$

⁶ This can be extended to handle more complex expressions [17].

Thus $\text{nid}.p$ states that p is *write idle* i.e., p does not have write access to any variable during the given (non-empty) interval; $\text{id}.p$ states that either p is write idle or the interval under consideration is empty; $[b].p$ holds iff b holds throughout the given interval and p is idle; $\bar{v} \Leftarrow \bar{e}$ denotes an instantaneous update, which holds iff \bar{e} evaluates to a vector of values \bar{k} as a pre-assertion, and \bar{v} is updated to \bar{k} , where p has write permission to each $v \in \bar{v}$; and $\boxed{\bar{v} \bullet C}.p$ holds iff $C.p$ holds and additionally no process other than p has permission to access $v \in \bar{v}$.

Example 4. The abstract specification is formalised as follows, where AAcq and ARel specify operations AAcq and ARel , respectively, while ATryOK and ATryFl specify execution of the ATry operation that succeed and fail to acquire the lock, respectively. We abbreviate $\mathbf{x} = 1$ and $\mathbf{x} = 0$ to x and $\neg x$, respectively. The return value of an execution of tryacquire in process p is modelled by a local variable r_p .

$$\begin{aligned} \text{AAcq}.p &\hat{=} \boxed{x \bullet [x]; (x \Leftarrow 0)} & \text{ATryOK}.p &\hat{=} \text{AAcq}.p ; \text{id} ; (r_p \Leftarrow \text{true}) \\ \text{ARel}.p &\hat{=} x \Leftarrow 1 & \text{ATryFl}.p &\hat{=} \boxed{x \bullet [\neg x]} ; \text{id} ; (r_p \Leftarrow \text{false}) \\ \text{AExec}.p &\hat{=} ((\text{AAcq} ; \text{id} ; \text{ARel}) \vee (\text{ATryOK} ; \text{id} ; \text{ARel}) \vee \text{ATryFl}).p \\ \text{Spec}.P &\hat{=} \{x\} \bigwedge_{p:P} ((\text{id} ; \text{AExec})^\omega ; \text{id}).p \end{aligned}$$

The concurrent execution of abstract clients is modelled by Spec , which begins in a state in which the lock x is available (i.e., x holds) and consists of a number of (truly) parallel processes. We assume that each client of the spinlock behaves as either Fig. 3 or Fig. 4, and furthermore, that the critical and non-critical sections do not modify variables x and r_p , and hence, both the critical and non-critical sections are modelled by id . Therefore, $\text{id} ; \text{AExec}$ models a single call to the abstract spinlock. Each process may make multiple (zero or more) calls, followed by no calls, and hence, all possible behaviours of an abstract client is given by $(\text{id} ; \text{AExec})^\omega ; \text{id}$.

We now explain how each operation is modelled. If $\text{AAcq}.p.\Delta.s$ holds for interval Δ and stream s , then only process p has access to x (i.e., no process $q \neq p$ may read or write to x) and either (i) Δ can be partitioned into Δ_1 and Δ_2 with $\Delta_1 < \Delta_2$ such that x holds in s throughout Δ_1 and x is updated to 0 in s within Δ_2 , or (ii) Δ is infinite and x (i.e., $\mathbf{x} = 1$) holds in s throughout Δ . Because $\text{await } b$ blocks until $\text{test}.b$ becomes true, there are no behaviours for $\text{AAcq}.p$ when $\text{test}.(\neg b)$ holds. Operation $\text{ARel}.p$ immediately sets x to 1, and by the definition of \Leftarrow together with assumption (1), we have that no other process reads or writes to x while this update occurs. Operation $\text{ATryOK}.p$ behaves as $\text{AAcq}.p$, performs some idling, then updates r_p to *true*. The idling between $\text{AAcq}.p$ and update to r_p provides scope for potential stuttering at the concrete level. Operation $\text{ATryFl}.p$ starts by behaving as $\boxed{x \bullet [\neg x]}$, which implies that x is not accessed by any process $q \neq p$ and that $\neg x$ holds throughout the given interval. Then, $\text{ATryFl}.p$ performs some idling and updates the return value r_p to *false*. \square

4.3 Reading Variables for Expression Evaluation with Buffer Effects

Section 4.2 provided an interval-based semantics for commands without buffers, which were in turn used to model the abstract specification. The concrete program executes under TSO memory and contains local buffers, whose effects on the program's behaviour

must be formalised. In this section, we present a method for evaluating expressions, i.e., when processes read variables, in the presence of local buffers. In particular, we formalise the fact that a TSO process first checks its buffer for pending writes; if a pending write exists, the last pending value is returned, and otherwise the value from memory is returned. Using interval-based methods enables one to formalise the effects of a buffer on the value of an expression at a high level of abstraction [15].

We assume that $B_p \in \text{Var}$ denotes the buffer for process p , whose value is of type $\text{seq.}(\text{Var} \times \text{Val})$, representing a pending write. Each buffer may contain multiple pending writes to the same location, and hence, we define a function cover that returns a set of mappings to the last pending write in a given buffer. Because $\text{seq.}X$ is a partial function of type $\mathbb{N} \rightarrow X$, we may use $\text{dom.}z$ to refer to the indices of $z \in \text{seq.}X$.

$$\text{cover.}B \hat{=} \{B.i \mid i: \text{dom.}B \wedge \forall j: \text{dom.}B \bullet j > i \Rightarrow (B.i \upharpoonright 1) \neq (B.j \upharpoonright 1)\}$$

When a process evaluates an expression, the values of pending writes in a process' buffer mask those in memory, which is modelled formally using functional override ' \oplus ' (see [24] for a formal definition).

