

A DENOTATIONAL SEMANTICS FOR SPARC TSO

RYAN KAVANAGH AND STEPHEN BROOKES

Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA
e-mail address: {rkavanagh, brookes}@cs.cmu.edu

ABSTRACT. The SPARC TSO weak memory model is defined axiomatically, with a non-compositional formulation that makes modular reasoning about programs difficult. Our denotational approach uses pomsets to provide a compositional semantics capturing exactly the behaviours permitted by SPARC TSO. It uses buffered states and an inductive definition of execution to assign an input-output meaning to pomsets. We show that our denotational account is sound and complete relative to the axiomatic account, that is, that it captures exactly the behaviours permitted by the axiomatic account. Our compositional approach facilitates the study of SPARC TSO and supports modular analysis of program behaviour.

1. INTRODUCTION

A memory model specifies what values can be read by a thread from a given memory location during execution. Traditional concurrency research has assumed *sequential consistency*, wherein memory actions operate atomically on a global state, and a read is guaranteed to observe the value most recently written to that location *globally*. Consequently, “the result of any execution is the same as if the operations of all the processors were executed in some sequential order” [Lam79]. However, sequential consistency negatively impacts performance, and modern architectures often provide much weaker guarantees. These weaker guarantees mean that classical concurrency algorithms, which often assume sequential consistency, can behave in unexpected ways. Consider, for example, the Dekker algorithm on a system using the SPARC instruction set. The Dekker algorithm seeks to ensure that at most one process enters a critical section at a time. Executing the following instance of the Dekker algorithm on a sequentially consistent system from an initial state where memory locations w , x , y , z are all zero will ensure that we end in a state where not both z and w are set to one:

$$(x := 1; \text{if } y = 0 \text{ then } z := 1 \text{ else skip}) \parallel (y := 1; \text{if } x = 0 \text{ then } w := 1 \text{ else skip}).$$

2012 ACM CCS: [Theory of computation]: Denotational semantics; Concurrency.

Key words and phrases: SPARC TSO, denotational semantics, pomsets, concurrency, weak memory models.

Funded in part by a Natural Sciences and Engineering Research Council of Canada (NSERC) Postgraduate Scholarship.

However, the SPARC ISA provides the weaker SPARC TSO (total store ordering) memory model. Under SPARC TSO, it is possible to start from the aforementioned initial state and end in a state where both z and w are set to one, thus violating mutual exclusion.

Weak memory models are often described in standards documents using natural language. This informality makes it difficult to reason about how programs will behave on systems that use these memory models. The SPARC Architecture Manual [SPA92] gives an axiomatic description of TSO using partial orders of actions. We present this description in Section 2; other axiomatic approaches are discussed in Section 5. Despite their formality, it is hard to use axiomatic accounts to reason about the behaviour of programs. This is because the axiomatic approach is non-compositional and precludes modular reasoning. We address this problem by presenting a denotational semantics for SPARC TSO in Section 3. Our denotational semantics assigns to each program a collection of pomsets. Pomsets are generalizations of traces and were first used by Pratt [Pra86] to give denotational semantics for concurrency, and later by Brookes [Bro16a], with some modifications, to study weak memory models. We illustrate our semantics by validating various *litmus tests* and expected program equivalences. To ensure our denotational semantics accurately captures the behaviour of SPARC TSO, we show in Section 4 that from every denotation of a program we can derive a collection of partial orders satisfying the axiomatic description of Section 2 and, moreover, that we can derive every such partial order from the denotation.

2. AN AXIOMATIC ACCOUNT

The SPARC manual [SPA92] gives an axiomatic description of TSO in terms of partial orders of actions. Unfortunately, this description is incomplete because it fails to specify the fork and join behaviour of TSO. In this section, we complete the SPARC manual’s account of TSO with an account of forking and joining. Before doing so, we give an informal description of TSO to help build intuition.

A system providing the TSO weak memory model can be thought of as a collection of processors, each with a write buffer. Whenever a processor performs a write, it places it in its write buffer. The buffer behaves as a queue, and writes migrate out of the buffers one at a time, and shared memory applies them according to a global total order. Whenever a processor tries to read a location, it first checks its buffer for a write to that location. If it finds one, it uses the value of the most recent such write; otherwise, it looks to shared memory for a value. Because of buffering, it is possible for writes and reads to be observed out of order relative to the program order.

2.1. Program Order Pomsets. To make the above intuition precise, we must formalize the notion of a *program order*, i.e., the ordering of read and write actions specified by a program. We do so by means of partially-ordered multisets.

Definition 2.1. A (strict) **partially-ordered multiset** or **pomset** $(P, <, \Phi)$ over a label set L consists of a strict poset $(P, <)$ of “action occurrences” and a function $\Phi : P \rightarrow L$ mapping each action occurrence to its label or “action”.

We frequently write just P for $(P, <, \Phi)$, in which case we let $<_P$ and Φ_P denote its obvious components. Denote by $\text{Pom}(L)$ the set of pomsets over L . We remark that pomsets are a natural generalization of traces. Indeed, all traces are simply pomsets where the underlying order is total.

We do not usually make the poset P explicit, because the structure of the pomset is invariant under relabellings of the elements of P . Consequently, we identify pomsets $(P, <, \Phi)$ and $(P', <', \Phi')$ if there exists an order isomorphism $\phi : P \rightarrow P'$ such that $\Phi = \Phi' \circ \phi$. We usually denote the elements of the pomset using just their labels, but we sometimes need to specify their exact occurrence, in which case we write l_p , where $l = \Phi(p)$.

It is useful to draw a pomset P as a labelled directed acyclic graph, where multiple vertices can have the same label and we have an edge $a \rightarrow b$ if $a <_P b$. For clarity, we always omit edges obtained by transitivity of $<_P$. For example, the following graph depicts the pomset where $P = \{0, 1, 2, 3\}$, the order is given by $0 < 1$, $1 < 2$, $0 < 2$, and $0 < 3$, and Φ is given by $\Phi(0) = a$, $\Phi(1) = b$, $\Phi(2) = a$, and $\Phi(3) = c$:

$$c \longleftarrow a \longrightarrow b \longrightarrow a .$$

We assume a countably infinite set of locations Loc , ranged over by metavariables x, y, z, \dots , and a set of values V , ranged over by v . In our examples, we will take V to be the set of integers. We call $x := v$ a *global write action*, $x = v$ a *read action*, and δ a *skip action*. Let \mathcal{A}_w and \mathcal{A}_r be the sets of global write actions and read actions, respectively.

Definition 2.2. A **program order** is a pomset P over the set $\mathcal{A}_{PO} = \mathcal{A}_w \cup \mathcal{A}_r \cup \{\delta\}$ of action labels that satisfies the (*locally*) *finite height property*, that is, such that for all $b \in P$, the set $\{a \in P \mid a <_P b\}$ is finite.

Informally, the finite height property guarantees that all actions in P can be executed after finitely many other actions, i.e., that the program order contains no unreachable actions.

Intuitively, the program order

$$x := 2 \longrightarrow y = 1 \quad y := 1 \longrightarrow y = 1$$

describes the parallel execution of writing 2 to x before reading 1 from y , and writing 1 to y before reading 1 from y , with no other ordering constraints.

2.2. TSO Axioms. We now turn our attention to our completed version of the axiomatic account given in the SPARC manual. To do so, we first introduce the notion of state and the requisite notation.

A global state is a finite partial function from locations Loc to values V . We let $\Sigma_{PO} = \text{Loc} \rightarrow_{\text{fin}} V$ be the set of global states, and use σ to range over Σ_{PO} .

Given any set S and partial order $<_S$ on it, every element $s \in S$ determines a set $s \downarrow_S = \{s' \in S \mid s' <_S s\} \cup \{s\}$ called its *lower closure*. Write $s \parallel_S s'$ to denote that s and s' are not comparable under the reflexive closure \leq_S of $<_S$, and write $s \leq_S s'$ to denote that they are comparable.

Definition 2.3. Let P be a program order and $<_T$ be a strict partial order on the elements of P . We say $<_T$ is **TSO-consistent** with P from (the initial state) σ if it satisfies the following six axioms:

- (O) **Ordering:** $<_T$ totally orders the write actions \mathcal{A}_w of P .
- (V) **Value:** for all reads $(x = v)_r$ in P , either
 - (A) there exists a write $(x := v')_w$ maximal under $<_T$ amongst all writes to x in $(x = v)_r \downarrow_T$, all writes to x in $(x = v)_r \downarrow_P$ are in $(x := v')_w \downarrow_T$, and $v = v'$; or
 - (B) there exists a write $(x := v')_w$ maximal under $<_P$ amongst all writes to x in $(x = v)_r \downarrow_P$, and both $(x = v)_r <_T (x := v')_w$ and $v = v'$; or

- (C) there are no writes to x in $(x = v)_{r \downarrow T}$ or $(x = v)_{r \downarrow P}$, and $\sigma(x) = v$.
(L) **LoadOp**: for all reads $r \in P$ and all actions $a \in P$, $r <_P a$ implies $r <_T a$.
(S) **StoreStore**: for all writes $w, w' \in P$, $w <_P w'$ implies $w <_T w'$.
(F) **Fork**: if $\alpha_1 <_P \alpha_2$, $\alpha_1 <_P \alpha_3$, and $\alpha_2 \parallel_P \alpha_3$, then $\alpha_1 <_T \alpha_2$ and $\alpha_1 <_T \alpha_3$.
(J) **Join**: if $\alpha_1 <_P \alpha_3$, $\alpha_2 <_P \alpha_3$, and $\alpha_1 \parallel_P \alpha_2$, then $\alpha_1 <_T \alpha_3$ and $\alpha_2 <_T \alpha_3$.

The fork axiom is easily understood by: if $\alpha_1 \leftarrow \alpha_1 \rightarrow \alpha_3$ in P , then $\alpha_2 \leftarrow \alpha_1 \rightarrow \alpha_3$ or $\alpha_1 \rightarrow \alpha_2 \rightarrow \alpha_3$ in $(P, <_T)$; the join axiom is symmetric. We simply say $<_T$ is TSO-consistent with P if there exists some initial state σ from which they are TSO-consistent. It will be useful to identify $<_T$ and the pomset $T = (P, <_T, \Phi_P)$. \square

Axioms (O), (VA), (VB), (L), and (S) are directly adapted from the formal specification given in Appendix K.2 of [SPA92]. We introduce axiom (VC) to simplify our presentation of examples. By requiring that programs first write to any locations from which they read, it can be omitted, and apart from examples, we will assume throughout that our TSO-consistent orders do not require (VC). Though the formal specification does not provide axioms (F) and (J), they are consistent with the behaviour intended by Appendix J.6 of [SPA92]. Intuitively, axiom (VB) requires that whenever a processor reads from a location, it must use the most recent write to that location in its write buffer (if it exists). If there is no such write in its write buffer, but we have observed a global write to that location, then (VA) requires that the most recent such write be the one read. Our presentation differs slightly from the formal specification. In particular, we do not consider instruction fetches or atomic load-store operations, and we do not consider flush actions, because they can be implemented as a derived action in our semantics by forking and immediately joining. To be consistent with our pomset development, we also assume the order to be strict.

As the following proposition's corollary shows, if $<_T$ is TSO-consistent for P , then there exists a (not necessarily unique) total order on P that is TSO-consistent with P and contains $<_T$. As a result, we can view all orders that are TSO-consistent with P as weakenings of total orders that are TSO-consistent with P . The proposition follows by a straightforward check of the axioms, where (V) is the only axiom that is not immediate.

Proposition 2.4. *Let $<_T$ be TSO-consistent for P and $a, b \in P$ be two actions such that $a \parallel_T b$. Let $<_{ab}$ be the transitive closure of $<_T \cup \{(a, b)\}$. If there exist maximum writes under a and b relative to \leq_T , call them μ_a and μ_b , respectively. If either μ_a and μ_b exist and $\mu_a <_T \mu_b$, or μ_a does not exist, then $<_{ab}$ is TSO-consistent for P . \square*

Corollary 2.5. *If P is a finite program order pomset and $<_T$ a partial order TSO-consistent with P , then there exists a total order \sqsubset_T TSO-consistent with P such that $<_T \subseteq \sqsubset_T$. \square*

Despite Corollary 2.5, one should be careful not to conflate the notion of linearisation with that of TSO-consistent total orders. Consider, for example, the program order

$$x := 2 \longrightarrow x = 2 \quad x := 3 \longrightarrow x = 3 .$$

The linearisation $x := 2 < x := 3 < x = 3 < x = 2$ is not TSO-consistent with the program order because it violates (VA); the order $x = 2 < x := 2 < x = 3 < x := 3$ is not a linearisation of the program order but is TSO-consistent with it.

