Brookes Is Relaxed, Almost!*

Radha Jagadeesan, Gustavo Petri, and James Riely

School of Computing, DePaul University

Abstract. We revisit the Brookes [1996] semantics for a shared variable parallel programming language in the context of the Total Store Ordering (TSO) relaxed memory model. We describe a denotational semantics that is fully abstract for Brookes’ language and also sound for the new commands that are specific to TSO. Our description supports the folklore sentiment about the simplicity of the TSO memory model.

1 Introduction

Sequential Consistency (SC), defined by Lamport [1979], enforces total order on memory operations — reads and writes to the memory — respecting the program order of each individual thread in the program. Operationally, SC is realized by traditional interleaving semantics, where shared memory is represented as a map from locations to values. For such an operational semantics, Brookes [1996] describes a fully abstract denotational view that identifies a process with its transition traces. This technique supports several approaches to program logics for shared memory concurrent programs based on separation logic (see Reynolds [2002] for an early survey). For example, O’Hearn [2007] and Brookes [2007] develop the semantics of Concurrent Separation Logic (CSL), an adaptation of separation logic to reason about concurrent threads operating on shared memory. CSL has been used to prove correctness of several concurrent data structures; for example, [Parkinson et al., 2007] and [Vafeiadis and Parkinson, 2007]. Similarly, Brookes [1996] gives the foundation for refinement approaches to prove the correctness of concurrent data structures such as in [Turon and Wand, 2011].

There are at least two motivations to consider memory models that are *weaker*, or more *relaxed*, than SC: First, modern multicore architectures permit executions that are not sequentially consistent. Second, SC disables some common compiler optimizations for sequential programs, such as the reordering of independent statements. This has led to a large body of work on relaxed memory models; Adve and Gharachorloo [1996] and Adve and Boehm [2010] provide a tutorial introduction with detailed bibliography on architectures and their impact on language design.

The operational semantics of programming languages in the presence of such relaxed memory models has now been explored. For example, Boudol and Petri [2009] explore the operational semantics of a process language with write buffers; Sevcík et al. [2011] explore the operational semantics of CLight executing with the TSO memory model; and Jagadeesan et al. [2010] describe the operational semantics of an object language under the Java Memory Model (JMM) of Manson et al. [2005].

* Research supported by NSF 0916741.

However, what has not been investigated in the literature is the denotational semantics of a language with a relaxed memory execution model. We solve this open problem in this paper.

1.1 Overview of the Paper

Our investigations are carried out in the context of the TSO memory model described in SPARC [1994], recently proposed as the model of x86 architectures by Sewell et al. [2010]. In TSO, each sequential thread carries its own write buffer that serves as the initial target of the writes executed by the thread. Thus, TSO permits executions that are not possible with SC.

To illustrate this relaxed behavior let us consider the canonical example depicted in 1 below. We have two sequential threads running in parallel. The left thread, with code $(x := 1; y)$, sets x to 1 and then reads y returning the value read. The thread on the right, with code $(y := 1; x)$, sets y to 1 and then reads and returns x . We consider that the initial state has $x = y = 0$. In the SC model, the execution where both threads read 0 is impermissible. It is however achieved by the following TSO execution with write buffers. Below we depict the initial configuration, where both threads have empty buffers (indicated by $\emptyset$) and the memory state is denoted by $\{x := 0, y := 0\}$.

$$(\{x := 0, y := 0\}, \langle \emptyset, x := 1; y \rangle \parallel \langle \emptyset, y := 1; x \rangle) \quad (1)$$

The writes performed by a thread go into its write buffer (rather than the shared memory). Thus, the above process configuration can evolve to

$$(\{x := 0, y := 0\}, \langle [x := 1], y \rangle \parallel \langle [y := 1], x \rangle)$$

where $\langle [x := 1], y \rangle$ stands for the thread with local buffer containing the assignment of 1 to x , which is not visible to the other thread, and similarly for $\langle [y := 1], x \rangle$. Now, both reads can return the value from the shared store, which is 0.

Of course, the usual SC executions are also available in a TSO model, which we demonstrate an example execution where both reads yield 1 starting from the initial process configuration. From the intermediate configuration above, both buffer updates can nondeterministically move into memory before the reads execute. Then, we get:

$$(\{x := 1, y := 1\}, \langle \emptyset, y \rangle \parallel \langle \emptyset, x \rangle)$$

leading to an execution where both reads yield 1.

We provide a precise formalization of the denotational semantics for the language of Brookes [1996] in the context of the TSO memory model. Our model includes the characteristic mfence instructions of TSO, which terminates only when the local buffer of the thread executing the instruction is empty.