Example 5. Suppose B and BB are buffers (BB is not shown below), p and q are processes, u, v , and w are variables, and σ is a state such that

$$B \hat{=} \langle (v, 11), (w, 33), (v, 44) \rangle \quad \sigma \hat{=} \{(B_p, B), (B_q, BB), (u, 0), (v, 1), (w, 2)\}$$

Then we have $\text{cover.}B = \{(w, 33), (v, 44)\}$, i.e., for each variable in B its last corresponding value in B is picked. Hence, we have

$$\sigma \oplus \text{cover.}B = \{(B_p, B), (B_q, BB), (u, 0), (v, 44), (w, 33)\}$$

which replaces each mapping in σ by those in $\text{cover.}B$. \square

We lift buffer effects to state predicates using $(\text{mask.}b.B).\sigma \hat{=} b.(\sigma \oplus \text{cover.}B)$, which states that b holds in a state σ covered by B . For the definitions in Example 5, both $(\text{mask.}(u = 0).B).\sigma$ and $(\text{mask.}(w < v).B).\sigma$ hold, but $(w < v).\sigma$ does not.

Processes evaluate state predicates (e.g., as part of a guard), however, in the presence of permissions and local buffers, evaluation is non-trivial. Firstly, one must ensure that a process p evaluating state predicate b is able to obtain read permission to each variable of b whenever the variable's value is fetched from memory. Note that this is only potentially problematic if the variable in question is shared (i.e., not a local variable of p) and not in p 's buffer (p may always access its local buffer). Secondly, the value of a variable v read by p must be the last value of v in p 's buffer if it exists, and the value of v in memory, otherwise. Assuming $\text{vars.}B \hat{=} \{(B.i) \upharpoonright 1 \mid i \in \text{dom.}B\}$ is the set of all variables in $B \in \text{seq.}(\text{Var} \times \text{Val})$, we define:

$$\Box_p b \hat{=} \Box(\text{mask.}b.B_p \wedge \forall v: \text{vars.}b \setminus \text{vars.}B_p \bullet \mathcal{R}.p.v)$$

Thus, $(\Box_p b).\Delta.s$ models the evaluation of state predicate b by the process p , by either reading variables from p 's buffer (if possible) or from main memory. Here $(\Box_p b).\Delta.s$ holds iff (i) b masked by B_p holds in s throughout Δ , and (ii) p has read permission to the free variables of b not in $\text{vars.}B_p$ throughout Δ in s .⁷ In Section 4.4, $(\Box_p b)$ is

⁷ In general, if b contains multiple shared variables, $\Box_p b$ is not an accurate model of evaluation because the variables in b may be read at different instants [17, 10, 12]. However, here, we assume that each expression/guard of each program under consideration contains at most one shared variable, in which case $\Box_p b$ is accurate (see [17, 10, 12]).

used to define expression evaluation, which in turn is used to model guards and the right hand side of assignments.

4.4 Commands under TSO

As already mentioned, processes that execute under TSO write only to their local buffers. The effects of these writes are not seen by other processes until a buffer is *flushed*, which moves the pending write from a buffer to shared memory. TSO buffers operate in a FIFO order, and hence, we define the following commands, where Φ models a single flush, $\tilde{\Phi}$ models a flush or a non-empty idle, and Φ_{all} models a complete buffer flush.

$$\begin{aligned} \Phi.p &\hat{=} \exists k \bullet \{B_p = k \wedge k \neq \langle \rangle\} \\ &B_p, (k.0 \uparrow 1) \Leftarrow tail.k, (k.0 \uparrow 2) \\ \tilde{\Phi} &\hat{=} \Phi \vee nid \\ \Phi_{all}.p &\hat{=} \Phi^+.p \{B_p = \langle \rangle\} \end{aligned}$$

Due to the fine-granularity of the concrete implementation, seemingly atomic statements become compound commands under TSO memory. Evaluation of a boolean expression b (e.g. a guard in an `if-then-else` block) is a compound statement that flushes or idles (zero or more times), evaluates b using the buffer-based evaluation semantics defined in Section 4.3, then flushes or idles again (zero or more times). A write of v with value k appends the pair (v, k) to the end of the local buffer. An assignment to a constant value k , potentially flushes or idles (zero or more times), appends the value to the buffer, then potentially flushes or idles (zero or more times). An assignment to a complex expression e , first evaluates the expression to a value k , then assigns k to v . Thus, we define:

$$\begin{aligned} \llbracket b \rrbracket.p &\hat{=} \tilde{\Phi}^*.p; (\Box_p b \wedge nid.p); \tilde{\Phi}^*.p \\ (v \leftarrow k).p &\hat{=} B_p \Leftarrow B_p \hat{\wedge} \langle (v, k) \rangle \\ v := e &\hat{=} \text{if } e \in Val \text{ then } \tilde{\Phi}^*; (v \leftarrow e); \tilde{\Phi}^* \text{ else } \exists k: Val \bullet \llbracket e = k \rrbracket; v := k \end{aligned}$$

There are several TSO instructions that force the entire buffer to be flushed. These additionally may lock certain variables from being accessed while the flush all is being executed. We therefore define commands $pre_{\Phi}.\bar{v}.C.p$ and $post_{\Phi}.\bar{v}.C.p$, where $pre_{\Phi}.\bar{v}.C.p$ flushes the entire buffer (locking \bar{v}) before C is executed (and similarly $post_{\Phi}.\bar{v}.C.p$ flushes all after C).