When we have a write followed by a read in the program order, but swapped in the linear order, as in this example, we can imagine the write having gotten stuck in the write buffer, and observing the read before the write.

3. A DENOTATIONAL ACCOUNT

So far we have dealt with program orders in the abstract. To make the rest of our development more concrete, we restrict our attention to program orders for well-defined programs in the simple imperative language given below. These program orders are defined in Section 3.2.

Restricting our attention to program orders of these well-defined programs raises the question of *compositionality*. The key is to find a way to derive TSO-consistent orders for a sequential composition $c_1; c_2$ or parallel composition $c_1 \parallel c_2$ given TSO-consistent orders for c_1 and c_2 . This is infeasible with the axiomatic approach, which requires reasoning about whole programs and is inherently non-compositional. In contrast, a denotational approach using pomsets is compositional: it allows us to derive the meaning of a program vis-à-vis a weak memory model from the meanings of its parts vis-à-vis the memory model.

Our denotational semantics has two components. The first associates to each program a set of *TSO pomsets*, which serves as the *abstract meaning* or *denotation* of the program. This component is described in Section 3.3. The second associates to each pomset a set of *executions*, which describe its input-output behaviours. This is described in Section 3.4.2.

3.1. A Simple Imperative Language. We express our programs using a simple imperative language. This formalism avoids the complexity of high-level languages, while still capturing the programs we are interested in. In the syntax below, e ranges over integer expressions, b over boolean expressions, c over commands, and p over programs. We distinguish between commands and programs, because although commands can be composed to form new commands, programs are assumed to be syntactically complete and executable. This distinction will be useful later when we consider executions, where we will assume programs are executed from initial states with empty buffers, but impose no such constraint on executions for commands.

$$\begin{aligned}
v &::= \dots, -2, -1, 0, 1, 2, \dots \\
e &::= v \mid x \mid e_1 + e_2 \mid e_1 * e_2 \mid \dots \\
b &::= \text{true} \mid \text{false} \mid \neg b \mid e_1 = e_2 \mid e_1 < e_2 \mid b_1 \vee b_2 \mid b_1 \wedge b_2 \mid \dots \\
c &::= \mathbf{skip} \mid x := e \mid c_1; c_2 \mid c_1 \parallel c_2 \mid \mathbf{if } b \mathbf{ then } c_1 \mathbf{ else } c_2 \mid \mathbf{while } b \mathbf{ do } c \\
p &::= c
\end{aligned}$$

Let \mathbf{VExp} denote the set of integer expressions, \mathbf{BExp} the set of boolean expressions, and \mathbf{Cmd} the set of commands.

3.2. PO Pomsets. Given a command c in our language, we must now compile it down to its set $\mathcal{P}_{PO}(c)$ of program order pomsets. We need operations for the sequential and parallel composition of pomsets over the same set of labels. When defining compositions of pomsets, we assume without loss of generality that the underlying posets are disjoint.

Definition 3.1. The **sequential composition** $(P_0, <_0, \Phi_0); (P_1, <_1, \Phi_1)$ is $(P_0, <_0, \Phi_0)$ whenever P_0 is infinite, and otherwise it is $(P_0 \cup P_1, <_0 \cup <_1 \cup P_0 \times P_1, \Phi_0 \cup \Phi_1)$. The **parallel composition** $(P_0, <_0, \Phi_0) \parallel (P_1, <_1, \Phi_1)$ of pomsets is $(P_0 \cup P_1, <_0 \cup <_1, \Phi_0 \cup \Phi_1)$. The empty pomset $\mathbf{0} = (\emptyset, \emptyset, \emptyset)$ is the unit for sequential and parallel composition. Given a pomset P on a set of labels L and a subset $L' \subseteq L$, the **restriction** $P \upharpoonright_{L'}$ of P to L' is the pomset on $\Phi^{-1}(L')$ whose ordering is induced by P . The **deletion** of L' from P is $P \upharpoonright_{L \setminus L'}$. We lift these operations to sets of pomsets in the obvious manner, e.g., $S_1; S_2 = \{P_1; P_2 \mid P_i \in S_i\}$. \square

Because the skip action δ has no effects, we identify program orders P and P' whenever there exists a non-empty pomset P_δ that can be obtained in two ways: by deleting a finite number of δ actions from P and also by deleting a finite number of δ actions from P' . This means, e.g., that we identify $\{\delta\}; P$, $\{\delta\} \parallel P$, and P whenever $P \neq \mathbf{0}$, but $\{\delta\}; \mathbf{0} = \{\delta\} \neq \mathbf{0}$.

We begin with the program order denotation of expressions. To each expression e , we assign a set $\mathcal{P}_{PO}(e)$ of tuples of program orders and corresponding values:

$$\begin{aligned} \mathcal{P}_{PO} &: \mathbf{VExp} \rightarrow \wp(\text{Pom}(\mathcal{A}_{PO}) \times V) \\ \mathcal{P}_{PO}(v) &= \{(\{\delta\}, v)\} \\ \mathcal{P}_{PO}(x) &= \{(\{x = v\}, v) \mid v \in V\} \\ \mathcal{P}_{PO}(e_1 \odot e_2) &= \{(P_1 \parallel P_2, v_1 \odot v_2) \mid (P_i, v_i) \in \mathcal{P}_{PO}(e_i)\} \end{aligned}$$

where \odot ranges over binary operations. Read expressions x are associated with arbitrary values in V for reasons of compositionality: we do not know with which writes the read may eventually be composed, and so we need to permit reading arbitrary values. We chose to evaluate binary operations $e_1 \odot e_2$ in parallel; one could just as legitimately have chosen to sequentialise the evaluation and written $P_1; P_2$. We assume $v_1 \odot v_2 \in V$ to be the result of applying the binary operation \odot to v_1 and v_2 . We handle program orders for unary expressions analogously, and assume $\neg b'$ is the result of negating the boolean value b' . To simplify the clauses involving conditionals, we give helper functions $\mathcal{P}_{true}(b)$ and $\mathcal{P}_{false}(b)$ to extract the pomsets corresponding to the given boolean values from $\mathcal{P}_{PO}(b)$.

$$\begin{aligned} \mathcal{P}_{PO} &: \mathbf{BExp} \rightarrow \wp(\text{Pom}(\mathcal{A}_{PO}) \times \text{Bool}) \\ \mathcal{P}_{PO}(b) &= \{(\{\delta\}, b)\} \quad (b \in \{true, false\}) \\ \mathcal{P}_{PO}(\neg b) &= \{(P, \neg b') \mid (P, b') \in \mathcal{P}_{PO}(b)\} \\ \mathcal{P}_{PO}(e_1 \odot e_2) &= \{(P_1 \parallel P_2, v_1 \odot v_2) \mid (P_i, v_i) \in \mathcal{P}_{PO}(e_i)\} \\ \mathcal{P}_{true}(b) &= \{P \mid (P, true) \in \mathcal{P}_{PO}(b)\} \\ \mathcal{P}_{false}(b) &= \{P \mid (P, false) \in \mathcal{P}_{PO}(b)\} \end{aligned}$$

Note that in the case of boolean binary operations, the e_i might be integer or boolean expressions, and the corresponding semantic clause for $\mathcal{P}_{PO}(e_i)$ should be used.

We give the program order denotation of commands in a similar manner, this time associating sets of program orders to each command phrase:

$$\begin{aligned} \mathcal{P}_{PO} &: \mathbf{Cmd} \rightarrow \wp(\text{Pom}(\mathcal{A}_{PO})) \\ \mathcal{P}_{PO}(\mathbf{skip}) &= \{\{\delta\}\} \\ \mathcal{P}_{PO}(x := e) &= \{P; \{x := v\} \mid (P, v) \in \mathcal{P}_{PO}(e)\} \\ \mathcal{P}_{PO}(c_1; c_2) &= \mathcal{P}_{PO}(c_1); \mathcal{P}_{PO}(c_2) \\ \mathcal{P}_{PO}(c_1 \parallel c_2) &= \mathcal{P}_{PO}(c_1) \parallel \mathcal{P}_{PO}(c_2) \\ \mathcal{P}_{PO}(\mathbf{if } b \mathbf{ then } c_1 \mathbf{ else } c_2) &= \mathcal{P}_{true}(b); \mathcal{P}_{PO}(c_1) \cup \mathcal{P}_{false}(b); \mathcal{P}_{PO}(c_2) \\ \mathcal{P}_{PO}(\mathbf{while } b \mathbf{ do } c) &= \bigcup_{n=0}^{\infty} I^n(b, c) \cup I^\omega(b, c), \end{aligned}$$

where $I^0(b, c) = \mathcal{P}_{false}(b)$ and $I^{n+1}(b, c) = \mathcal{P}_{true}(b); \mathcal{P}_{PO}(c); I^n(b, c)$.

The only interesting clause is for **while** b **do** c . Here, we take union of all of the finite unrollings $I^n(b, c)$ of the loop. We must also consider the case of an infinite loop. This is captured by $I^\omega(b, c)$, which describes the infinite pomset obtained by unrolling the loop countably infinitely many times. The **while** b **do** c clause also illustrates why we associate the pomset $\{\delta\}$ instead of $\mathbf{0}$ to values: otherwise, we would have $\mathcal{P}_{PO}(\mathbf{while\ false\ do\ }c) = \{\mathbf{0}\}$, and this would break our intuition that this program should be denotationally equivalent to **skip**. It would also have no executions under the formal account of Section 3.4.

To illustrate the above semantic clauses, we return to the Dekker program from the introduction. This program has pomsets of each of the following forms, for each choice of $v \neq 0$ and $v' \neq 0$:

$$\begin{array}{cccc}
 x := 1 & y := 1, & x := 1 & y := 1, & x := 1 & y := 1, & x := 1 & y := 1. \\
 \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\
 y = 0 & x = 0 & y = v & x = 0 & y = 0 & x = v & y = v & x = v' \\
 \downarrow & \downarrow & & \downarrow & \downarrow & & & \\
 z := 1 & w := 1 & & w := 1 & z := 1 & & &
 \end{array}$$

The first program order describes an execution where we read both $y = 0$ and $x = 0$ and where Dekker fails. The next three forms of program order describe executions in which one or both reads obtain a non-zero value.

3.3. TSO Pomsets. In this subsection, we assign a set $\mathcal{P}_{TSO}(p)$ of *TSO pomsets* to each program p , serving as the abstract meaning or denotation of the program p under the TSO memory model. To do so, we will need to carefully model write buffers. For compactness, we will write \mathcal{P} instead of \mathcal{P}_{TSO} in this section's semantic clauses.

We introduce a set $\text{BLoc} = \{\bar{x} \mid x \in \text{Loc}\}$ of buffer locations, assumed to be in bijection with Loc , and let the set of buffer write actions be $\mathcal{A}_b = \{\bar{x} := v \mid \bar{x} \in \text{BLoc}, v \in V\}$. An action $\bar{x} := v$ by a thread denotes adding a write $x := v$ to the thread's write buffer. The set of TSO actions \mathcal{A}_{TSO} then consists of \mathcal{A}_{PO} extended with \mathcal{A}_b . A *TSO pomset* will then be a pomset in $\text{Pom}(\mathcal{A}_{TSO})$ satisfying the finite height property.

To capture the effects of buffers, we parametrize our semantic clauses with lists of global write actions, which represent the writes currently in our buffer. We let $\text{Ls} = \mathcal{A}_w$ list be the set of all lists. The intuition is that write buffers behave as queues under TSO, and we can use a list $L \in \text{Ls}$ to model a queue by dequeuing from the head of the list and enqueueing at the end of the list. For expository convenience, we identify lists and *linear* pomsets, where we say a pomset is linear if its underlying poset is linear. Explicitly, we identify $[\]$ with the empty pomset $\mathbf{0}$, and $[\lambda_1, \dots, \lambda_n]$ with the pomset $\{\lambda_1\}; \dots; \{\lambda_n\}$. To minimize notation, we leverage this identification and write $L; L'$ to denote the concatenation of L and L' .