Our formalization satisfies the Data Race Free (DRF) property Adve and Hill [1990]. Informally, a program is DRF if no SC execution of the program leads to a state in which a write happens concurrently with another operation on the same location. A DRF *model* requires that the programmer view of computation coincides with SC for programs that satisfy the DRF property.

Let us review [Brookes, 1996] before adapting it to a TSO setting. We use the metavariable s to stand for a shared memory, that is a partial map of variables to values, and C for commands (possibly partially executed). Brookes [1996] views the denotation of a command, $\mathcal{T}\llbracket C \rrbracket$, as a set of completed transition traces, ranged by the metavariable α , and with the form $\alpha = (s_0, s'_0) \cdot (s_1, s'_1) \dots (s_n, s'_n)$. These traces describe the interaction between a *system* and its *environment*, where the following conditions hold.

- The execution starts with the command under consideration, so $C_0 = C$.
- Transitions from s_k to s'_k model a *system step*, i.e. $\forall k \in [0, n] . s_k, C_k \rightarrow s'_k, C_{k+1}$.
- Transitions from s'_k to s_{k+1} model an *environment step*.
- The transition trace represents a terminated execution, so $C_n = \text{skip}$.

As in any sensible semantics, skip must be a unit for sequential composition.

$$\text{skip}; C; \text{skip} \equiv C \quad (2)$$

This equation motivates the *stuttering* and *mumbling* closure properties. Closure by stuttering accommodates the case when the system does not move at all but just observes the current state, i.e. $s'_i = s_i$. Closure by mumbling permits the combination of system steps that have no intervening environment step.

We can now describe the model for TSO. The type of command denotations, $\mathcal{T}\llbracket C \rrbracket$, changes to a function that takes an input buffer b and yields a set of pairs of the form $\langle \alpha, b' \rangle$ where α is a transition trace as before, and b' is the resulting buffer. The pair $\langle \alpha, b' \rangle$ is to be understood as follows, where we use P 's as metavariables for threads, and letting $\alpha = (s_0, s'_0) \cdot (s_1, s'_1) \dots (s_n, s'_n)$.

- The execution of the command starts with the input buffer b , so $P_0 = \langle b, C \rangle$.
- The state pairs still represent system steps, i.e. $\forall k \in [0, n] . s_k, P_k \rightarrow s'_k, P_{k+1}$.
- The change from s'_k to s_{k+1} still represents an environment step.
- The transition trace represents a terminated execution leaving b' as the resulting buffer, so $P_n = \langle b', \text{skip} \rangle$. Thus, the pending updates in the resulting buffer b' are yet to reach the shared memory even though there is no command left to be executed.

Our TSO semantics has analogues of the stuttering and mumbling properties for the same reasons as discussed above. In addition, it has two buffer closure properties.

Buffer update closure. Consider the program skip. Executions in $\mathcal{T}\llbracket \text{skip} \rrbracket(b)$ can result in a smaller buffer b' , because buffer updates can propagate into the shared memory. Furthermore, the change from b to b' can be done piecemeal, one buffer update at a time. Thus, skip should permit any executions of $\text{upd}(b)$ defined as the stuttering and mumbling closure of the set

$$\{ \langle (s_0, s'_0) \dots (s_n, s'_n), b' \rangle \mid b = [x_0 := v_0, \dots, x_n := v_n] \uparrow b' \ \& \ \forall i \in [0, n] . s'_i = s_i[x_i := v_i] \}$$

Each step in the above trace corresponds to the addition of one buffer update into memory. Mumbling closure introduces the possibility of multiple buffer updates in one atomic step.

To validate Equation 2, *buffer-update closure* permits data to potentially move from the buffers to shared state before and after any command executes:

$$\frac{\langle \alpha_1, b_1 \rangle \in \text{upd}(b), \langle \alpha_2, b_2 \rangle \in \mathcal{T}\llbracket C \rrbracket(b_1), \langle \alpha_3, b' \rangle \in \text{upd}(b_2)}{\langle \alpha_1 \cdot \alpha_2 \cdot \alpha_3, b' \rangle \in \mathcal{T}\llbracket C \rrbracket(b)}$$

Buffer reduction closure. The program $(x := 1; x := 1)$ simulates the program $(x := 1)$ (taking the two steps uninterruptedly), whereas the converse is not true. In buffer terms, this motivates the idea that two identical contiguous writes can be replaced by one copy of the write without leading to any new behaviors. We formalize this notion of buffer simulation as a binary relation $b_1 \triangleright b'$ and demand:

$$\frac{\langle \alpha, b_1 \rangle \in \mathcal{T} \llbracket C \rrbracket(b), b_1 \triangleright b'}{\langle \alpha, b' \rangle \in \mathcal{T} \llbracket C \rrbracket(b)}$$

Results. We present the following results.