$$\begin{aligned} pre_{\Phi}.\bar{v}.C.p &\hat{=} \{B_p = \langle \rangle\}C.p \vee (\overline{\bar{v} \bullet \Phi_{all}}; C).p \\ post_{\Phi}.\bar{v}.C.p &\hat{=} C.p\{B_p = \langle \rangle\} \vee (C; \overline{\bar{v} \bullet \Phi_{all}}).p \end{aligned}$$

Some TSO instructions do not lock the memory while the buffer is being flushed. These may be modelled using $pre_{\Phi}.\emptyset.C$ and $post_{\Phi}.\emptyset.C$, which we abbreviate to $pre_{\Phi}.C$ and $post_{\Phi}.C$, respectively. A `lock` (e.g., `a2` in Fig. 2) acquires the memory lock then flushes the entire buffer; an `unlock` (e.g., `a5` in Fig. 2) flushes the entire buffer then releases the memory lock. Therefore, executing a command C within a `lock-unlock` block is modelled by

$$\overline{\bar{v} \bullet C}_{\Phi} \hat{=} pre_{\Phi}.\bar{v}.(post_{\Phi}.\bar{v}.\overline{\bar{v} \bullet C})$$

which executes C and ensures the buffer is empty before and after executing C . In addition it ensures that no reads and writes to \bar{v} by other processes occur while C is being executed. Note however, that if a process p executes $\overline{\bar{v} \bullet C}_{\Phi}$ and a process $q \neq p$ has a pending write to \bar{v} in its local buffer, then q may read this value of \bar{v} even while p is executing $\overline{\bar{v} \bullet C}_{\Phi}$.

Example 6. Our modelling notation is used to formalising the behaviour of the concrete implementation as follows, where $Lck.p \hat{=} \boxed{x \bullet \llbracket x \rrbracket ; (x := 0)}_{\Phi}$.

$$\begin{aligned}
Acq.p &\hat{=} \left(\boxed{x \bullet \llbracket \neg x \rrbracket}_{\Phi} ; \llbracket \neg x \rrbracket^{\omega} ; \llbracket x \rrbracket \right)^{\omega} ; Lck.p & Rel.p &\hat{=} x := 1 \\
TryOK.p &\hat{=} Lck.p ; (r_p := true) & TryFl.p &\hat{=} \boxed{x \bullet \llbracket \neg x \rrbracket}_{\Phi} ; (r_p := false) \\
Exec.p &\hat{=} (Acq ; id ; Rel \vee TryOK ; id ; Rel \vee TryFl).p \\
Prog.P &\hat{=} \{x\} \bigwedge_{p:P} (\{B_p = \langle \rangle\} (\tilde{\Phi}^+ ; Exec)^{\omega} ; post_{\Phi}.x.(\tilde{\Phi}^+)).p
\end{aligned}$$

Within $Acq.p$, command $\boxed{x \bullet \llbracket \neg x \rrbracket}_{\Phi}$ models an execution consisting of the lock at a2, failed test at a3, unlock at a7 (see Fig. 2). Command $\llbracket \neg x \rrbracket^{\omega} ; \llbracket x \rrbracket$ models the while loop at a8. Therefore, the outermost $^{\omega}$ iteration in $Acq.p$ models executions of the outermost loop of acquire that fail to acquire the lock. Command $Lck.p$ models the lock at a2, successful test at a3, assignment at a4, and unlock at a5 followed by the return at a6. The other operations are similar. \square

5 Refinement and Local Refinement for TSO

5.1 Interval-Based Refinement

In this section, we develop a theory for proving that a command C *refines* another command A , providing a formal link between the behaviours of C and A . Here, A is an abstraction and therefore admits more behaviours than C , or conversely, any behaviour of C must also be a behaviour of A . In an interval-based setting, we use the following definition of refinement [11]. In the context of our example, if refinement holds, then whenever a spinlock client is able to enter its critical section, it must also be possible for the abstract client to enter the critical section.

Definition 1. *If C and A are commands, then C refines A with respect to $P \subseteq Proc$, (denoted $C \sqsupseteq_P A$) iff for any interval Δ and stream s , $C.P \Rightarrow A.P$. We say C is equivalent to A with respect to P (denoted $C \sqsubseteq_P A$) iff both $C \sqsupseteq_P A$ and $A \sqsupseteq_P C$.*

Refinement is defined in terms of implication, and hence, relation \sqsupseteq_P is both reflexive and transitive. In this paper, we use \sqsupseteq_P as a basis for transforming the abstract specification $Spec$ and the concrete program $Prog$ individually. We use a notion of local buffer refinement (Definition 2) to relate concrete behaviours (with buffers) to abstract behaviours (without buffers).