The semantic clauses are given in two strata. The semantic clauses \mathcal{B} for “basic TSO pomsets” capture the meaning of the syntactic phrases in a manner very similar to the program order definitions in Section 3.2. \mathcal{B} assign to each command phrase a function from buffer lists to a set of pairs of TSO pomsets and buffer lists. We present these clauses using the abbreviation $\mathcal{B}_L(c) = \mathcal{B}(c)(L)$. The pomset component of $\mathcal{B}_L(c)$ captures the meaning of the phrase in the presence of the buffer L , while the buffer component describes the state of the buffer after performing the actions associated with the phrase. In the second stratum, we use \mathcal{P} clauses to capture the meaning of the phrase in the presence of buffer flushing. Flushing a write from a buffer L consists of dequeuing a global write $x := v$ from L and inserting it in the pomset. $\mathcal{P}_L(c) = \mathcal{P}(c)(L)$ is again a subset of $\text{Pom}(\mathcal{A}_{TSO}) \times \text{Ls}$.

To generate TSO pomsets, we modify the semantic clauses generating program orders in four key places to get our basic pomsets. The first is for write commands $x := e$. Starting from a buffer $L \in \text{Ls}$, we get the pomset P and associated value v for e from the denotation $\mathcal{P}_L(e)$ instead of $\mathcal{P}_{PO}(e)$. The buffer L may have changed to a buffer L' while we were evaluating e , and $\mathcal{P}_L(e)$ also gives us this L' . Instead of immediately making a global write to x as we would have in the program order clause, we enqueue the global write on the buffer L' :

$$\mathcal{B}_L(x := e) = \{(P; \{\bar{x} := v\}, L'; \{x := v\}) \mid (P, v, L') \in \mathcal{P}_L(e)\}.$$

We must also change the semantic clauses for read expressions. By axiom (VB), whenever we read from a location x , we must use the most recent value available for it in the write buffer, if available. We use the following helper function to convert a buffer $L \in \text{Ls}$ to a partial function $\beta_L : \text{Loc} \rightarrow_{fin} V$ giving us the value of the most recent write in L to a given location:

$$\beta_{[\]}(x) = \text{undefined}, \quad \beta_{L; \{x:=v\}}(y) = \begin{cases} v & \text{if } x = y \\ \beta_L(y) & \text{otherwise.} \end{cases} \quad (3.1)$$

Then, the semantic clause giving us the basic pomsets for reads is

$$\mathcal{B}_L(x) = \{(\{x = v\}, v, L) \mid \beta_L(x) = v\} \cup \{(\{x = v\}, v, L) \mid x \notin \text{dom}(\beta_L), v \in V\}.$$

The first part tells us to use the value associated with x in the buffer L , if available. The second part uses arbitrary values if the value is unavailable, as with program orders.

The third major change involves parallel composition. We explain parallel composition of expressions; parallel composition of commands is analogous. By axioms (F) and (J), we must flush our buffers before every fork and join. We therefore begin by flushing our entire buffer, i.e., by taking L and placing it at the beginning of our pomset. Having flushed the buffer, we then evaluate the e_i with empty buffers and get back pomsets P_i and v_i . Because we can only join threads if their buffers are empty, we require that these P_i and v_i be associated with empty buffers in $\mathcal{P}_{[\]}(e_i)$. We then proceed as for the program order, and add the parallel composition of the P_i to our pomset, and compute the value $v_1 \odot v_2$. Because we just joined two empty buffers, our resulting buffer is empty:

$$\mathcal{B}_L(e_1 \odot e_2) = \{(L; (P_1 \parallel P_2), v_1 \odot v_2, [\]) \mid (P_i, v_i, [\]) \in \mathcal{P}_{[\]}(e_i)\}.$$

Finally, when we sequentially compose two commands c_1 and c_2 (assuming no forking or joining), c_2 continues executing from the buffer c_1 finished with. $\mathcal{P}(c_1)$ and $\mathcal{P}(c_2)$ are both functions of type $\text{Ls} \rightarrow \wp(\text{Pom}(\mathcal{A}_{TSO}) \times \text{Ls})$ and are not composable *qua* functions. Consequently, we need to define a composition operation capturing the above the operational intuition. This composition is the polymorphic function \curvearrowright , where $S \in \wp(\text{Pom} \times A)$ and $f \in A \rightarrow \wp(\text{Pom} \times B)$:

$$\begin{aligned} \curvearrowright &: \forall A. \forall B. \wp(\text{Pom} \times A) \rightarrow (A \rightarrow \wp(\text{Pom} \times B)) \rightarrow \wp(\text{Pom} \times B) \\ S \curvearrowright f &= \{(P; P', b) \mid (P, a) \in S \wedge (P', b) \in f(a)\}. \end{aligned}$$

Taking $A = B = \text{Ls}$, sequential composition can be expressed using as

$$\mathcal{B}_L(c_1; c_2) = \mathcal{P}_L(c_1) \curvearrowright \mathcal{P}(c_2).$$

Explicitly, this means $\mathcal{B}_L(c_1; c_2) = \{(P_1; P_2, L_2) \mid (P_1, L_1) \in \mathcal{P}_L(c_1), (P_2, L_2) \in \mathcal{P}_{L_1}(c_2)\}$. This idiom of chaining pairs of pomsets and buffers together using \curvearrowright will be useful throughout. We make \curvearrowright polymorphic so that we can handle, e.g., the case of $A = \text{Ls}$ and $B = V \times \text{Ls}$ below.

The remainder of the basic clauses are analogous to those for program order pomsets, subject to the modifications described above:

$$\begin{aligned}
\mathcal{B} &: \mathbf{VExp} \rightarrow \mathbf{Ls} \rightarrow \wp(\mathbf{Pom}(\mathcal{A}_{TSO}) \times V \times \mathbf{Ls}) \\
\mathcal{B}_L(v) &= \{(\{\delta\}, v, L)\} \\
\mathcal{B} &: \mathbf{BExp} \rightarrow \mathbf{Ls} \rightarrow \wp(\mathbf{Pom}(\mathcal{A}_{TSO}) \times \mathbf{Bool} \times \mathbf{Ls}) \\
\mathcal{B}_L(\neg e) &= \{(P, \neg b, L') \mid (P, b, L') \in \mathcal{B}_L(e)\} \\
\mathcal{P}_{L,true}(b) &= \{(P, L') \mid (P, true, L') \in \mathcal{P}_L(b)\} \\
\mathcal{P}_{L,false}(b) &= \{(P, L') \mid (P, false, L') \in \mathcal{P}_L(b)\} \\
\mathcal{B} &: \mathbf{Cmd} \rightarrow \mathbf{Ls} \rightarrow \wp(\mathbf{Pom}(\mathcal{A}_{TSO}) \times \mathbf{Ls}) \\
\mathcal{B}_L(\mathbf{skip}) &= \{(\{\delta\}, L)\} \\
\mathcal{B}_L(c_1 \parallel c_2) &= \{(L; (P_1 \parallel P_2), []) \mid (P_i, []) \in \mathcal{P}_{[]}(c_i)\} \\
\mathcal{B}_L(\mathbf{if } b \mathbf{ then } c_1 \mathbf{ else } c_2) &= (\mathcal{P}_{L,true}(b) \curvearrowright \mathcal{P}(c_1)) \cup (\mathcal{P}_{L,false}(b) \curvearrowright \mathcal{P}(c_2)) \\
\mathcal{B}_L(\mathbf{while } b \mathbf{ do } c) &= \bigcup_{n=0}^{\infty} I_L^n(b, c) \cup I_L^\omega(b, c)
\end{aligned}$$

where $I_L^0(b, c) = \mathcal{P}_{L,false}(b)$ and $I_L^{n+1}(b, c) = \mathcal{P}_{L,true}(b) \curvearrowright \mathcal{P}(c) \curvearrowright I^n(b, c)$. The set $I_L^\omega(b, c) \subseteq \mathbf{Pom}(\mathcal{A}_{TSO}) \times \{[]\}$ contains all infinite pomsets obtained through countably infinitely many unfoldings; because we can never observe the buffer at the end, we treat it as empty to simplify presentation.

We now turn our attention to flushing. The intent is that a thread can flush arbitrarily many of its writes at any point in its execution. Thus, the pomsets associated with flushes for a buffer L are the prefixes L' of L , and the resulting buffers are the remainders of L . We use $\mathit{split}(L)$ to denote these prefix-suffix pairs:

$$\begin{aligned}
\mathit{split} &: \mathbf{Ls} \rightarrow \wp(\mathbf{Pom}(\mathcal{A}_{TSO}) \times \mathbf{Ls}) \\
\mathit{split}(L) &= \{(L', L'') \mid L = L'; L''\}
\end{aligned}$$

We introduce a variant of \curvearrowright to cope with triples of pomsets, values, and buffers, and will rely on types to disambiguate the version needed in any given situation:

$$\begin{aligned}
\curvearrowright &: \forall A. \forall B. \wp(\mathbf{Pom} \times A \times B) \rightarrow (B \rightarrow \wp(\mathbf{Pom} \times B)) \rightarrow \wp(\mathbf{Pom} \times A \times B) \\
S \curvearrowright f &= \{(P; P', A, B') \mid (P, A, B) \in S \wedge (P', B') \in f(B)\}
\end{aligned}$$

We define the TSO pomsets \mathcal{P} in terms of \mathcal{B} and split . \mathcal{P} composes split and \mathcal{B} in a manner that we can flush some writes from the buffer, then evaluate e or perform c , and then flush some writes at the end:

$$\begin{aligned}
\mathcal{P}_L(e) &= \mathit{split}(L) \curvearrowright \mathcal{B}(e) \curvearrowright \mathit{split} \\
\mathcal{P}_L(c) &= \mathit{split}(L) \curvearrowright \mathcal{B}(c) \curvearrowright \mathit{split}
\end{aligned}$$

We can validate various expected equivalences by unfolding these definitions. For example, sequential composition of commands is associative, because

$$\mathcal{P}_L(c_1; (c_2; c_3)) = \mathit{split}(L) \curvearrowright \mathcal{B}(c_1) \curvearrowright \mathit{split} \curvearrowright \mathcal{B}(c_2) \curvearrowright \mathit{split} \curvearrowright \mathcal{B}(c_3) \curvearrowright \mathit{split} = \mathcal{P}_L((c_1; c_2); c_3).$$

Using the identity $\mathbf{0}; P = P$ and the fact that parallel composition of pomsets is associative, one can show that parallel composition of commands is associative, i.e., that

$\mathcal{P}_L(c_1 \parallel (c_2 \parallel c_3)) = \mathcal{P}_L((c_1 \parallel c_2) \parallel c_3)$. The parallel composition of pomsets commutes, so the parallel composition of commands commutes, i.e., $\mathcal{P}_L(c_1 \parallel c_2) = \mathcal{P}_L(c_2 \parallel c_1)$.

To illustrate the effects of flushing and the effect of buffers on reads, we consider the expression x in the presence of the buffer $L = [x := 3, y := 2]$. The triples $(P, v, L') \in \mathcal{P}_L(x)$ are of the form

$$\begin{array}{lll}
P = (x = 3) & v = 3 & L' = [x := 3, y := 2] \\
P = (x = 3 \rightarrow x := 3) & v = 3 & L' = [y := 2] \\
P = (x = 3 \rightarrow x := 3 \rightarrow y := 2) & v = 3 & L' = [] \\
P = (x := 3 \rightarrow x = v) & v \in V & L' = [y := 2] \\
P = (x := 3 \rightarrow x = v \rightarrow y := 2) & v \in V & L' = [] \\
P = (x := 3 \rightarrow y := 2 \rightarrow x = v) & v \in V & L' = []
\end{array}$$

In the first case, the resulting P denotes performing the read without also doing any flushing. In this case, because we have a write to x in the buffer, the read from x must use its value. The second and third case give rise to pomsets denoting performing the read before doing one or two flushes. In the fourth and fifth cases, we first flush the write to x from the buffer; the value v read from x is then free to range over all possible values because there are no other writes to x in the buffer. Finally, in the last case, we flush all of the writes from the buffer before reading x , and the value read from x is again unconstrained.

Definition 3.2. The **TSO pomsets** of a program p are $\mathcal{P}_{TSO}(p) = \{P \mid (P, []) \in \mathcal{P}_{[]} (p)\}$. \square

This definition selects pomsets with empty buffers from $\mathcal{P}_{[]} (p)$ because we expect programs to be run from an empty buffer and to only stop after emptying all buffers.