- We describe operational and denotational semantics for the language that accommodate the extra executions permitted by TSO.
- We prove that our denotational semantics is fully abstract when we observe termination of programs.
- We use the model to identify some equational principles that hold for parallel programs under the TSO memory model.

Our results provide some formal validation for the “folklore” sentiment about the simplicity of the TSO memory model.

Organization of paper. We eschew a separate related works section since we cite the related work in context. In Section 2 we discuss the transition system for the programming language. We develop the model theory in Section 3, and prove the correspondence between operational and denotational semantics in Section 4. In Section 5, we illustrate the differences from Brookes [1996] by describing some laws that hold for programs. More detailed proof sketches are found in a fuller version of the paper.¹

2 Operational Semantics

We assume disjoint sets of *variables*, x, y and *values*, v . The only values we consider are natural numbers. In conditionals, we interpret non-zero (resp. zero) integers as true (resp. false). As usual we denote by $FV(C)$ the set of *free variables* of command C .

$$E ::= x \mid v \mid E_1 + E_2 \mid \neg E \mid \dots \quad (\text{Expression})$$

$$C, D ::= \text{skip} \mid x := E \mid C; D \mid C \parallel D \mid \text{if } E \text{ then } C \text{ else } D \quad (\text{Command})$$

$$\mid \text{while } E \text{ do } C \mid \text{local } x \text{ in } C \mid \text{await } E \text{ then } C \mid \text{mfence}$$

$$P, Q ::= \langle b, C \rangle \mid P; D \mid P \parallel Q \mid \text{new } x := v \text{ in } P \quad (\text{Process})$$

$$\mathbb{P}, \mathbb{Q} ::= [-] \mid \mathbb{P}; D \mid \mathbb{P} \parallel \mathbb{Q} \mid P \parallel Q \mid \text{new } x := v \text{ in } \mathbb{P} \quad (\text{Process context})$$

A *buffer*, $b \in \text{Buff}$, is a list of variable/value pairs, with Buff the domain of all buffers. If $b = [x_1 := v_1, \dots, x_n := v_n]$, then $\text{dom}(b) \triangleq \{x_1, \dots, x_n\}$. We write $++$ for concatenation, $\emptyset$ for the empty buffer and $b|_x$ for the buffer that results from removing x from b . We consider buffer rewrites ($\triangleright : \text{Buff} \times \text{Buff}$) that can merge contiguous identical writes, e.g. $[x_1 := v_1, \dots, x_n := v_n, x_n := v_n] \triangleright [x_1 := v_1, \dots, x_n := v_n]$.

¹ <http://fpl.cs.depaul.edu/jjriely/papers/2011brookes.pdf>

Definition 1. The relation $\triangleright : \text{Buff} \times \text{Buff}$ is defined inductively as follows.

$$\frac{}{\forall x, v. [x := v, x := v] \triangleright [x := v]} \quad \frac{}{b \triangleright b} \quad \frac{b \triangleright b_1, b_1 \triangleright b'}{b \triangleright b'} \quad \frac{b_1 \triangleright b'_1, b_2 \triangleright b'_2}{b_1 \uparrow b_2 \triangleright b'_1 \uparrow b'_2}$$

A memory, $s \in \Sigma$, is a partial map from variables to values, where Σ is the domain of all memories. We adopt several notation conventions for partial maps: if $s = \{x_1 := v_1, \dots, x_n := v_n\}$, then $\text{dom}(s) \triangleq \{x_1, \dots, x_n\}$. We write $s[x := v]$ for the memory s with the value of reference x substituted for v , and $s[b]$ to denote the memory which results from applying the updates contained in b from left to right.

As usual, we suppose a semantic function which maps expressions to functions from memories to values (in notation $\llbracket E \rrbracket s = v$). In the forthcoming transition rules, the memory passed to this function is already updated with (any) relevant buffer. The function is defined by induction on e as

$$\frac{s(x) = v}{\llbracket x \rrbracket(s) = v} \quad \frac{\llbracket E_1 \rrbracket(s) = v_1, \llbracket E_2 \rrbracket(s) = v_2}{\llbracket E_1 + E_2 \rrbracket(s) = v_1 + v_2} \quad \dots$$

In this paper, we consider that expressions evaluate atomically, following the first language considered in Brookes [1996]. There are two standard approaches to formalizing finer grain semantics; either 1. a compilation of complex expressions to a sequence of simpler commands that only perform a single read or add local variables, or 2. a direct formalization in terms of a transition system as done in the later sections of Brookes [1996]. Our presentation can accommodate either of these changes. We elide details in the interest of space.