Example 7. We transform the TSO implementation $Prog$ into a form that is more amenable to verification. In particular, a difficulty encountered when verifying $Prog$ in Example 6 directly is that for each process, p , command $Exec.p$ is not guaranteed to end in a flush, and hence changes to x may not be globally visible until the start of the next iteration (which starts with a lock that performs a flush all). In particular, it is not immediately possible to match the behaviour of Rel with abstract $ARel$ because Rel only places a pending write in the buffer, whereas $ARel$ modifies the value of x in memory. Therefore, we aim to transform $Prog$ to $Prog'$ below (see Proposition 1), where the flush occurs at the end of execution. We use notation

$$\overline{\overline{\overline{v \bullet C}}}_{\Phi} \hat{=} post_{\Phi}. \overline{\overline{v \bullet C}}$$

Unlike $\boxed{\bar{v} \bullet C}_{\Phi}$, command $\boxed{\bar{v} \bullet \bar{C}}_{\Phi}$ only flushes the buffer at the end of execution. Below, we have used the property $g^{\omega} \equiv \varepsilon \vee g^{\omega+}$ to split Acq into two cases. Note that Acq'_1 is defined in terms of Acq .

$$\begin{aligned} Acq'_1.p &\hat{=} \boxed{x \bullet \boxed{\bar{v} \bullet \bar{C}}_{\Phi}} ; \boxed{\bar{v} \bullet \bar{C}}_{\Phi} ; \llbracket \bar{v} \bullet \bar{C} \rrbracket^* ; \llbracket x \rrbracket ; Acq.p & Acq'_2.p &\hat{=} \boxed{x \bullet \llbracket x \rrbracket} ; \boxed{\bar{v} \bullet \bar{C}}_{\Phi} ; (x := 0)_{\Phi} \\ TryOK'.p &\hat{=} \boxed{x \bullet \llbracket x \rrbracket} ; (x := 0)_{\Phi} ; (r_p := true) & TryFl'.p &\hat{=} \boxed{x \bullet \llbracket \bar{v} \bullet \bar{C} \rrbracket} ; (r_p := false) \\ Exec'.p &\hat{=} ((Acq'_1 ; id ; Rel) \vee (Acq'_2 ; id ; Rel) \vee (TryOK' ; id ; Rel) \vee TryFl').p \\ Prog'.P &\hat{=} \{x\} \wedge_{p:P} \left(id ; \left(\{B_p = \langle \rangle\} Exec' ; post_{\Phi}.x.(\tilde{\Phi}^+) \right)^{\omega} \right).p \quad \square \end{aligned}$$

Clearly, transforming $Prog$ to $Prog'$ by reasoning at trace-based level of Definition 1 is infeasible. Therefore, we develop a number of refinement laws that are applied to our example. First, we have the following; the proof of each equivalence is straightforward.

Law 2 *If $p \in Proc$, C and D are commands, each C_i is a command and \bar{v} is a vector of variables, then*

$$\boxed{\bar{v} \bullet C}_{\Phi} \sqsubseteq_p pre_{\Phi}.\bar{v}.\boxed{\bar{v} \bullet \bar{C}}_{\Phi} \quad (4)$$

$$pre_{\Phi}.C \sqsubseteq_p pre_{\Phi}.\{\{B_p = \langle \rangle\}\}C \quad (5)$$

$$pre_{\Phi}.\bar{v}.\boxed{\bar{v} \bullet \bar{C}}_{\Phi} ; D \sqsubseteq_p pre_{\Phi}.\bar{v}.\left(\boxed{\bar{v} \bullet \bar{C}}_{\Phi} ; D\right) \quad \text{provided } C.p \Rightarrow \neg\varepsilon \quad (6)$$

$$\bigvee_i pre_{\Phi}.C_i \sqsubseteq_p pre_{\Phi}.\left(\bigvee_i C_i\right) \quad (7)$$

$$C ; pre_{\Phi}.D \sqsubseteq_p post_{\Phi}.C ; D \quad \text{provided } C.p \vee D.p \Rightarrow \neg\varepsilon \quad (8)$$

To transform $Prog$ to $Prog'$, we develop a leapfrog theorem analogous to Law 1, whose proof uses the equivalences defined in Law 2 as well as $PF.\bar{v} \hat{=} post_{\Phi}.\bar{v}.\tilde{\Phi}^+$.

Theorem 1 (Leapfrog flush). *Suppose $p \in Proc$, each C_i and D_i is a command such that $C_i.p \Rightarrow \neg\varepsilon$, and \bar{v} is a vector of variables. Then*

$$\left(\tilde{\Phi}^+ ; \left(\bigvee_i \boxed{\bar{v} \bullet C_i}_{\Phi} ; D_i\right)\right)^{\omega} ; PF.\bar{v} \sqsubseteq_p PF.\bar{v} ; \left(\{B_p = \langle \rangle\} \left(\bigvee_i \boxed{\bar{v} \bullet \bar{C}_i}_{\Phi} ; D_i\right) ; PF.\bar{v}\right)^{\omega}$$

The left hand side of Theorem 1 contains a disjunction that executes $\boxed{\bar{v} \bullet C_i}_{\Phi}$, which ensures the buffer is empty (via flushes if necessary) both before and after execution of C_i . After the end of the iteration, command $PF.\bar{v}$ is executed, which ensures the buffer is empty when the process terminates; flushes may be necessary due to the behaviour of D_i . On the right hand side, each iteration is guaranteed to start with an empty buffer and each disjunct starts with the weaker $\boxed{\bar{v} \bullet \bar{C}_i}_{\Phi}$, which only flushes the buffer at the end of execution. However, each iteration ends with $PF.\bar{v}$. Further note that on the right hand side, each iteration is guaranteed to begin in a state where the buffer of p is empty.

Proposition 1. $Prog \sqsubseteq_p Prog'$

Proof. Applying Theorem 1 to $Prog$, we obtain:

$$Prog''.P \hat{=} \{x\} \wedge_{p:P} \left(\{B_p = \langle \rangle\} PF.x ; (\{B_p = \langle \rangle\} Exec' ; PF.x)^{\omega} \right).p$$

Then, because $\{B_p = \langle \rangle\} PF.x \sqsubseteq_p id$, we have $Prog'' \sqsubseteq_p Prog'$. \square

For the proof in Section 5.3, we find the following laws to be useful, each of which is proved in a straightforward manner. Note that for (11) and (14), the refinement only holds in one direction. Of these, (14) states that an assignment $\{B_p = \langle \rangle\}v := k$ either ends with $B_p = \langle (v, k) \rangle$, or the buffer B_p is flushed as part of the assignment semantics.