We illustrate the constructions by giving four example families of TSO pomsets for the Dekker program from the introduction, again assuming $v \neq 0$ and $v' \neq 0$:

$$\begin{array}{cccc}
\bar{x} := 1 & \bar{y} := 1, & \bar{x} := 1 & \bar{y} := 1, & \bar{x} := 1 & \bar{y} := 1, & \bar{x} := 1 & \bar{y} := 1. \\
\downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\
x := 1 & y := 1 & y = 0 & x = 0 & x := 1 & y := 1 & y = v & y := 1 \\
\downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\
y = 0 & x = 0 & x := 1 & y := 1 & y = v & x = 0 & x := 1 & x = v' \\
\downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\
\bar{z} := 1 & \bar{w} := 1 & \bar{z} := 1 & \bar{w} := 1 & & \bar{w} := 1 & & \\
\downarrow & \downarrow & \downarrow & \downarrow & & \downarrow & & \\
z := 1 & w := 1 & z := 1 & w := 1 & & w := 1 & &
\end{array}$$

In the first family of pomsets, we flush the writes immediately after inserting them in the buffers, while in the second, we flush the writes to x and y after reading y and x . In the third family, we flush x right after placing its write in the buffer, but fall into the *false* case of the conditional after reading some value $v \neq 0$, thus taking the **skip** branch. In the fourth pomset, we read y after placing the write $x := 1$ in the buffer, but before it gets flushed, and both threads fall into the **skip** branch.

3.3.1. Laws of parallel programming. Because parallel composition of pomsets is associative and commutative, and because sequential composition of pomsets is associative, we satisfy

many laws of parallel programming from [Bro96] and [JPR12] for free. Let $C \equiv C'$ whenever $\mathcal{P}_L(C) = \mathcal{P}_L(C')$ for all L . Our semantics satisfies:

$$\begin{aligned}
& \mathbf{skip}; c \equiv c \equiv c; \mathbf{skip} \\
& (c_1; c_2); c_3 \equiv c_1; (c_2; c_3) \\
& c_1 \parallel c_2 \equiv c_2 \parallel c_1 \\
& (c_1 \parallel c_2) \parallel c_3 \equiv c_1 \parallel (c_2 \parallel c_3) \\
& (\mathbf{if } b \mathbf{ then } c_1 \mathbf{ else } c_2); c_3 \equiv \mathbf{if } b \mathbf{ then } c_1; c_3 \mathbf{ else } c_2; c_3 \\
& \mathbf{while } b \mathbf{ do } c \equiv \mathbf{if } b \mathbf{ then } (c; \mathbf{while } b \mathbf{ do } c) \mathbf{ else skip}
\end{aligned}$$

Because we must flush buffers before every fork and join, **skip** is not a unit for parallel composition under TSO: $\mathbf{skip} \parallel c \not\equiv c$.

3.4. Executions. Our TSO pomset semantics gives an abstract account capturing families of possible executions. However, compositionality comes with its price: we associate to programs some pomsets that cannot in any real sense be “executed”. Consider for example, the pomset $\bar{x} := 2 \rightarrow x := 2 \rightarrow x = 1 \rightarrow \dots$ for the program $c = (x := 2; \mathbf{if } x = 1 \mathbf{ then } c_1 \mathbf{ else } c_2)$. In no circumstance do we expect to execute c_1 when this program is run alone, and so the above TSO pomset has, in a sense made precise later, no executional meaning. However, compositionality requires this pomset be associated with the command c , because one *could* execute c_1 if our program were instead $c \parallel x := 1$. Our notion of *execution* filters out these pomsets with no executional meaning and yields an input/output behaviour for programs built from their pomset semantics.

3.4.1. Buffered States. Our notion of execution uses buffered global state, i.e., a global state with a write buffer per thread. We execute threads individually. Each thread’s execution starts from a state with a buffer, which in combination reflect that thread’s view of shared memory. Let $\text{Locs} = \text{BLoc} \cup \text{Loc}$ be the set of all locations. We use elements of $\Sigma = (\text{BLoc} \rightarrow_{fin} (V \times \mathbb{N})/\approx) \times (\text{Loc} \rightarrow_{fin} V)$ to model the combination of a global state and a buffer, where \approx is given by $(v, n) \approx (v', m)$ if and only if both $n = m$, and $v = v'$ or $n = 0$. Because Loc and BLoc are disjoint, we identify Σ with its obvious inclusion in $\text{Locs} \rightarrow_{fin} ((V \times \mathbb{N})/\approx \cup V)$. The intuition is that if $\sigma(\bar{x}) = (v, n)$, then there are n writes to x in σ ’s write buffer, and the most recent buffer write to σ had the value v . We need to keep track of the number n of writes to x still in the buffer to know whether we should continue reading x from the buffer after a flush. We identify $(v, 0)$ and $(v', 0)$ for all v and v' because one should not be able to observe a value for a write that is no longer in the buffer, and this identification allows us to “forget” the value by setting n to 0. For $x, y, z, \dots \in \text{Locs}$ and u, v, w, \dots in the corresponding subset of $((V \times \mathbb{N})/\approx) \cup V$, we denote by $[x : u, y : v, z : w, \dots]$ the buffered state with graph $\{(x, u), (y, v), (z, w), \dots\}$. For compactness of notation, we write v_n for the equivalence class of (v, n) in $(V \times \mathbb{N})/\approx$.

3.4.2. *Footprints.* Footprints are the first step towards formalizing execution and filtering out unexecutable pomsets. Informally, a footstep $(\sigma, \tau) \in \Sigma \times \Sigma$ of an action λ is a minimal piece of state σ required to be able to perform λ , and a description τ of the effects of performing λ . For example, to perform a global write $x := v$, x must be in the domain of the initial state and present in the buffer, so $\sigma = [x : v', \bar{x} : v''_{n+1}]$ for some v' and v'' , and the result is setting the global value of x to v while removing one occurrence of x from the buffer, so $\tau = [x : v, \bar{x} : v''_n]$. Though v and v'' are unrelated, this gives the correct behaviour in the context of command pomsets because global writes to x occur in the same order as buffer writes to x . To perform a read action $x = v$, we must either have no entries for x in the buffer and have $x : v$ in the global state, or we must have x in the buffer with value v , i.e., $\bar{x} : v_n$ for some $n > 0$. We call the set of footsteps associated with an action its *footprint*. Pomsets also have footsteps and footprints.

Definition 3.3. TSO footprints for actions are given as follows:

$$\begin{aligned} \llbracket x = v \rrbracket &= \{([x : v, \bar{x} : v_0], []), ([\bar{x} : v_{n+1}], []) \mid n \in \mathbb{N}\} \\ \llbracket \bar{x} := v \rrbracket &= \{([\bar{x} : v'_n], [\bar{x} : v_{n+1}]) \mid v' \in V \wedge n \in \mathbb{N}\} \\ \llbracket x := v \rrbracket &= \{([x : v', \bar{x} : v''_{n+1}], [x : v, \bar{x} : v''_n]) \mid v', v'' \in V \wedge n \in \mathbb{N}\} \\ \llbracket \delta \rrbracket &= \{([], [])\} \end{aligned} \quad \square$$

To give footprints to pomsets, we need to know when it makes sense to combine two footsteps sequentially or in parallel. We say two states σ_1 and σ_2 are **consistent**, $\sigma_1 \uparrow \sigma_2$, if for all $x \in \text{dom}(\sigma_1) \cap \text{dom}(\sigma_2)$, $\sigma_1(x) = \sigma_2(x)$. In this case, $\sigma_1 \cup \sigma_2$ is also a state. Let $\sigma \setminus \text{dom}(\tau) = \sigma \upharpoonright_{\text{dom}(\sigma) \setminus \text{dom}(\tau)}$. Then the result of updating σ by τ is $[\sigma \mid \tau] = (\sigma \setminus \text{dom}(\tau)) \cup \tau$. To sequence the footprint (σ_1, τ_1) before the footprint (σ_2, τ_2) , we must ensure that the result of the first computation from its initial state, $[\sigma_1 \mid \tau_1]$, is consistent with the requirements σ_2 of the second computation, i.e., we must ensure $[\sigma_1 \mid \tau_1] \uparrow \sigma_2$. In this case, sequentially performing both computations requires that the initial state provide everything required by the first computation, plus everything required by the second computation not already provided by the first one. This is $\sigma_1 \cup (\sigma_2 \setminus \text{dom}(\tau_1))$. The effects of the two computations are the first's effects updated by the effects of the second one, $[\tau_1 \mid \tau_2]$. Consequently, for states (σ_i, τ_i) such that $[\sigma_1 \mid \tau_1] \uparrow \sigma_2$, we define their sequential composition to be:

$$(\sigma_1, \tau_1) \triangleleft (\sigma_2, \tau_2) = (\sigma_1 \cup (\sigma_2 \setminus \text{dom}(\tau_1)), [\tau_1 \mid \tau_2]).$$

This lifts to an associative operation on $\Sigma \times \Sigma$:

$$S_1 \triangleleft S_2 = \{(\sigma_1, \tau_1) \triangleleft (\sigma_2, \tau_2) \mid (\sigma_i, \tau_i) \in S_i \wedge [\sigma_1 \mid \tau_1] \uparrow \sigma_2\}.$$

To account for global writes occurring elsewhere during the program, we parametrize the rules assigning footprints $\llbracket P \rrbracket_\Lambda$ to pomsets P by a list Λ containing a linearisation of the pomset as a subsequence, combined with any number of other global writes that represent flushes from buffers belonging to other threads. Because these global writes are performed, in principle, by other threads, they should not affect the buffers in the footprints associated with P . The exact mechanics of how Λ and P interact in producing footprints will be made clear below. Formally, given a pomset P , we let $\text{Lin}(P)$ be the set of its linearisations. Then the clauses are parametrized by $\Lambda \in \Gamma(P) = \bigcup_{L \in \text{Ls}} \text{Lin}(P \parallel L)$, where given some global-write environment $\Lambda \in \Gamma(P)$, we identify P with its image in Λ . So Λ is a linearisation of P interspersed with global writes that will not involve buffers. In particular, when P is $\{\lambda\}$, then every $\Lambda \in \Gamma(P)$ will be of the form $\Lambda_1; \{\lambda\}; \Lambda_2$ for some unique $\Lambda_1, \Lambda_2 \in \text{Ls}$. This

fact will be important for understanding the (ACT) rule below. Because Λ is a linearisation, it totally orders the writes it contains, including those of P .

Given some $L \in \text{Ls}$, let $\llbracket L \rrbracket^*$ be inductively defined on the structure of L as follows:

$$\begin{aligned} \llbracket [] \rrbracket^* &= \{([], [])\} \\ \llbracket \{x := v\}; L \rrbracket^* &= \{([x : v'], [x : v]) \mid v' \in V\} \triangleleft \llbracket L \rrbracket^* \end{aligned}$$

The intuition here is that the foreign buffer flushes in L should only affect the global part of the state and have no effect on our buffer. Given a list $[x_1 := v_1, \dots, x_n := v_n]$, we write $\llbracket x_1 := v_1, \dots, x_n := v_n \rrbracket^*$ for $\llbracket [x_1 := v_1, \dots, x_n := v_n] \rrbracket^*$.

We will need to know if buffered states have empty buffers. We let $\zeta(\sigma)$ hold if and only if for all $x \in \text{dom}(\sigma \upharpoonright_{\text{BLoc}})$, $\sigma(x) = v_0$. In other words, $\zeta(\sigma)$ holds if and only if σ “has an empty buffer”. Thus, for example, both $\zeta([])$ and $\zeta([x : 1, \bar{y} : 2_0])$, but neither $\zeta([x : 1, \bar{x} : 2_5])$ nor $\zeta([\bar{y} : 1_1])$.