Each sequential thread has its own buffer. Process are parallel compositions of commands. A *configuration* is a pair of a memory and a process. In Figure 1 we define the evaluation relation $s, P \rightarrow s', P'$, where $\rightarrow^*$ is the reflexive and transitive closure of the relation $\rightarrow$, and $C\{y/x\}$ denotes the command derived from C by replacing every occurrence of x with y .

The buffers grow larger in ASSIGN that adds a new update to the buffer, and become smaller in COMMIT that moves thread local buffer updates into the shared memory. CTEXT-BUF allows contiguous and identical updates in the buffer to be collapsed.

The command skip captures our notion of termination. For example, in SKIP-SEQ, the succeeding command moves into the evaluation context when the preceding process evaluates to skip. When a process terminates, its associated buffer is not necessarily empty; e.g. when $x := E$ terminates, the update to x might still be in the buffer and not yet reflected in the shared memory.

The rule FENCE implements mfence as an assertion that can terminate only when the threads buffer is empty; e.g. $x := E; \text{mfence}$ terminates only when the update to x has been moved to the shared memory, thus making it visible to every other parallel thread.

The rule PAR-CMD enables the initiation of a parallel composition only when the buffer is empty. This restriction is in conformance with Appendix J of SPARC [1994] to ensure that the newly created threads can be scheduled on different processors. For similar reasons, SKIP-PAR ensures that a parallel composition terminates only when the buffers of both parallel processes are empty.

$$\begin{array}{c}
\frac{}{s, \langle b, \text{while } E \text{ do } C \rangle \rightarrow s, \langle b, \text{if } E \text{ then } (C; \text{while } E \text{ do } C) \text{ else skip} \rangle} \text{ (WHILE)} \\
\frac{\llbracket E \rrbracket(s[b]) \neq 0}{s, \langle b, \text{if } E \text{ then } C \text{ else } D \rangle \rightarrow s, \langle b, C \rangle} \text{ (THEN)} \quad \frac{\llbracket E \rrbracket(s[b]) = 0}{s, \langle b, \text{if } E \text{ then } C \text{ else } D \rangle \rightarrow s, \langle b, D \rangle} \text{ (ELSE)} \\
\frac{y \notin \text{dom}(b) \cup \text{FV}(C)}{s, \langle b, \text{local } x \text{ in } C \rangle \rightarrow s, \text{new } y := 0 \text{ in } \langle b, C[y/x] \rangle} \text{ (LOCAL)} \\
\frac{\llbracket E \rrbracket s \neq 0 \quad s, \langle \emptyset, C \rangle \rightarrow^* s', \langle \emptyset, \text{skip} \rangle}{s, \langle \emptyset, \text{await } E \text{ then } C \rangle \rightarrow s', \langle \emptyset, \text{skip} \rangle} \text{ (AWAIT)} \quad \frac{\llbracket E \rrbracket(s[b]) = v}{s, \langle b, x := E \rangle \rightarrow s, \langle b \dashv\vdash [x := v], \text{skip} \rangle} \text{ (ASSIGN)} \\
\frac{}{s, \langle [x := v] \dashv\vdash b, C \rangle \rightarrow s[x := v], \langle b, C \rangle} \text{ (COMMIT)} \quad \frac{}{s, \langle \emptyset, \text{mfence} \rangle \rightarrow s, \langle \emptyset, \text{skip} \rangle} \text{ (FENCE)} \\
\frac{}{s, \langle \emptyset, (C \parallel D) \rangle \rightarrow s, \langle \emptyset, C \rangle \parallel \langle \emptyset, D \rangle} \text{ (PAR-CMD)} \quad \frac{}{s, \langle \emptyset, \text{skip} \rangle \parallel \langle \emptyset, \text{skip} \rangle \rightarrow s, \langle \emptyset, \text{skip} \rangle} \text{ (SKIP-PAR)} \\
\frac{s, P \rightarrow s', P'}{s, P \parallel Q \rightarrow s', P' \parallel Q} \text{ (CTXT-LEFT)} \quad \frac{s, Q \rightarrow s', Q'}{s, P \parallel Q \rightarrow s', P \parallel Q'} \text{ (CTXT-RIGHT)} \\
\frac{}{s, \text{new } y := v \text{ in } \langle b, \text{skip} \rangle \rightarrow s, \langle b|_y, \text{skip} \rangle} \text{ (SKIP-NEW)} \quad \frac{}{s, \langle \emptyset, \text{skip} \rangle; D \rightarrow s, \langle \emptyset, D \rangle} \text{ (SKIP-SEQ)} \\
\frac{b \triangleright b'}{s, \langle b, C \rangle \rightarrow s', \langle b', C \rangle} \text{ (CTXT-BUF)} \quad \frac{s, \langle b, C \rangle \rightarrow s, \langle b', C' \rangle}{s, \langle b, C; D \rangle \rightarrow s', \langle b', C'; D \rangle} \text{ (CTXT-CMD)} \\
\frac{s, P \rightarrow s', P'}{s, P; D \rightarrow s', P'; D} \text{ (CTXT-SEQ)} \\
\frac{s[y := v], P \rightarrow s', P' \quad s'(y) = v'}{s, \text{new } y := v \text{ in } P \rightarrow s'[y := s(y)], \text{new } y := v' \text{ in } P'} \text{ (CTXT-NEW)}
\end{array}$$