Law 3 Suppose C and D are commands, p is a process, v is a variable and k a value. Then each of the following holds:

$$\{B_p = \langle \rangle\} \overline{v \bullet C} \sqsubseteq_p \{B_p = \langle \rangle\} \overline{v \bullet C} \quad (9)$$

$$\overline{v \bullet C}; \overline{v \bullet D} \sqsubseteq_p \overline{v \bullet C}; \overline{v \bullet D} \quad \text{provided } C.p \vee D.p \Rightarrow \neg \varepsilon \quad (10)$$

$$\{B_p = \langle \rangle\} \overline{v \bullet v := k} \sqsubseteq_p \overline{v \bullet \text{id}; v \leftarrow k; \text{id}} \{B_p = \langle \rangle\} \quad (11)$$

$$\overline{v \bullet C}; \overline{v \bullet D} \sqsubseteq_p \overline{v \bullet C; D} \quad (12)$$

$$\{B_p = \langle \rangle\} \overline{v \bullet [b]} \sqsubseteq_p \{B_p = \langle \rangle\} \overline{v \bullet [b]} \quad (13)$$

$$\{B_p = \langle \rangle\} v := k \sqsubseteq_p (\text{id}; v \leftarrow k; \text{id}) \{B_p = \langle (v, k) \rangle\}; (\varepsilon \vee (\Phi; \text{id})) \quad (14)$$

5.2 Local Buffer Refinement

In this section, we develop a novel method for proving refinement for TSO architectures, where buffer effects are taken into account. The method allows one to prove that a (concrete) process with buffer effects has the same behaviour as an (abstract) process without the buffer. In essence, one may pretend that the effect of local changes have already been flushed at the abstract level. This is allowed because in the context of the overall behaviour of a program, it makes no difference whether a variable local to process p has a pending write in p 's local buffer, or in shared memory. This essentially removes the potential non-determinism that arises from reasoning about flushes for local variables. We let

$$\begin{aligned} L\text{Cover}.P.(\sigma, \pi) &\hat{=} (\sigma \oplus \bigcup_{p,p} L\text{Var}.p \triangleleft \text{cover}.(\sigma.B_p), \pi) \\ L\text{Buffer}.P.s &\hat{=} \lambda t: \mathbb{Z} \bullet L\text{Cover}.P.(s.t) \end{aligned}$$

Here $\text{cover}.(\sigma.B_p)$ generates a set of pairs from $\sigma.B_p$ and $L\text{Var}.p \triangleleft \text{cover}.(\sigma.B_p)$ restricts these to the local variables of p . Within $L\text{Cover}.P.(\sigma, \pi)$ such a localised mapping is generated for each $p \in P$, and then, σ is overwritten by this mapping.

Definition 2. If C and A are commands, and P is a non-empty set of processes, we say C buffer refines A with respect to P , (denoted $C \ni_P A$) iff for any interval Δ and stream s , $(C.P).\Delta.s \Rightarrow (A.P).\Delta.(L\text{Buffer}.P.s)$.

Thus, whenever C holds for a set of processes P over interval Δ in stream s , command A must hold for P in Δ for the masked stream $L\text{Buffer}.P.s$. In particular, this implies that A behaves as if the local buffer effects of the concrete command have already been applied. We write \ni_p for $\ni_{\{p\}}$.

Clearly, reasoning at the level of Definition 2 is infeasible. Instead, we explore some higher level properties for the programming constructs we use. In general, buffer refinement is neither reflexive nor transitive⁸, e.g., if C always reads from main memory regardless of the state of the buffer, then reflexivity does not hold. However, \ni_P may be combined with standard refinement as follows, which holds trivially after expanding the definitions of \sqsubseteq_P and \ni_P .

Theorem 2. If $C' \sqsubseteq_P C$, $C \ni_P A$ and $A \sqsubseteq_P A'$ then $C' \ni_P A'$.

⁸ Fully exploring the properties of \ni_P lie outside the scope of this paper.

We immediately have the following laws, which help simplify verification of assignments and buffer flushes.

Law 4 *If $v \in \text{Var}$, $p \in \text{Proc}$, $k \in \text{Val}$, $P \hat{=} B_p \neq \langle \rangle \wedge (B_p.0 \upharpoonright 1) \in \text{LVar}.p$, and $Q \hat{=} B_p \neq \langle \rangle \wedge (B_p.0 \upharpoonright 1) \notin \text{LVar}.p$, then*

$$\{v \in \text{LVar}.p\}v \leftarrow k \ni_p v \Leftarrow k \quad (15) \quad \{P\} \Phi \ni_p \text{id} \quad (17)$$

$$\{v \notin \text{LVar}.p\}v \leftarrow k \ni_p \text{id} \quad (16) \quad \{Q\} \Phi \ni_p (B_p.0 \upharpoonright 1) \Leftarrow (B_p.0 \upharpoonright 2) \quad (18)$$

By conditions (15) and (16), adding a pending write (v, k) to p 's buffer is a local buffer refinement of a global update to v whenever $v \in \text{LVar}.p$, and of id , otherwise. On the other hand, (17) states that flushing a local variable of process p from p 's buffer has the same effect as executing id abstractly (because the effect of the variable has already occurred), and condition (18) states that flushing a global variable has the same effect as executing the corresponding write in memory at the abstract level.