Definition 3.4. The **footprint** $\llbracket P \rrbracket_\Lambda$ of a pomset P under a global-write environment $\Lambda \in \Gamma(P)$ is the smallest set closed under the following three rules:

- (ACT) If $P = \{\lambda\}$ for some action λ and $\Lambda = \Lambda_1; P; \Lambda_2$ for some $\Lambda_1, \Lambda_2 \in \text{Ls}$, then $\llbracket P \rrbracket_\Lambda = \llbracket \Lambda_1 \rrbracket^* \triangleleft \llbracket \lambda \rrbracket \triangleleft \llbracket \Lambda_2 \rrbracket^*$.
- (SEQ) If $P = P_1; P_2$ and $\Lambda = \Lambda_1; \Lambda_2$, then $\llbracket P_1 \rrbracket_{\Lambda_1} \triangleleft \llbracket P_2 \rrbracket_{\Lambda_2} \subseteq \llbracket P \rrbracket_\Lambda$.
- (PAR) If $P = P_1 \parallel P_2$, Λ_1 is the result of deleting the read and buffer write actions of P_2 from Λ , Λ_2 is the symmetric restriction, $(\sigma_i, \tau_i) \in \llbracket P_i \rrbracket_{\Lambda_i}$, $\zeta(\sigma_i)$ and $\zeta(\tau_i)$ ($i = 1, 2$), and $\sigma_1 \uparrow \sigma_2$, then $(\sigma_1 \cup \sigma_2, \tau_1 \cup \tau_2) \in \llbracket P \rrbracket_\Lambda$. \square

This definition is inspired by Lamport’s “happened before” relation [Lam78]. In the case of (ACT), for a given $P = \{\lambda\}$ and $\Lambda \in \Gamma(P)$, the intuition is that Λ specifies that the global writes in Λ_1 happened before λ , and that λ happened before the global writes in Λ_2 . In the case of (SEQ), for $\llbracket P_i \rrbracket_{\Lambda_i}$ to be well-defined, we are implicitly assuming that $\Lambda_i \in \Gamma(P_i)$. (SEQ) tells us that the result of sequentially executing a program in the presence of global writes should be the same as executing the pieces sequentially in the presence of the appropriate subset of global writes. Finally, in (PAR), the restrictions of Λ are such that both parallel components observe *all* writes in the same order, and this is how we simulate the effects of writes to a global state. One can show that the set $\llbracket P \rrbracket_\Lambda$ is well-defined with regards to our identification of pomsets up-to-isomorphism. In particular, $\llbracket (P_1; P_2); P_3 \rrbracket_\Lambda = \llbracket P_1; (P_2; P_3) \rrbracket_\Lambda$ and $\llbracket (P_1 \parallel P_2) \parallel P_3 \rrbracket_\Lambda = \llbracket P_1 \parallel (P_2 \parallel P_3) \rrbracket_\Lambda$.

We give various results below that simplify computing the footprint of pomsets. By Proposition 3.5, a pomset has a footprint only if it is series-parallel. A pomset is **series-parallel** or **SP** if it is linear, or if it is the sequential or parallel composition of SP pomsets.

The maximal linear segments of a SP pomset are its *SP components*; any SP pomset can be uniquely decomposed into these. To compute the footprint of a pomset P , we can use Proposition 3.6 and 3.8 to first decompose P into its series-parallel components, each of which will be linear. We can then compute the footprints of these linear components using Corollary 3.7 and Proposition 3.9, and combine them using the appropriate applications of (SEQ) and (PAR).

We use these results to illustrate how global-write environments simulate the effects of global writes. Let P_1 and P_2 be the pomsets

$$\bar{x} := 2 \longrightarrow x := 2 \longrightarrow x = 3 \quad \text{and} \quad \bar{x} := 3 \longrightarrow x = 3 \longrightarrow x := 3,$$

respectively, and let $P = P_1 \parallel P_2$. Consider the global-write environments

$$\begin{aligned}\Lambda &= [\bar{x} := 2, \bar{x} := 3, x := 2, x = 3, x := 3, x = 3] \\ \Lambda_1 &= [\bar{x} := 2, x := 2, x := 3, x = 3] \\ \Lambda_2 &= [\bar{x} := 3, x := 2, x = 3, x := 3].\end{aligned}$$

We see that Λ is a global-write environment for P , and that Λ_i is the restriction of Λ given by the (PAR) rule for P_i and is again a global-write environment for P_i . To compute the footprint of P , we must begin by applying the (PAR) rule and recursively compute the footprint of P_i under Λ_i . We first consider $\llbracket P_1 \rrbracket_{\Lambda_1}$. By Corollary 3.7, it is given by

$$\llbracket \bar{x} := 2 \rrbracket \triangleleft \llbracket x := 2 \rrbracket \triangleleft \llbracket x := 3 \rrbracket^* \triangleleft \llbracket x = 3 \rrbracket,$$

where we omitted the instances of the unit $\llbracket [] \rrbracket^*$ for the \triangleleft operation. Simplifying this expression, we get that $\llbracket P_1 \rrbracket_{\Lambda_1} = \{([x : v, \bar{x} : v'_n], [x : 3, \bar{x} : 2_n]) \mid v, v' \in V \wedge n \in \mathbb{N}\}$. Despite there being no write of 3 to x in P_1 , the presence of such a global write in Λ_1 and its position in Λ_1 mean that the read $x = 3$ can be executed by P_1 . More interesting, perhaps, is the footprint $\llbracket P_2 \rrbracket_{\Lambda_2}$, and in particular, how the global write $x := 2$ does not interfere with the read $x = 3$ thanks to buffering. Again by Corollary 3.7, it is given by

$$\llbracket \bar{x} := 3 \rrbracket \triangleleft \llbracket x := 2 \rrbracket^* \triangleleft \llbracket x = 3 \rrbracket \triangleleft \llbracket x := 3 \rrbracket.$$

The subexpression $\llbracket \bar{x} := 3 \rrbracket \triangleleft \llbracket x := 2 \rrbracket^*$ simplifies to

$$\{([x : v, \bar{x} : v'_n], [x : 2, \bar{x} : 3_{n+1}]) \mid v, v' \in V \wedge n \in \mathbb{N}\} \quad (3.2)$$

Because for each $(\sigma, \tau) \in \llbracket \bar{x} := 3 \rrbracket \triangleleft \llbracket x := 2 \rrbracket^*$ we have $\tau(\bar{x}) = 3_n$ for some $n > 0$, i.e., because τ 's buffer has a write to x , the footsteps $([x : 3, \bar{x} : v_0], []) \in \llbracket x = 3 \rrbracket$ describing a read of 3 from shared memory are ignored when combining $\llbracket \bar{x} := 3 \rrbracket \triangleleft \llbracket x := 2 \rrbracket^*$ with $\llbracket x = 3 \rrbracket$ using the \triangleleft operation. Instead, $\llbracket \bar{x} := 3 \rrbracket \triangleleft \llbracket x := 2 \rrbracket^*$ gets combined with the subset $\{([\bar{x} : 3_{n+1}], []) \mid n \in \mathbb{N}\} \subseteq \llbracket x = 3 \rrbracket$ to give (3.2) again. Then combining (3.2) with $\llbracket x := 3 \rrbracket$ gives us $\llbracket P_2 \rrbracket_{\Lambda_2} = \{([x : v, \bar{x} : v'_n], [x : 3, \bar{x} : 3_n]) \mid v, v' \in V \wedge n \in \mathbb{N}\}$. Finally, we get $\llbracket P \rrbracket_{\Lambda} = \{([x : v, \bar{x} : v_0], [x : 3, \bar{x} : v_0]) \mid v \in V\}$ by combining these footprints using the (PAR) rule.

One can show by induction on the rules defining footprints that:

Proposition 3.5. *For all $P \in \text{Pom}(\mathcal{A}_{TSO})$ and $\Lambda \in \Gamma(P)$, if $\llbracket P \rrbracket_{\Lambda} \neq \emptyset$, then P is SP. \square*

The following proposition shows that the execution of a pomset P in a global-write environment Λ can be thought of as executions of parts of P interspersed with environment steps from Λ . Its corollary shows how to efficiently compute the footprint of linear pomsets.

Proposition 3.6. *If $P = P_1; \dots; P_n$ and $\Lambda = E_0; \Lambda_1; E_1; \dots; \Lambda_n; E_n \in \Gamma(P)$ such that $\Lambda_i \in \Gamma(P_i)$ and $E_i \in \text{Ls}$ for all i , then*

$$\llbracket P \rrbracket_{\Lambda} = \llbracket E_0 \rrbracket^* \triangleleft \llbracket P_1 \rrbracket_{\Lambda_1} \triangleleft \llbracket E_1 \rrbracket^* \triangleleft \dots \triangleleft \llbracket P_n \rrbracket_{\Lambda_n} \triangleleft \llbracket E_n \rrbracket^*.$$

Without loss of generality, the Λ_i and E_i can be chosen such that the first and last elements of Λ_i are in P_i . \square

Corollary 3.7. *If $P = [\lambda_1, \dots, \lambda_n]$ and $\Lambda = E_0; \{\lambda_1\}; E_1; \dots; \{\lambda_n\}; E_n \in \Gamma(P)$, then*

$$\llbracket P \rrbracket_{\Lambda} = \llbracket E_0 \rrbracket^* \triangleleft \llbracket \lambda_1 \rrbracket \triangleleft \llbracket E_1 \rrbracket^* \triangleleft \dots \triangleleft \llbracket \lambda_n \rrbracket \triangleleft \llbracket E_n \rrbracket^*. \quad \square$$

We can similarly decompose parallel compositions of pomsets when computing footprints.

Proposition 3.8. *If $P = P_1 \parallel \cdots \parallel P_n$, $\Lambda \in \Gamma(P)$, and for all i , $\Lambda_i \in \Gamma(P_i)$ is obtained from Λ by deleting the read and buffer write actions not in P_i , then*

$$\llbracket P \rrbracket_\Lambda = \llbracket P_1 \rrbracket_{\Lambda_1} \parallel \cdots \parallel \llbracket P_n \rrbracket_{\Lambda_n},$$

where the associative binary operator \parallel on subsets of $\Sigma \times \Sigma$ is given by

$$S_1 \parallel S_2 = \{(\sigma_1 \cup \sigma_2, \tau_1 \cup \tau_2) \mid (\sigma_i, \tau_i) \in S_i, \zeta(\sigma_i), \zeta(\tau_i), \sigma_1 \uparrow \sigma_2, \tau_1 \uparrow \tau_2\}. \quad \square$$

Finally, we can characterize the footprints of the above “environment steps” using the following lemma, which can be shown by induction on the length of Λ :

Proposition 3.9. *For all $\Lambda \in \mathbf{Ls}$ and $(\sigma, \tau) \in \llbracket \Lambda \rrbracket^*$, we have $x := v \in \Lambda$ for some v if and only if $x \in \text{dom}(\sigma)$. For no $\bar{x} \in \text{BLoc}$ do we have $\bar{x} \in \text{dom}(\sigma) \cup \text{dom}(\tau)$. Moreover, $\text{dom}(\sigma) = \text{dom}(\tau)$. For all $x \in \text{dom}(\sigma)$ and $v \in V$, we have $([\sigma \mid x : v], \tau) \in \llbracket \Lambda \rrbracket^*$. If $x := v$ is the maximal write to x in Λ , then $\tau(x) = v$. \square*

To better understand and characterize pomset executions, we introduce the **characteristic** $\chi_x(P) \in \mathbb{N} \times \mathbb{N}$ of a location x in a series-parallel pomset P . Intuitively, when $\chi_x(P) = (g, b)$, g is the number of global writes to x in P that do not have a buffer entry witnessing them, and b describes the number of buffer entries left to be flushed after executing P . The characteristic is inductively given on the structure of P as follows:

$$\begin{aligned} \chi_x(\{y := v\}) &= \begin{cases} (1, 0) & x = y \\ (0, 0) & x \neq y \end{cases} \\ \chi_x(\{\bar{y} := v\}) &= \begin{cases} (0, 1) & x = y \\ (0, 0) & x \neq y \end{cases} \\ \chi_x(y = v) &= (0, 0) \\ \chi_x(\delta) &= (0, 0) \\ \chi_x(P_1 \parallel P_2) &= (g_1 + g_2, b_1 + b_2) \\ \chi_x(P_1; P_2) &= (g_1 + \max\{g_2 - b_1, 0\}, b_2 + \max\{b_1 - g_2, 0\}), \end{aligned}$$

where $\chi_x(P_i) = (g_i, b_i)$. It is not hard to show that the characteristic $\chi_x(P)$ is well-defined.

As shown in Proposition 3.10 below, the concept of characteristic is closely related to that of the *differential*. Where g_x and b_x are the number of global and buffer write actions to x in a pomset P , respectively, call $\Delta_x(P) = b_x - g_x$ the **differential** of x in P .

Proposition 3.10. *For all $P \in \text{Pom}(\mathcal{A}_{TSO})$, if $\chi_x(P) = (g, b)$, then $\Delta_x(P) = b - g$. \square*

The following proposition tells us that the number of buffer entries for x in the final state of a footprint of a pomset P is determined by the differential of x in P and by the number of buffer entries for x in the initial state. For compactness, let $\sigma[\bar{x}] = n$ whenever $\sigma(\bar{x}) = v_n$.