Fig. 1. Evaluation: $s, P \rightarrow s', P'$

Our sole use of the local construct is to provide a model of thread-local registers in the special case when C is a sequential thread. However, our more general formalization permits the description of state that is shared among parallel processes. The process context $\text{new } y := v \text{ in } \mathbb{P}$ carries the shared state of this variable. The hypothesis on the initial buffer in LOCAL ensures that any mfence in C does not affect the global x . The renaming ensures that the updates of CTXT-NEW do not affect the global x . SKIP-NEW discards any remaining updates to the local y . The commands IF and WHILE are standard. The AWAIT construct from Brookes [1996] is a conditional critical region. It provides atomic protection to the entire command C which in our use will be generally be a series of assignments. The compare-and-set instruction of TSO architectures is programmable as follows:

$$\text{cas}(x, v, w) = \text{await } 1 \text{ then if } x = v \text{ then } x := w \text{ else } x := v$$

And similarly for the other atomic instruction of TSO. Following the semantics of `cas` in x86-TSO given by Owens et al. [2009], AWAIT ensures that the buffers are empty

$$\begin{array}{cc}
\left[\begin{array}{l} flag_0 := 1; \\ \text{if } flag_1 = 0 \text{ then} \\ CS_0 \end{array} \right] \parallel \left[\begin{array}{l} flag_1 := 1; \\ \text{if } flag_0 = 0 \text{ then} \\ CS_1 \end{array} \right] & \left[\begin{array}{l} data := 1; \\ flag := 1 \end{array} \right] \parallel \left[\begin{array}{l} \text{local } r \text{ in} \\ \text{if } flag = 0 \text{ then} \\ r := data \end{array} \right] \\
\text{(a) Dekker Mutual Exclusion} & \text{(b) Safe Publication}
\end{array}$$

Fig. 2. Examples of TSO Programs

before and after the command executes and prevents buffer updates from other threads cf. the LOKD modifier of Owens et al. [2009]. While TSO does not directly support such multi-instruction atomic conditional critical regions, our semantics continues to be sound for a traditional TSO programming model, only providing the simpler `cas` and the single-word atomics alluded to above. We use this construct to permit a direct comparison with Brookes [1996] and use it (as in that work) to construct discriminating contexts in the proof of full abstraction.

Let us revise some examples of TSO in Figure 2. Dekker’s mutual exclusion algorithm 2a fails under TSO. In initial memories that contain 0 for $flag_0$ and $flag_1$, the initial write of both threads can be put in their internal buffers, remaining inaccessible to the other thread while the reads can proceed before the updates are performed. Thus, both threads can get values 0 for their respective reads and execute their critical sections concurrently. On the other hand, the standard safe publication idiom of Figure 2b is safe under TSO, since the updates of $flag$ and $data$ will proceed in order. Thus, if $flag$ is seen to have value 1 in the thread to the right, the update of 1 on $data$ has also propagated to the memory.

We end this section by remarking that our programming language satisfies the standard DRF guarantee, following traditional proofs, e.g. see Adve and Gharachorloo [1996]; Boudol and Petri [2009]; Owens et al. [2009].

3 Denotational Semantics

We use α, β etc. for elements of $(\Sigma \times \Sigma)^*$, the sequences of state pairs, and ε for the empty trace. We will consider $\mathcal{P}((\Sigma \times \Sigma)^*)$, the powerset of sequences of state pairs, with the subset ordering. Similar assumptions are made for $\mathcal{P}((\Sigma \times \Sigma)^* \times \text{Buff})$, ranged by the metavariable $\mathcal{U}$. Commands yield functions in $\text{Buff} \rightarrow \mathcal{P}((\Sigma \times \Sigma)^* \times \text{Buff})$.