The laws below allow local buffer refinement to be decomposed, and the pre/post assertions under local buffer refinement to be strengthened.

Law 5 *If $C \ni_p A$, $C_1 \ni_p A_1$ and $C_2 \ni_p A_2$ for a set of processes P , and b, c are state predicates such that $b \Rightarrow c$ and $(\text{vars}.b \cup \text{vars}.c) \cap (\bigcup_{p:P} \text{LVar}.p) = \emptyset$, then*

$$C_1 ; C_2 \ni_p A_1 ; A_2 \quad (19) \quad C_1 \vee C_2 \ni_p A_1 \vee A_2 \quad (22)$$

$$C^\omega \ni_p A^\omega \quad (20) \quad \{b\}C \ni_p \{c\}A \quad (23)$$

$$C_1 \wedge C_2 \ni_p A_1 \wedge A_2 \quad (21) \quad C\{b\} \ni_p A\{c\} \quad (24)$$

Law 6 (Parallel composition) *If $C'.P \hat{=} \bigwedge_{q:P} C.q$, $A'.P \hat{=} \bigwedge_{q:P} A.q$ and $C \ni_q A$ holds for each for each $q \in P$, then $C' \ni_p A'$.*

5.3 Application: Spinlock Example

We now apply our rules to the running spinlock example (modelled by *Prog*), and prove it to be a refinement of the abstract program (modelled by *Spec*). Our notion of refinement is local buffer refinement (Definition 2), i.e., we show

$$\text{Prog} \ni_p \text{Spec} \quad (25)$$

for an arbitrarily chosen non-empty set of processes P . Various refinement rules may be introduced to generalise the theory as needed. By Theorem 2 and Proposition 1, (25) immediately reduces to a proof of $\text{Prog}' \ni_p \text{Spec}$. Using (23), followed by Law 6, proof of $\text{Prog}' \ni_p \text{Spec}$ again reduces to the following for some arbitrarily chosen $p \in P$.

$$\text{id} ; (\{B_p = \langle \rangle\} \text{Exec}' ; \text{post}_{\Phi}.x.(\tilde{\Phi}^+))^\omega \ni_p (\text{id} ; \text{AExec})^\omega ; \text{id} \quad (26)$$

We use $\text{id} \sqsubseteq_p \text{id} ; \text{id}$ to split the first id on the right hand side of (26), then Law 1, to obtain $\text{id} ; (\text{id} ; \text{AExec}.p ; \text{id})^\omega$. Hence, using (19) followed by (20), the proof of (26) reduces again to a proof of

$$\{B_p = \langle \rangle\} \text{Exec}' ; \text{post}_{\Phi}.x.(\tilde{\Phi}^+) \ni_p \text{id} ; \text{AExec} ; \text{id} \quad (27)$$

Then using (22), condition (2) and the fact that $;$ distributes over \vee , we are left with a number of proof obligations for each disjunct of Exec' . Of these, we present the most complex: the proof obligation for Acq'_1 .

$$\{B_p = \langle \rangle\} Acq'_1 ; \text{id} ; Rel ; \text{post}_{\Phi}.x.(\tilde{\Phi}^+) \ni_p \text{id} ; AAcq ; \text{id} ; ARel ; \text{id} \quad (28)$$

It is trivial to show $\{B_p = \langle \rangle\} \boxed{x \bullet \llbracket \neg x \rrbracket}_{\Phi} ; \llbracket \neg x \rrbracket^* ; \llbracket x \rrbracket ; Acq \sqsupseteq_p \text{id} ; \{B_p = \langle \rangle\} Acq$, and using $\text{id} \sqsubseteq_p \text{id} ; \text{id}$ the proof of (28) reduces as follows, where the initial part of Acq'_1 is refined to id .

$$\{B_p = \langle \rangle\} Acq ; \text{id} ; Rel ; \text{post}_{\Phi}.x.(\tilde{\Phi}^+) \ni_p \text{id} ; AAcq ; \text{id} ; ARel ; \text{id} \quad (29)$$

We now focus on the initial part of the left hand side, where we distinguish between $ALoop \hat{=} \boxed{x \bullet \llbracket \neg x \rrbracket}_{\Phi} ; \llbracket \neg x \rrbracket^{\omega} ; \llbracket x \rrbracket$ and $ADo \hat{=} \boxed{x \bullet \llbracket x \rrbracket ; (x := 0)}_{\Phi}$.

$$\begin{aligned} & \{B_p = \langle \rangle\} Acq \\ \sqsubseteq_p & \{B_p = \langle \rangle\} ADo \vee \{B_p = \langle \rangle\} ALoop^{\omega+} ; ADo \quad \text{defn of } Acq, \text{ then } g^{\omega} \equiv \varepsilon \vee g^{\omega+} \end{aligned}$$

Using $\{B_p = \langle \rangle\} ALoop^{\omega+} \ni_p \text{id} \{B_p = \langle \rangle\}$ (proof elided) condition (29) reduces to

$$\{B_p = \langle \rangle\} ADo ; \text{id} ; Rel ; \text{post}_{\Phi}.x.(\tilde{\Phi}^+) \ni_p AAcq ; \text{id} ; ARel ; \text{id} \quad (30)$$