Proposition 3.11. *For all $P \in \text{Pom}(\mathcal{A}_{TSO})$, $\Lambda \in \Gamma(P)$, $(\sigma, \tau) \in \llbracket P \rrbracket_\Lambda$, and $\bar{x} \in \text{dom}(\sigma)$, if $\bar{x} \in \text{dom}(\tau)$, then $\tau[\bar{x}] = \sigma[\bar{x}] + \Delta_x(P)$. \square*

Corollary 3.12. *Let $P \in \text{Pom}(\mathcal{A}_{TSO})$, $\Lambda \in \Gamma(P)$, $(\sigma, \tau) \in \llbracket P \rrbracket_\Lambda$, and $\chi_x(P) = (g, b)$. Then $\tau[\bar{x}] = \sigma[\bar{x}] + b - g$. In particular, when $g = 0$, then $\tau[\bar{x}] = \sigma[\bar{x}] + b$. \square*

The following lemma is useful in bounding the characteristic of P . In particular, it implies that if there exists a $\Lambda \in \Gamma(P)$ and a $(\sigma, \tau) \in \llbracket P \rrbracket_\Lambda$ such that $\zeta(\sigma)$ and $\zeta(\tau)$, then $\chi_x(P) = (0, 0)$.

Lemma 3.13. *Let $P \in \text{Pom}(\mathcal{A}_{TSO})$, $\Lambda \in \Gamma(P)$, $(\sigma, \tau) \in \llbracket P \rrbracket_\Lambda$, and $\chi_x(P) = (g, b)$. Then $\sigma[\bar{x}] \geq g$, and if $\bar{x} \notin \text{dom}(\sigma)$, then $g = 0$. Moreover, $\tau[\bar{x}] \geq b$, and if $\bar{x} \notin \text{dom}(\tau)$, then $b = 0$. \square*

Corollary 3.14. *If $(\sigma, \tau) \in \llbracket P_1 \parallel P_2 \rrbracket_\Lambda$, then $\chi_x(P_i) = (0, 0)$ and $\chi_x(P) = (0, 0)$. \square*

In certain cases, we can “read off” from P and Λ what certain values in a footprint should be. We can show by induction on the derivation of $(\sigma, \tau) \in \llbracket P \rrbracket_\Lambda$ that:

Proposition 3.15. *Let $P \in \text{Pom}(\mathcal{A}_{TSO})$, $\Lambda \in \Gamma(P)$, $(\sigma, \tau) \in \llbracket P \rrbracket_\Lambda$ and let $\chi_x(P) = (g, b)$. Let $\{\rho^{(i)}\}_{i=1..k}$ be the minimal reads $\{x = v^{(i)}\}_{i=1..k}$ from x under $<_P$.*

- (1) *If there exist i_1, \dots, i_l , $l \geq 1$, such that there are no writes to x or \bar{x} in $\bigcup_{j=1}^l \rho^{(i_j)} \downarrow_\Lambda$, then $\bar{x} \in \text{dom}(\sigma)$, there exists a v such that $v = v^{(i_1)} = \dots = v^{(i_l)}$, $\sigma(\bar{x}) = v_n$ for some $n \geq g$, and $\sigma(x) = v$ when $n = 0$.*
- (2) *If there exists a maximal global write $x := v$ to x under $<_\Lambda$ (it may be in P or not), then $\tau(x) = v$. Moreover, $x \in \text{dom}(\tau)$ if and only if there exists a global write to x in Λ .*
- (3) *If there exists a maximal buffer write $\bar{x} := v$ to x under $<_\Lambda$, then $\tau(\bar{x}) = v_n$ for some $n \geq b$. Moreover, $\bar{x} \in \text{dom}(\tau)$ if and only if there exists a global write to x or a buffer write to \bar{x} in P .*
- (4) *If there exist no buffer writes $\bar{x} := v$ and $\bar{x} \in \text{dom}(\tau)$, then where $\sigma(\bar{x}) = v_n$, we have $\tau(\bar{x}) = v_m$ for some m .*

As a corollary, we validate the intuition that the final state should be determined by the total order imposed by Λ on the writes, where β is as in (3.1):

Corollary 3.16. *For all P , $\Lambda \in \Gamma(P)$, and $(\sigma, \tau) \in \llbracket P \rrbracket_\Lambda$, we have $\tau \upharpoonright_{\text{Loc}} = \beta_{\Lambda \upharpoonright_{\mathcal{A}_w}}$. \square*

3.4.3. Executions. Let $\llbracket P \rrbracket_\Lambda^\zeta = \{(\sigma, \tau) \in \llbracket P \rrbracket_\Lambda \mid \zeta(\sigma) \wedge \zeta(\tau)\}$ be the subset of footsteps with empty buffers. The set of **TSO executions** of a finite pomset P is given by the set $\mathcal{E}(P) = \{(\sigma, [\sigma \mid \tau]) \mid \Lambda \in \text{Lin}(P), (\sigma', \tau) \in \llbracket P \rrbracket_\Lambda^\zeta, \sigma' \subseteq \sigma\}$; infinite pomsets are considered in Section 3.6. These executions take all of the states σ containing a minimal fragment σ' required to execute P to that state updated with the effects τ of P . The set of TSO executions for a program p is then $\mathcal{E}(p) = \bigcup_{P \in \mathcal{P}_{TSO}(p)} \mathcal{E}(P)$. We say that a finite pomset P and a program p are **TSO executable** if $\mathcal{E}(P)$ and $\mathcal{E}(p)$ are non-empty, respectively.

We illustrate TSO pomset executions by validating the IRIW litmus test, i.e., by showing that all writes appear in the same order to all threads. For example, starting from a state initialized to zero, executing the program

$$x := 1 \parallel y := 1 \parallel (w_1 := x; w_2 := y) \parallel (z_1 := y; z_2 := x)$$

under TSO should never give a state consistent with $[w_1 : 1, w_2 : 0, z_1 : 1, z_2 : 0]$. To show this, it is sufficient to show that the following pomset P is not executable:

$$\begin{array}{cccc} \bar{x} := 1 & \bar{y} := 1 & x = 1 & y = 1 \\ \downarrow & \downarrow & \downarrow & \downarrow \\ x := 1 & y := 1 & y = 0 & x = 0. \end{array}$$

Consider some $\Lambda \in \text{Lin}(P)$. Without loss of generality, assume $x := 1 <_\Lambda y := 1$. To get an execution, we must apply (PAR), and eventually we will need to compute $\llbracket P_4 \rrbracket_{\Lambda_4}$ where P_4 is $y = 1 \rightarrow x = 0$ and Λ_4 is such that $x := 1 <_{\Lambda_4} y := 1$. To be able to execute $y = 1$ and still get a footprint with an empty initial buffer, we need Λ_4 to satisfy $y := 1 < y = 1$. But then

$\llbracket P_4 \rrbracket_{\Lambda_4} = \llbracket x := 1, y := 1 \rrbracket^* \triangleleft \llbracket \{y = 1\} \rrbracket_{[y=1]} \triangleleft \llbracket \{x = 0\} \rrbracket_{[x=0]}$, and there are no footsteps in $\llbracket x := 1, y := 1 \rrbracket^* \triangleleft \llbracket \{y = 1\} \rrbracket_{[y=1]}$ that can be combined with those in $\llbracket \{x = 0\} \rrbracket_{[x=0]}$ to get states with empty buffers. This means we cannot combine footsteps from P_4 to get footsteps for P using the (PAR) rule, and so $\llbracket P \rrbracket_{\Lambda}$ will be empty.

As discussed in the introduction, the Dekker mutual exclusion algorithm fails under TSO. Indeed, the second pomset for Dekker on page 10 can be executed from an initial state having both x and y set to zero. To do so, we take a Λ such that $y = 0 <_{\Lambda} y := 1$ and $x = 0 <_{\Lambda} x := 1$, and apply (PAR) followed by (SEQ).

In contrast, the Peterson algorithm successfully enforces mutual exclusion under TSO. Consider the following instance of the Peterson algorithm:

$(x := 1; \mathbf{if } x = 2 \mathbf{ then } l := 1 \mathbf{ else skip}) \parallel (x := 2; \mathbf{if } x = 1 \mathbf{ then } r := 1 \mathbf{ else skip}).$

Starting from the initial state $[x : 0, l : 0, r : 0]$, one cannot execute the above under TSO and reach a state where both l and r are 1. In showing this, we can safely ignore all pomsets where a read from x appears before the global write to x , because whenever we have $\bar{x} := v \rightarrow x = v' \rightarrow x := v$ in a command's TSO pomset, we must have $v = v'$. This implies that if a thread reads x before it does the global write to x , it will take the **skip** branch of the conditional. It is then sufficient to show that the following pomset is not executable:

$$\begin{aligned} \bar{x} := 1 &\rightarrow x := 1 \rightarrow x = 2 \\ \bar{x} := 2 &\rightarrow x := 2 \rightarrow x = 1. \end{aligned}$$

Consider some $\Lambda \in \text{Lin}(P)$. Without loss of generality, assume $x := 1 <_{\Lambda} x := 2$. To get an execution, we must apply (PAR) and derive a footprint for the bottom row P_2 under some Λ_2 where $x := 1 <_{\Lambda_2} x := 2$. To compute this footprint, we must repeatedly apply (SEQ), and will eventually reach the stage where $\llbracket P_2 \rrbracket_{\Lambda_2} = \{([x : 0], [x : 2])\} \triangleleft \llbracket x = 1 \rrbracket_{[x=1]}$. But this footprint must be empty, because $[x : 0 \mid x : 2]$ is not consistent with $[x : 1]$. We thus cannot apply (PAR) and we conclude that the pomset is not executable. It follows that the Peterson algorithm enforces mutual exclusion under TSO.

3.5. Fences. We can extend the above semantics to deal with fences. This extension will not be referenced in subsequent sections, and we mention it here merely to emphasize the flexibility of our general development.

A *fence* constrains the reordering of memory actions. To capture fences, we first introduce a command **fence**. Under TSO, fences cause all actions before the fence to be observed before any actions after the fence. It is sufficient to flush the thread's buffer to ensure this, giving rise to the semantic clause $\mathcal{P}_L(\mathbf{fence}) = \{(L; \{\delta\}, [])\}$.

3.6. Infinite executions. Though we cannot capture the input-output behaviour of infinite pomsets (such executions have no “final state”), we can describe their executions in terms of executions of finite pomsets. We first characterize the “phased executions” describing infinite executions. Then we describe when a phased execution is an execution of a given infinite pomset.

The set $\mathcal{P}_{\mathcal{E}}(\sigma) \subseteq \text{Pom}(\Sigma \times \mathcal{P}_{TSO} \times \Sigma)$ of **execution-from- σ pomsets** is coinductively defined by:

- (1) if P is finite and $(\sigma, \tau) \in \mathcal{E}(P)$, then $\{(\sigma, P, \tau)\} \in \mathcal{P}_{\mathcal{E}}(\sigma)$; and
- (2) if P is finite, $(\sigma, \tau) \in \mathcal{E}(P)$, and $E \in \mathcal{P}_{\mathcal{E}}([\sigma \mid \tau])$, then $\{(\sigma, P, \tau)\}; E \in \mathcal{P}_{\mathcal{E}}(\sigma)$.

All pomsets in $\mathcal{P}_{\mathcal{E}}(\sigma)$ are linear, satisfy the finite-height property, and are series-parallel. Given an $E \in \mathcal{P}_{\mathcal{E}}(\sigma)$, its underlying TSO pomset $\pi_2(E)$ is obtained by sequentially composing its pomsets in the obvious manner. Explicitly, $\pi_2(E) = (\bigcup_{e \in E} \{e\} \times (\pi_2 \circ \Phi_E)(e), <, \Phi)$, where $\Phi(e, p) = \Phi_{\Phi_E(e)}(p)$, and $(e_0, p_0) < (e_1, p_1)$ if and only if $e_0 <_E e_1$, or $e_0 = e_1$ and $p_0 <_{\Phi(e_0)} p_1$.

An execution pomset E is a **phased execution** of $P \in \mathcal{P}_{TSO}$ if $P = \pi_2(E)$. Let $\mathcal{P}_{\mathcal{E}}(P) = \bigcup_{\sigma \in \Sigma} \{E \in \mathcal{P}_{\mathcal{E}}(\sigma) \mid \pi_2(E) = P\}$ be the set of execution pomsets of P .

This definition captures the “termination” axiom [SPA92, p. 283] that guarantees fair executions. A pomset P satisfies this axiom if for all write actions w to x and all infinite sequences $r_1 <_P r_2 <_P r_3 <_P \dots$ of read actions from x , there exists a j such that $w <_P r_j$. This axiom prohibits, e.g., the non-terminating execution of the program $y := 1 \parallel \mathbf{while} \ y = 0 \ \mathbf{do} \ \mathbf{skip}$. For every $E \in \mathcal{P}_{\mathcal{E}}(\sigma)$, $\pi_2(E)$ satisfies the termination axiom. Indeed, consider a write action w and an infinite sequence $\{r_i\}_{i \in \mathbb{N}}$ in $\pi_2(E)$. There exists an $e \in E$ such that $w \in e$. Because each of the pomsets in E is finite and E is linear, there exists an e' and an i such that $e <_E e'$ and $r_i \in e'$. Then $w <_{\pi_2(E)} r_i$ by definition.