Definition 2. For any $b \in \text{Buff}$, define $\mathcal{T}[[C]](b) \in \mathcal{P}((\Sigma \times \Sigma)^* \times \text{Buff})$ as follows.

$$\begin{aligned}
\mathcal{T}[[C]](b) = \{ \langle (s_0, s'_0) \dots (s_n, s'_n), b' \rangle \mid \forall k \in [0, n-1]. s_k, P_k \xrightarrow{*} s'_k, P_{k+1} \ \& \\
P_0 = \langle b, C \rangle \ \& \ P_n = \langle b', \text{skip} \rangle \}
\end{aligned}$$

Thus, we only consider transition traces where the residual left of the command is skip, albeit with potentially unfinished buffer updates.

As in [Brookes, 1996], the transition traces are closed under stuttering and mumbling, to capture the reflexivity and transitivity of the operational transition relation.

$$\frac{\langle \alpha \cdot \beta, b \rangle \in \mathcal{U}}{\langle \alpha \cdot (s, s) \cdot \beta, b \rangle \in \mathcal{U}} \text{ STUTTERING} \qquad \frac{\langle \alpha \cdot (s, s') \cdot (s', s'') \cdot \beta, b \rangle \in \mathcal{U}}{\langle \alpha \cdot (s, s'') \cdot \beta, b \rangle \in \mathcal{U}} \text{ MUMBLING}$$

Let $\mathcal{U} \in \mathcal{P}((\Sigma \times \Sigma)^* \times \text{Buff})$, we define $\mathcal{U}^\dagger$ to be the smallest set containing $\mathcal{U}$ such that is stuttering and mumbling closed.

Definition 3. Define $\text{upd}(b)$ to be the stuttering and mumbling closure of

$$\{(s_0, s'_0) \cdots (s_n, s'_n), b'\} \mid b = [x_0 := v_0, \dots, x_n := v_n] \uparrow b' \ \& \ \forall k \in [0, n] . s'_k = s_k[x_k := v_k]\}$$

And then we can deduce the inclusion: $\forall b \in \text{Buff} . \text{upd}(b) \subseteq \mathcal{T}[\llbracket \text{skip} \rrbracket](b)$.

We now let $f : \text{Buff} \rightarrow \mathcal{P}((\Sigma \times \Sigma)^* \times \text{Buff})$, and consider the following closure properties.

$$\frac{\langle \alpha_1, b_1 \rangle \in \text{upd}(b), \langle \alpha_2, b_2 \rangle \in f(b_1), \langle \alpha_3, b' \rangle \in \text{upd}(b_2)}{\langle \alpha_1 \cdot \alpha_2 \cdot \alpha_3, b' \rangle \in f(b)} \text{ BUFF-UPD}$$

$$\frac{\langle \alpha, b_1 \rangle \in f(b), b_1 \triangleright b'}{\langle \alpha, b' \rangle \in f(b)} \text{ BUFF-RED}$$

Definition 4. Let $f : \text{Buff} \rightarrow \mathcal{P}((\Sigma \times \Sigma)^* \times \text{Buff})$. Then $f^\dagger$ is the smallest function (in the pointwise order) such that:

1. For all b , $f(b)$ is stuttering and mumbling closed.
2. f is buffer-update and buffer-reduction closed.

If $f = f^\dagger$, we say f is closed. Any command yields a closed function.

Lemma 5. For every command C , $(\mathcal{T}[\llbracket C \rrbracket])^\dagger = \mathcal{T}[\llbracket C \rrbracket]$.

The following auxiliary definitions enable us to describe the equations satisfied by the transition traces semantics. Let h be a partial function from buffers to sets of transition traces such that $\forall b \in \text{Buff} . (\exists b_1 \in \text{dom}(h) . (\exists b' \in \text{Buff} . b = b' \uparrow b_1))$; then, there is a unique closed function that contains h . Formally, we overload the closure notation and write:

$$h^\dagger = \lambda b . \{ \langle \alpha \cdot \beta, b' \rangle \mid \langle \alpha, b_1 \rangle \in \text{upd}(b), \langle \beta, b' \rangle \in h(b_1) \}^\dagger$$

We define the operator $\parallel : (\Sigma \times \Sigma)^* \times (\Sigma \times \Sigma)^* \rightarrow \mathcal{P}^+((\Sigma \times \Sigma)^*)$ that yields the set of all interleavings of its arguments. We write it infix and define it inductively.