Again focusing on the initial part of the left hand side, we have:

$$\begin{aligned} & \{B_p = \langle \rangle\} ADo \\ \sqsubseteq_p & \{B_p = \langle \rangle\} \boxed{x \bullet \llbracket x \rrbracket ; (x := 0)}_{\Phi} && \text{defn of } ADo \text{ and by (9)} \\ \sqsubseteq_p & \{B_p = \langle \rangle\} \boxed{x \bullet \llbracket x \rrbracket} ; \boxed{x \bullet (x := 0)}_{\Phi} && \text{by (10)} \\ \sqsupseteq_p & \{B_p = \langle \rangle\} \boxed{x \bullet \llbracket x \rrbracket} ; \{B_p = \langle \rangle\} \boxed{x \bullet (x := 0)}_{\Phi} && \text{by logic} \\ \sqsupseteq_p & \{B_p = \langle \rangle\} \boxed{x \bullet \llbracket x \rrbracket} ; \boxed{x \bullet \text{id} ; (x \Leftarrow 0)} ; \text{id} \{B_p = \langle \rangle\} && \text{by (11)} \\ \sqsubseteq_p & \{B_p = \langle \rangle\} \boxed{x \bullet \llbracket x \rrbracket ; \text{id}} ; \boxed{x \bullet (x \Leftarrow 0)} ; \text{id} \{B_p = \langle \rangle\} && \text{using (12) twice} \\ \sqsupseteq_p & \{B_p = \langle \rangle\} \boxed{x \bullet \llbracket x \rrbracket} ; \boxed{x \bullet (x \Leftarrow 0)} ; \text{id} \{B_p = \langle \rangle\} && \llbracket x \rrbracket ; \text{id} \sqsubseteq_p \llbracket x \rrbracket \\ \sqsupseteq_p & \{B_p = \langle \rangle\} \boxed{x \bullet \llbracket x \rrbracket} ; \boxed{x \bullet (x \Leftarrow 0)} ; \text{id} \{B_p = \langle \rangle\} && \text{by (13)} \\ \sqsubseteq_p & \boxed{x \bullet \llbracket x \rrbracket ; (x \Leftarrow 0)} ; \text{id} \{B_p = \langle \rangle\} && \text{by (12) and weakening} \end{aligned}$$

It is straightforward to show $\boxed{x \bullet \llbracket x \rrbracket (x \Leftarrow 0)} ; \text{id} \ni_p AAcq ; \text{id}$, which expanding Rel and using (14) further reduces (30) to a proof of:

$$(\text{id} ; x \Leftarrow 1 ; \text{id}) \{B_p = \langle (x, 1) \rangle\} ; (\varepsilon \vee \Phi ; \text{id}) ; \text{post}_{\Phi}.x.(\tilde{\Phi}^+) \ni_p \text{id} ; ARel ; \text{id}$$

Using $\text{id} \sqsubseteq_p \text{id} ; \text{id}$ the right hand side transforms to $\text{id} ; \text{id} ; ARel ; \text{id}$. Then, using (16) and Theorem 2, we have $\text{id} ; x \Leftarrow 1 ; \text{id} \ni_p \text{id}$, and hence, using (19), we obtain $\{B_p = \langle (x, 1) \rangle\} (\text{post}_{\Phi}.x.(\tilde{\Phi}^+) \vee (\Phi ; \text{id} ; \text{post}_{\Phi}.x.(\tilde{\Phi}^+)) \ni_p \text{id} ; ARel ; \text{id}$. Finally, because $\Phi ; \text{id} ; \text{post}_{\Phi}.x.(\tilde{\Phi}^+) \sqsupseteq_p \text{post}_{\Phi}.x.(\tilde{\Phi}^+)$, we are left with

$$\{B_p = \langle (x, 1) \rangle\} \text{post}_{\Phi}.x.(\tilde{\Phi}^+) \ni_p \text{id} ; ARel ; \text{id} \quad (31)$$

Because there is exactly one item in the buffer, and because $\text{post}_{\Phi}.x.(\tilde{\Phi}^+)$ guarantees that a flush occurs, the left hand side reduces to $\{B_p = \langle (x, 1) \rangle\} \text{id} ; \Phi ; \text{id}$. Finally, using the fact that B_p is a local variable of p and is not modified by id followed by (18), our proof is completed.

Notable in our verification is that concurrency aspects hardly need to be considered. The fact that the locking mechanisms guarantee safety is understood at the level of *Spec*. The refinement proof only requires consideration of the local buffer. This is in

contrast to existing methods which require global conditions to be checked, e.g., [20] checks race conditions, [7, 16, 25] check linearizability, and [9] checks reduction. We conjecture that more complex examples will indeed require consideration of the behaviour of other processes. To this end, we will integrate compositional methods such as rely/guarantee into our framework [13, 11].

6 Conclusions

Existing approaches to relaxed memory verification (e.g., [6, 21, 22, 9, 20, 7, 16]) focus on a low-level language (i.e., individual reads/writes), and hence, to perform a verification, programs need to be observed and understood in their (verbose) low-level representation. We are not aware of any approach that tries to lift memory model effects to a higher level of abstraction; our work here is hence unique in this sense [15].

The basic idea is to think of a statements as being executed over an interval of time or an execution window. Such execution windows can overlap if programs are executed concurrently and overlapping windows correspond to program instructions that can be executed in any order, representing the effect of concurrent executions and reorderings due to TSO. Overlapping execution windows may also interfere with each other and fixing the outcome of an execution within a window can influence the outcome within another. This paper presents several advances to the semantics in [15] by simplifying the interval logic, and program semantics, as well as developing buffer-specific rules for expression evaluation and refinement. The underlying rules are algebraic in nature, and hence, we provide generic transformation laws, which are in turn applied to our running example.