Proposition 3.17. *We can characterize phased executions of pomsets P as follows:*

- (1) *if P is finite, then every E in $\mathcal{P}_{\mathcal{E}}(P)$ is finite;*
- (2) *if some $E \in \mathcal{P}_{\mathcal{E}}(P)$ is finite, then P is finite;*
- (3) *if P is finite, $(\sigma, \tau) \in \mathcal{E}(P)$, then $(\sigma, P, \tau) \in \mathcal{P}_{\mathcal{E}}(P)$;*
- (4) *if $\{(\sigma_1, P_1, \tau_1)\}; \dots; \{(\sigma_n, P_n, \tau_n)\} \in \mathcal{P}_{\mathcal{E}}(P)$ is finite, then $P = P_1; \dots; P_n$, $\sigma_{i+1} = [\sigma_i \mid \tau_i]$ for $1 \leq i < n$, and $(\sigma_1, [\sigma_1 \mid \tau_1 \mid \dots \mid \tau_n]) \in \mathcal{E}(P)$.*

4. SOUNDNESS AND COMPLETENESS

We show that our denotational account of TSO in Section 3 is sound and complete relative to the axiomatic account of Section 2. Soundness implies that we capture only behaviours permitted by the axiomatic account; completeness implies that we capture all behaviours permitted by the axiomatic account. Because all TSO-consistent orders are contained in TSO-consistent total orders by Corollary 2.5 and can be obtained by weakening these, it is sufficient to show that we capture all TSO-consistent total orders. We identify total orders and lists.

4.1. Soundness. We call a function $f : \text{Pom}(\mathcal{A}_{PO}) \rightarrow \wp(\mathcal{A}_{PO} \text{ list})$ **sound** when for every program p and finite pomset $P \in \mathcal{P}_{PO}(p)$, if $L \in f(P)$, then L is TSO-consistent with P .

We will construct such an f and show that it is sound in this subsection, and we will show that it is complete in the next subsection. We begin by giving an important characterization of pomsets for programs.

Lemma 4.1. *For every program p and pomset $P \in \mathcal{P}_{TSO}(p)$, there exists an order isomorphism $\omega : P \upharpoonright_{\mathcal{A}_b} \rightarrow P \upharpoonright_{\mathcal{A}_w}$ such that if $\omega(\bar{x} := v) = (y := w)$, then $y = x$ and $w = v$, and such that for all $b \in P \upharpoonright_{\mathcal{A}_b}$, we have $b <_P \omega(b)$. \square*

For programs, this means that global writes appear after the corresponding write to the buffer, that global writes occur in the same order as the writes to the buffer, and that all writes to the buffer give rise to global writes. We call a series-parallel TSO pomset for which such an ω exists for each of its SP components *well-balanced*; these “patch together” to form an ω for the entire pomset. Given a program p , all TSO pomsets in $\mathcal{P}_{TSO}(p)$ are

well-balanced, and finite $P \in \mathcal{P}_{TSO}(p)$ have a unique such ω . We assume in the rest of this section that our TSO pomsets are well-balanced.

Consider the function $U : \text{Pom}(\mathcal{A}_{TSO}) \rightarrow \text{Pom}(\mathcal{A}_{PO})$ that takes each TSO pomset to its underlying program order. It does so by first deleting all global write actions, and then relabelling all buffer write actions $\bar{x} := v$ by corresponding global write actions $x := v$. We identify all reads in P with the corresponding reads in $U(P)$ and $(x := v) = \omega(\bar{x} := v)$ with the write $x := v$ below $\bar{x} := v$. We can imagine U and the identifications as being given in the following diagram, where dashed arrows indicate identifications, solid arrows indicate the pomset orders, and the λ_j are arbitrary read actions:

$$\begin{array}{ccccccccccc}
 \dots & \rightarrow & \lambda_i & \rightarrow & \bar{x} := v & \rightarrow & \lambda_{i+2} & \rightarrow & \dots & \rightarrow & \lambda_k & \rightarrow & x := v & \rightarrow & \lambda_{k+2} & \rightarrow & \dots & & P \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & & & \downarrow & U \\
 \dots & \rightarrow & \lambda_i & \rightarrow & x := v & \rightarrow & \lambda_{i+2} & \rightarrow & \dots & \rightarrow & \lambda_k & \rightarrow & \lambda_{k+2} & \rightarrow & \dots & & & & U(P)
 \end{array}$$

By observing that the PO pomset clauses are essentially special cases of the TSO pomset clauses, we have that for all programs p , $U(\mathcal{P}_{TSO}(p)) \subseteq \mathcal{P}_{PO}(p)$. This inclusion is actually an equality, because given any program order P for p , we can construct a TSO pomset P' such that $U(P') = P$ by immediately flushing the buffer with split after every write, i.e., by replacing all occurrences of $x := v$ in P with $\bar{x} := v \rightarrow x := v$ to get P' .

Let the set $\mathcal{T}(P)$ of TSO-consistent total orders of $P \in \text{Pom}(\mathcal{A}_{TSO})$ be given by

$$\mathcal{T}(P) = \bigcup_{P' \in U^{-1}(P)} \{ \Lambda \upharpoonright_{\mathcal{A}_{PO}} \mid \Lambda \in \text{Lin}(P') \wedge \llbracket P' \rrbracket_{\Lambda}^{\zeta} \neq \emptyset \}.$$

Informally, $\mathcal{T}(P)$ captures the linearisations of TSO pomsets in $U^{-1}(P)$ that give rise to TSO executions of pomsets.

We will eventually show that \mathcal{T} is sound. First, we prove a few technical lemmas that will be useful in showing axiom (V) is satisfied. The first implies that (VA) is satisfied.

Lemma 4.2. *Let P be a TSO pomset and $\Lambda \in \Gamma(P)$ such that $\llbracket P \rrbracket_{\Lambda}^{\zeta} \neq \emptyset$. Let $(x = v)_r \in P$ be a read action such that for all buffered write actions $b <_P (x = v)_r$, its corresponding global write action $\omega(b)$ satisfies $\omega(b) <_P (x = v)_r$. If there exists a write $(x := v')_w$ maximal under $<_{\Lambda}$ amongst all writes to x in $(x = v)_r \downarrow_{\Lambda}$ and all global writes to x in $(x = v)_r \downarrow_P$ are in $(x := v')_w \downarrow_{\Lambda}$, then $v' = v$.*

Proof. By the remarks on page 13, it is sufficient to consider only linear pomsets P . Let $\Lambda_L = r \downarrow_{\Lambda} \setminus r$, Λ_R be such that $\Lambda = \Lambda_L; \Lambda_R$, $L = P \cap \Lambda_L$, and $R = P \cap \Lambda_R$. Let $(g, b) = \chi_x(P)$, then by Lemma 3.13, we have $g = b = 0$. Consider some arbitrary footstep $(\sigma, \tau) \in \llbracket P \rrbracket_{\Lambda}^{\zeta}$, and let $(\sigma_L, \tau_L) \in \llbracket L \rrbracket_{\Lambda_L}$ and $(\sigma_R, \tau_R) \in \llbracket R \rrbracket_{\Lambda_R}$ be the footsteps combined by (SEQ) to form (σ, τ) . Then because $\sigma_L \subseteq \sigma$ and $\zeta(\sigma)$, we have $\zeta(\sigma_L)$. Because L has the same number of buffer writes to \bar{x} as global writes to x , we have $\chi_x(L) = (0, 0)$ by Lemma 3.13 and Proposition 3.10. So by Corollary 3.12, we get $\tau_L[\bar{x}] = \sigma[\bar{x}] = 0$. By Proposition 3.9, we then get $\tau_L(x) = v'$, and also that $\sigma_R(x) = v$. Because $[\sigma_L \mid \tau_L] \upharpoonright \sigma_R$ by (SEQ), we conclude $v' = v$. \square

The following lemma implies that (VB) is satisfied.

Lemma 4.3. *Let P be a TSO pomset and $\Lambda \in \Gamma(P)$ such that $\llbracket P \rrbracket_{\Lambda}^{\zeta} \neq \emptyset$. Let $(x = v)_r \in P$ be a read action such that there exists a buffered write action $b' <_P (x = v)_r$ writing to x .*

Let $(\bar{x} := v')_b$ be the maximal such b' with regards to $<_P$. If $(x = v)_r <_P \omega((\bar{x} := v')_b)$, then $v = v'$.

Proof. It is again sufficient to consider only linear P . Let $\Lambda_L, \Lambda_R, L, R, (\sigma_L, \tau_L)$, and (σ_R, τ_R) be as in the proof of Lemma 4.2. By Proposition 3.11, we have $\tau_L[\bar{x}] = \sigma_L[\bar{x}] + b_x - g_x$, where b_x and g_x are the number of buffer writes and global writes to x in L , respectively. But $\sigma_L[\bar{x}] = \sigma[\bar{x}] = 0$, and well-balancedness and the last hypothesis imply $b_x > g_x$, so $\tau_L[\bar{x}] > 0$ and $\tau_L(\bar{x}) = v'_m$ for some v and $m > 0$. But $\sigma_R(\bar{x}) = v_n$ for some v and n as well. Because $[\sigma_L \mid \tau_L] \uparrow \sigma_R$ by (SEQ), we conclude $v' = v$. \square

Theorem 4.4. *The function \mathcal{T} is sound.*

Proof. Let p be an arbitrary program, and let $P' \in \mathcal{P}_{PO}(p)$, $P \in U^{-1}(P')$, and $\Lambda \in \text{Lin}(P)$ be arbitrary such that there exists a $(\sigma, \tau) \in \llbracket P \rrbracket_{\Lambda}^{\zeta}$. Let $L = \Lambda \upharpoonright_{\mathcal{A}_{PO}}$. We check the six axioms to show that L is TSO-consistent for $U(P) = P'$ from σ .

Axiom (O). This is clearly satisfied, because U leaves reads untouched, and there is an order isomorphism between global writes in P and global writes in $U(P)$ by Lemma 4.1, and linearisations preserve all elements.

Axiom (VA). Let $(x = v)_r \in U(P)$ be an arbitrary read, and assume we have some write $(x := v')_w$ maximal amongst all writes to x in $(x = v)_r \downarrow_L$, and that all writes to x in $(x = v)_r \downarrow_{U(P')}$ are in $(x := v')_w \downarrow_L$. That's to say, assume the hypotheses to (VA) hold for $(x = v)_r$. We must show $v = v'$. Because all writes to x in $(x = v)_r \downarrow_{U(P')}$ are in $(x := v')_w \downarrow_L$, this means all buffer writes to x in $(x = v)_r \downarrow_{\Lambda}$ are in $(x := v')_w \downarrow_L$. Because w appears in L , it is a global write, and so all of the corresponding global writes for the buffer writes to x in $(x = v)_r \downarrow_{\Lambda}$ are also in $(x := v')_w \downarrow_L$. Then we can apply Lemma 4.2 and conclude $v = v'$.

Axiom (VB). Let $(x = v)_r \in U(P)$ be an arbitrary read. Assume there exists a write $(x := v')_w$ to x maximal under $<_P$ amongst all writes to x in $(x = v)_r \downarrow_{U(P)}$, but that $(x = v)_r <_L (x := v')_w$. We must show that $v' = v$. That $(x := v')_w <_{U(P)} (x = v)_r$ implies $(\bar{x} := v')_w <_P (x = v)_r$, and $(x = v)_r <_L (x := v')_w$ implies $(x = v)_r <_P \omega((\bar{x} := v')_w)$. So by Lemma 4.3, we have $v' = v$.

Axiom (VC). Let $(x = v)_r \in U(P)$ be an arbitrary read. If there exist no writes $(x := v')_w <_{U(P)} (x = v)_r$ in $U(P)$, then there exist no buffered writes $(\bar{x} := v')_w <_P (x = v)_r$ in P , and so by Proposition 3.15, axiom (VC) is satisfied.

Axiom (L). Consider some arbitrary read action $r \in U(P)$ and action a such that $r <_{U(P)} a$. Because linearisation preserves order, it is sufficient to show that $r <_P a$. If a is also a read action, then U leaves a untouched, and so $r <_P a$. If a is a write action, then there exists a corresponding buffered write a' in P such that $r <_P a'$ and $\omega(a') = a$. By Lemma 4.1, $a' < \omega(a')$, so $r <_P a$ by transitivity.