$$\alpha \parallel \varepsilon = \{ \alpha \} \qquad \frac{\beta \in \alpha_1 \parallel \alpha_2}{\beta \in \alpha_2 \parallel \alpha_1} \qquad \frac{\beta \in \alpha_1 \parallel \alpha_2}{(s_0, s'_0) \cdot \beta \in ((s_0, s'_0) \cdot \alpha_1) \parallel \alpha_2}$$

We say that the system does not alter x in $(s_0, s'_0) \cdots (s_n, s'_n)$ if $\forall k \in [1, n] . s_k(x) = s'_k(x)$ and we use $(\Sigma \times \Sigma)_{x^+}^*$ for the set of such transition sequences. We say that the environment does not alter x in $(s_0, s'_0) \cdots (s_n, s'_n)$, if $\forall k \in [1, n-1] . s'_k(x) = s_{k+1}(x)$ and we use $(\Sigma \times \Sigma)_{x^-}^*$ for the set of such transition sequences. We write $\alpha|_x = \beta|_x$ if traces α and β are identical except for the values of reference x . We write $\text{Buff}|_x$ for the set of buffers that do not have x in their domain. We let $\llbracket E=0 \rrbracket = \lambda b . \{ \langle (s, s), b \rangle \mid \llbracket E \rrbracket(s[b]) = 0 \}^\dagger$ and similarly for $\llbracket E \neq 0 \rrbracket$.

The transition traces semantics from Theorem 2 satisfies the equations of Figure 3.

Lemma 6. For every command C , $\llbracket C \rrbracket = \mathcal{T}[\llbracket C \rrbracket]$

The proof is a straightforward structural induction on the command, and we elide it in the interests of space. In this light, we are able to freely interchange $\llbracket C \rrbracket$ and $\mathcal{T}[\llbracket C \rrbracket]$ in the rest of the paper.

$$\begin{aligned}
\llbracket \text{skip} \rrbracket &= \lambda b . \{ \langle \varepsilon, b \rangle \}^\dagger \\
\llbracket C; D \rrbracket &= \lambda b . \{ \langle \alpha \cdot \beta, b' \rangle \mid \exists b_1 \in \text{Buff} . \langle \alpha, b_1 \rangle \in \llbracket C \rrbracket(b), \langle \beta, b' \rangle \in \llbracket D \rrbracket(b_1) \}^\dagger \\
\llbracket \text{mfence} \rrbracket &= \lambda b . \{ \langle \alpha, \emptyset \rangle \in \llbracket \text{skip} \rrbracket(b) \} \\
\llbracket x := E \rrbracket &= \lambda b . \{ \langle (s, s), b \vdash [x := v] \rangle \mid \llbracket E \rrbracket(s[b]) = v \}^\dagger \\
\llbracket \text{if } E \text{ then } C \text{ else } D \rrbracket &= \llbracket E=0 \rrbracket; \llbracket D \rrbracket \cup \llbracket E \neq 0 \rrbracket; \llbracket C \rrbracket \\
\llbracket \text{while } E \text{ do } C \rrbracket &= (\llbracket E \neq 0 \rrbracket; \llbracket C \rrbracket)^*; \llbracket E=0 \rrbracket \\
\llbracket \text{await } E \text{ then } C \rrbracket &= \lambda b \in \{ \emptyset \} . \{ \langle (s, s'), \emptyset \rangle \mid \llbracket E \rrbracket(s) \neq 0, \langle (s, s'), \emptyset \rangle \in \llbracket C \rrbracket(\emptyset) \}^\dagger \\
\llbracket C_1 \parallel C_2 \rrbracket &= \lambda b \in \{ \emptyset \} . \{ \langle \beta, \emptyset \rangle \mid \beta \in \beta_1 \parallel \beta_2, \forall i \in [1, 2] . \langle \beta_i, \emptyset \rangle \in \llbracket C_i \rrbracket(\emptyset) \}^\dagger \\
\llbracket \text{local } x \text{ in } C \rrbracket &= \lambda b \in \text{Buff}_{|x} . \{ \langle \beta, b'_{|x} \rangle \mid \beta \in (\Sigma \times \Sigma)_{x^+}^*, \\
&\quad \exists \langle \beta_1, b' \rangle \in \llbracket C \rrbracket(b) . \beta_1 \in (\Sigma \times \Sigma)_{x^-}^* \ \& \ \beta_{|x} = \beta_{1|x} \}^\dagger
\end{aligned}$$

Fig. 3. Denotational semantics of TSO + await

4 Full Abstraction

In this section we follow Brookes [1996] as closely as possible in order to highlight the differences caused by TSO.