A difficulty when reasoning about TSO memory is that in addition to the normal non-determinism caused by concurrency, an additional level of non-determinism is introduced via use of local buffers. The methods in this paper allow one to reduce the non-determinism that must be considered when reasoning about local updates. In particular, we develop a notion of *local buffer refinement*, which allows one to proceed as if pending writes to local variables have already occurred in the abstract level. In particular, this means that local writes do not appear out of order. A similar observation is used for local transformation in the context of compilers for weak memory [8], however, these do not consider higher-level synchronisation instructions such as `lock`.

As part of future work, we aim to study the connections between local buffer refinement, and existing notions such as triangular race freedom [20] and reduction [9].

References

1. S. V. Adve and K. Gharachorloo. Shared memory consistency models: A tutorial. *IEEE Computer*, 29(12):66–76, 1996.
2. J. Alglave. A formal hierarchy of weak memory models. *Formal Methods in System Design*, 41(2):178–210, 2012.
3. R. J. R. Back and J. von Wright. Reasoning algebraically about loops. *Acta Informatica*, 36(4):295–334, July 1999.
4. D. Bovet and M. Cesati. *Understanding the Linux Kernel*. OReilly, 3rd edition, 2005.

5. J. Boyland. Checking interference with fractional permissions. In R. Cousot, editor, *SAS*, volume 2694 of *LNCS*, pages 55–72. Springer, 2003.
6. S. Burckhardt, R. Alur, and M. M. K. Martin. Checkfence: Checking consistency of concurrent data types on relaxed memory models. In *PLDI*, pages 12–21, 2007.
7. S. Burckhardt, A. Gotsman, M. Musuvathi, and H. Yang. Concurrent library correctness on the TSO memory model. In H. Seidl, editor, *ESOP*, volume 7211 of *LNCS*, pages 87–107. Springer, 2012.
8. S. Burckhardt, M. Musuvathi, and V. Singh. Verifying local transformations on relaxed memory models. In R. Gupta, editor, *CC*, volume 6011 of *LNCS*, pages 104–123. Springer, 2010.
9. E. Cohen and B. Schirmer. From total store order to sequential consistency: A practical reduction theorem. In M. Kaufmann and L. C. Paulson, editors, *ITP*, volume 6172 of *LNCS*, pages 403–418. Springer, 2010.
10. B. Dongol and J. Derrick. Data refinement for true concurrency. In J. Derrick, E. A. Boiten, and S. Reeves, editors, *Refine*, volume 115 of *EPTCS*, pages 15–35, 2013.
11. B. Dongol, J. Derrick, and I. J. Hayes. Fractional permissions and non-deterministic evaluators in interval temporal logic. *ECEASST*, 53, 2012.
12. B. Dongol and I. J. Hayes. Deriving real-time action systems in a sampling logic. *Sci. Comput. Program.*, 78(11):2047–2063, 2013.
13. B. Dongol, I. J. Hayes, and J. Derrick. Deriving real-time action systems with multiple time bands using algebraic reasoning. *Sci. of Comp. Prog.*, 85, Part B(0):137 – 165, 2014.
14. B. Dongol, I. J. Hayes, L. Meinicke, and K. Solin. Towards an algebra for real-time programs. In W. Kahl and T.G. Griffin, editors, *RAMiCS*, volume 7560 of *LNCS*, pages 50–65, 2012.
15. B. Dongol, O. Travkin, J. Derrick, and H. Wehrheim. A high-level semantics for program execution under total store order memory. In Z. Liu, J. Woodcock, and H. Zhu, editors, *ICTAC*, volume 8049 of *LNCS*, pages 177–194. Springer, 2013.
16. A. Gotsman, M. Musuvathi, and H. Yang. Show no weakness: Sequentially consistent specifications of TSO libraries. In M. K. Aguilera, editor, *DISC*, volume 7611 of *LNCS*, pages 31–45. Springer, 2012.
17. I. J. Hayes, A. Burns, B. Dongol, and C. B. Jones. Comparing degrees of non-determinism in expression evaluation. *Comput. J.*, 56(6):741–755, 2013.
18. L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691, 1979.
19. B. C. Moszkowski. A complete axiomatization of Interval Temporal Logic with infinite time. In *LICS*, pages 241–252, 2000.
20. S. Owens. Reasoning about the implementation of concurrency abstractions on x86-TSO. In T. D’Hondt, editor, *ECOOP*, volume 6183 of *LNCS*, pages 478–503. Springer, 2010.
21. S. Park and D. L. Dill. An executable specification, analyzer and verifier for RMO (relaxed memory order). In *SPAA*, pages 34–41, 1995.
22. P. Sewell, S. Sarkar, S. Owens, F. Z. Nardelli, and M. O. Myreen. x86-TSO: a rigorous and usable programmer’s model for x86 multiprocessors. *Commun. ACM*, 53(7):89–97, 2010.
23. D. J. Sorin, M. D. Hill, and D. A. Wood. *A Primer on Memory Consistency and Cache Coherence*. Synthesis Lectures on Computer Architecture. Morgan & Claypool, 2011.
24. J.M. Spivey. *The Z Notation: A Reference Manual*. Prentice Hall, 1992.
25. O. Travkin, A. Mütze, and H. Wehrheim. SPIN as a linearizability checker under weak memory models. In V. Bertacco and A. Legay, editors, *HVC’13*, volume 8244 of *LNCS*, pages 311–326. Springer, 2013.