Axiom (S). Lemma 4.1 implies there exists an order isomorphism between $P \upharpoonright_{\mathcal{A}_w}$ and $U(P) \upharpoonright_{\mathcal{A}_w}$, so $w <_P w'$ implies $w <_{U(P)} w'$. Linearisation preserves order, so $w <_{U(P)} w'$ implies $w <_{\Lambda} w'$. Because w and w' are global writes, $w <_L w'$ as desired.

Axiom (F). Let $\alpha_1, \alpha_2, \alpha_3 \in U(P)$ be arbitrary such that $\alpha_1 <_{U(P)} \alpha_3$, $\alpha_1 <_{U(P)} \alpha_3$, and $\alpha_2 \parallel_{U(P)} \alpha_3$. If α_1 is a write, then it is sufficient to observe that by the definition of TSO pomsets for parallel compositions of expressions or commands, for any buffered write b such that $b <_P \alpha_2$ and $b <_P \alpha_3$ for some $\alpha_2 \parallel_P \alpha_3$, we have $\omega(b) <_P \alpha_2$ and $\omega(b) <_P \alpha_3$. For then $\alpha_1 <_L \alpha_2$ and $\alpha_1 < \alpha_3$. If α_1 is a read, then we are done by (L).

Axiom (J). Let $\alpha_1, \alpha_2, \alpha_3 \in U(P)$ be arbitrary such that $\alpha_1 <_{U(P)} \alpha_3$, $\alpha_2 <_{U(P)} \alpha_3$, and $\alpha_1 \parallel_{U(P)} \alpha_2$. If α_1 and α_2 are reads, then we are done by (L), so assume without loss

of generality that α_1 is some write $(x := v)_w$. Then α_1 corresponds to some buffered write $(\bar{x} := v)_w$ in P . By the dual observation to that in (F), we have that if b is a buffered write such that for some a and c , $a <_P c$ and $b <_P a$ but $a \parallel_P b$, then $\omega(b) <_P a$. So we get $\alpha_1 <_L \alpha_3$ as desired. A symmetric argument applies when α_2 is a write. \square

4.2. Completeness. We call a function $f : \text{Pom}(\mathcal{A}_{PO}) \rightarrow \wp(\mathcal{A}_{PO} \text{ list})$ **complete** when for every program p and finite pomset $P \in \mathcal{P}_{PO}(p)$, if $L \in \mathcal{A}_{PO} \text{ list}$ is TSO-consistent with P , then $L \in f(P)$.

Our goal is to show that \mathcal{T} is complete. To do so, we first construct a pomset $s(P, L) \in U^{-1}(P)$ for any $P \in \text{Pom}(\mathcal{A}_{PO})$ and total order L that is TSO-consistent with P . Let the underlying set $s(P, L)$ be given by $\{0\} \times P \cup \{1\} \times P \upharpoonright_{\mathcal{A}_w}$. Let $\Phi_{s(P, L)}$ be given by $\Phi_{s(P, L)}((i, p)) = (\bar{x} := v)$ if both $i = 0$ and $\Phi_P(p) = (x := v)$, and $\Phi_{s(P, L)}((i, p)) = \Phi_P(p)$ otherwise. Intuitively, $<_{s(P, L)}$ merges the global writes into the program order in the places specified by L . Let $<_{s(P, L)}$ be the least strict partial order generated by the following collection of inequalities: (i) $(0, p) <_{s(P, L)} (0, p')$ if $p <_P p'$; (ii) $(1, w) <_{s(P, L)} (1, w')$ if $w \leq_{\geq_P} w'$ and $w <_L w'$; (iii) $(0, w) <_{s(P, L)} (1, w)$ if $w \in P \upharpoonright_{\mathcal{A}_w}$; (iv) $(0, p) <_{s(P, L)} (1, w)$ if $p \leq_{\geq_P} w$ and $p <_L w$; and (v) $(1, w) <_{s(P, L)} (0, p)$ if $w \leq_{\geq_P} p$ and $w <_L p$. We identify the program order actions in P with their corresponding instances in $s(P, L)$, that is, we identify all read actions r of P with $(0, r)$ in $s(P, L)$, and all global write actions w of P with the global write action $(1, w)$ in $s(P, L)$. We leverage this identification below to lift L from P to $s(P, L)$.

The proof of completeness can be broken down into three lemmas as follows:

Lemma 4.5. *If p is a program, $P \in \mathcal{P}_{PO}(p)$, and L is TSO-consistent with P , then $s(P, L) \in \mathcal{P}_{TSO}(p)$ and $U(s(P, L)) = P$.* \square

Lemma 4.6. *If $P \in \text{Pom}(\mathcal{A}_{PO})$ and L is TSO-consistent with P , there exists a $\Lambda \in \text{Lin}(s(P, L))$ such that $\Lambda \upharpoonright_{\mathcal{A}_{PO}} = L$.* \square

Lemma 4.7. *If p is a program, $P \in \mathcal{P}_{TSO}(p)$, $L_0 \in \text{Ls}$, L is TSO-consistent with $U(P) \parallel L_0$, $\Lambda \in \text{Lin}(P \parallel L_0)$, and $\Lambda \upharpoonright_{\mathcal{A}_{PO}} = L$, then $\llbracket P \rrbracket_{\Lambda}^{\zeta} \neq \emptyset$.* \square

The presence of L_0 in the statement of Lemma 4.7 lets us use an induction hypothesis in the case where P is a parallel composition of pomsets, because when we want to apply the induction hypothesis to one of the pomsets, we need to be able to reference the global writes in the other pomset.

Theorem 4.8. *The function \mathcal{T} is complete.*

Proof. Given a program p , a $P \in \mathcal{P}_{PO}(P)$ and an L that is TSO-consistent with P , we have by Lemma 4.5 that $s(P, L) \in U^{-1}(P)$. By Lemma 4.6, we have a $\Lambda \in \text{Lin}(s(P, L))$ such that $\Lambda \upharpoonright_{\mathcal{A}_{PO}} = L$. By Lemma 4.7 with $L_0 = []$, $\llbracket s(P, L) \rrbracket_{\Lambda}^{\zeta} \neq \emptyset$. So $L \in \mathcal{T}(P)$, and we conclude completeness. \square

5. RELATED WORK

Other approaches to semantics for weak memory models mostly use execution graphs and operational semantics. Execution graphs [BOS⁺11, BA08] serve to describe the executional behaviour of an entire program, an inherently non-modular approach. We see our denotational framework as offering an alternative basis for program analysis, compositional and modular by design. Boudol and Petri [BP09] gave an operational semantics framework for weak memory models that uses buffered states. Jagadeesan et al. [JPR12] adapted a fully abstract, trace-based semantics by Brookes [Bro96] to give a fully abstract denotational semantics for TSO. Higham and Kawash [HK00] gave an axiomatic semantics for TSO using linearizations. They showed their semantics to be equivalent to an abstract machine with buffers based on the informal description in [SPA92]. Their account does not consider forking and joining of processes. Demange et al. [DLZ⁺13] proposed a tractable memory model consistent with the Java Memory Model. They give an operational semantics using buffers and an axiomatic account and show the two to be equivalent. Sewell et al. [SSO⁺10] gave a TSO-like memory model for x86, formalized in HOL4. Alglave et al. [AMT14] gave a generic axiomatic framework for describing weak memory models and showed how to specialize it to various memory models. Jeffrey and Riely [JR19] give a denotational account of relaxed memory models using event structures.

Pratt [Pra86] was the first to generalize from traces to pomsets in the study of concurrency. He introduced the parallel composition operations we presented in Section 3 and he used pomsets in the study of concurrent processes. Building on Pratt’s work, Brookes [Bro15, Bro16a, Bro16b] introduced a pomset framework to study weak memory. This framework used Pratt’s parallel composition operator, and its sequential composition is a variant of Pratt’s concatenation operation. The TSO semantics given above builds on Brookes’s work. The key technical differences involve adapting pomset semantics to incorporate state equipped with abstract buffers, with careful accounting to deal properly with order relaxations allowed by TSO. Our formal axiomatization of SPARC TSO, including full treatment of forks and joins, is a crucial part of the set-up, allowing us to be precise about the relationship between our abstract denotational semantics and the more concrete and informal characterization of TSO that appears in the manual.

6. CONCLUSION

Our denotational semantics accurately captures the behaviours of SPARC TSO, and its compositionality enables us reason modularly about programs. The main strength of the pomset approach is its conceptual simplicity. Unlike trace semantics, which include irrelevant orderings of actions, our pomset semantics specifies only relevant orderings of actions. Its simplicity also makes it readily adaptable to other memory models. For example, to capture SPARC PSO, which relaxes TSO to provide only per-location global orders on writes, we conjecture that it is sufficient to replace the single buffer parameter L in the \mathcal{B}_L semantic clauses with families $\{L_x\}_{x \in \text{Loc}}$ of buffers, and to then modify the buffer flushing clauses in the obvious manner and to handle store barrier “`stbar`” commands. Memory models provided by modern processors are often much weaker than TSO and PSO. To model these, we conjecture that it is sufficient to specify the correct set of actions, and to modify the semantic clauses generating pomsets and footsteps accordingly. We believe pomset semantics provide fertile ground for future research in semantics for weak memory models.

ACKNOWLEDGEMENTS

This paper is an extended version of one [KB18] presented at MFPS XXXIII. The authors gratefully acknowledge feedback from anonymous reviewers and from André Platzer.

REFERENCES

- [AMT14] Jade Alglave, Luc Maranget, and Michael Tautschnig. Herding cats: Modelling, simulation, testing, and data mining for weak memory. *ACM Trans. Program. Lang. Syst.*, 36(2):7:1–7:74, July 2014.
- [BA08] Hans-J. Boehm and Sarita V. Adve. Foundations of the C++ concurrency memory model. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '08, pages 68–78, New York, NY, USA, 2008. ACM.
- [BOS⁺11] Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. Mathematizing C++ concurrency. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '11, pages 55–66, New York, NY, USA, 2011. ACM.
- [BP09] Gérard Boudol and Gustavo Petri. Relaxed memory models: An operational approach. In *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '09, pages 392–403, New York, NY, USA, 2009. ACM.
- [Bro96] Stephen Brookes. Full abstraction for a shared variable parallel language. *Information and Computation*, 127:145–163, 1996.
- [Bro15] Stephen Brookes. Partial order semantics and weak memory. Invited talk, Domains XII, Boole Symposium, University of Cork, August 2015.
- [Bro16a] Stephen Brookes. A denotational semantics for weak memory concurrency. Invited talk, Mathematical Foundations of Computer Science, May 2016.
- [Bro16b] Stephen Brookes. A denotational semantics for weak memory concurrency. Midlands Graduate School in the Foundations of Computing Science, April 2016.
- [DLZ⁺13] Delphine Demange, Vincent Laporte, Lei Zhao, Suresh Jagannathan, David Pichardie, and Jan Vitek. Plan B: A buffered memory model for Java. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '13, pages 329–342, New York, NY, USA, 2013. ACM.
- [HK00] Lisa Higham and Jalal Kawash. Memory consistency and process coordination for SPARC multiprocessors. In *Proceedings of the 7th International Conference on High Performance Computing*, HiPC '00, pages 355–366, London, UK, 2000. Springer-Verlag.
- [JPR12] Radha Jagadeesan, Gustavo Petri, and James Riely. Brookes is relaxed, almost! In *Proceedings of the 15th International Conference on Foundations of Software Science and Computational Structures*, FOSSACS'12, pages 180–194, Berlin, Heidelberg, 2012. Springer-Verlag.
- [JR19] Alan Jeffrey and James Riely. On Thin Air Reads: Towards an Event Structures Model of Relaxed Memory. *Logical Methods in Computer Science*, 15(1), March 2019.
- [KB18] Ryan Kavanagh and Stephen Brookes. A denotational semantics for SPARC TSO. *Electronic Notes in Theoretical Computer Science*, 336:223–239, 2018. The Thirty-third Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXIII).
- [Lam78] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, July 1978.
- [Lam79] Leslie Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Comput.*, 28(9):690–691, September 1979.
- [Pra86] Vaughan Pratt. Modeling concurrency with partial orders. *Int. J. Parallel Program.*, 15(1):33–71, February 1986.
- [SPA92] SPARC International Inc. *The SPARC Architecture Manual*. Menlo Park, CA, USA, 1992. version 8, revision SAV080SI9308.
- [SSO⁺10] Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. X86-TSO: A rigorous and usable programmer's model for x86 multiprocessors. *Commun. ACM*, 53(7):89–97, July 2010.