The input-output relation of a program is defined using only the shared memory, i.e. the program is started with an empty buffer and the output state is observed when the buffer is empty.

Definition 7 (IO). For every command C , $\text{IO}\llbracket C \rrbracket = \{ (s, s') \mid \langle (s, s'), \emptyset \rangle \in \mathcal{T}\llbracket C \rrbracket(\emptyset) \}$

Definition 8. The trace $\alpha = (s_0, s'_0) \cdots (s_n, s'_n)$ is *Interference Free (IF)* if and only if for all $i \in [0, n-1]$ we have $s'_i = s_{i+1}$.

Notice that every $(s, s') \in \text{IO}\llbracket C \rrbracket$ arises from the mumbling closure of IF traces.

We add the following notations for technical convenience:

$$\begin{aligned}
\text{IO}\llbracket C \rrbracket / s &= \{ (s, s') \mid (s, s') \in \text{IO}\llbracket C \rrbracket \} \\
\llbracket C \rrbracket(b) / s &= \{ \langle \alpha, b' \rangle \mid \langle \alpha, b' \rangle \in \llbracket C \rrbracket(b) \ \& \ \alpha = (s, s') \cdot \alpha' \}
\end{aligned}$$

Definition 9. The operational ordering compares the IO relation of commands in all possible command contexts $\mathbb{C}[-]$,

$$C \leq_{\text{IO}} D \iff \forall \mathbb{C}[-], s . \text{FV}(\mathbb{C}[C]) \cup \text{FV}(\mathbb{C}[D]) \subseteq \text{dom}(s) \Rightarrow \text{IO}\llbracket \mathbb{C}[C] \rrbracket / s \subseteq \text{IO}\llbracket \mathbb{C}[D] \rrbracket / s$$

Definition 10. There is a natural ordering induced by the denotational semantics,

$$C \sqsubseteq D \iff \forall b, s . \text{FV}(C) \cup \text{FV}(D) \cup \text{dom}(b) \subseteq \text{dom}(s) \Rightarrow \llbracket C \rrbracket(b) / s \subseteq \llbracket D \rrbracket(b) / s$$

In the rest of this section, we prove that $\sqsubseteq$ and $\leq_{\text{IO}}$ coincide.

From Figure 3, it is evident that all the program combinators are monotone with respect to set inclusion. Thus, we deduce the following lemma.

Lemma 11 (Compositional Monotonicity). For all commands C and D ,

$$C \sqsubseteq D \Rightarrow \forall \mathbb{C}[-] . \mathbb{C}[C] \sqsubseteq \mathbb{C}[D]$$

Since $\llbracket \cdot \rrbracket$ and $\mathcal{T}\llbracket \cdot \rrbracket$ coincide by Theorem 6 we obtain:

Corollary 12 (Adequacy). *For all commands C and D , we have $C \sqsubseteq D \Rightarrow C \leq_{IO} D$.*

We now introduce some macros that we will use for the following developments. For all memories s, s' and buffer b , there is evidently an expression IS_s such that

$$\llbracket IS_s \rrbracket(s'[b]) \neq 0 \Leftrightarrow \text{dom}(s) = \text{dom}(s') \ \& \ (\forall x \in \text{dom}(s) . s(x) = s'[b](x))$$

Moreover, for all memories s, s' and buffer b , there is evidently a program consisting of a sequence of assignments $MAKE_s$ such that

$$s', \langle b, MAKE_s \rangle \rightarrow^* s, \langle \emptyset, \text{skip} \rangle$$

Finally, for each buffer b , there is evidently a program consisting of a sequence of assignments $MAKE_b$ such that for any s, b'

$$s, \langle b', MAKE_b \rangle \rightarrow^* s[b'], \langle b, \text{skip} \rangle$$

The program $MAKE_b$ can be used to encode input buffers as a command context.

Lemma 13. *For any command C and buffers b and b' , $\llbracket MAKE_{b'}; C \rrbracket(b) = \llbracket C \rrbracket(b ++ b')$.*

Proof (SKETCH). By induction on the length of b' . The base case is immediate and the inductive case follows from the definition of sequential composition.

Corollary 14. *For all C_1 and C_2 , $C_1 \not\sqsubseteq C_2 \Rightarrow \exists C, s. \llbracket C; C_1 \rrbracket(\emptyset)/_s \not\sqsubseteq \llbracket C; C_2 \rrbracket(\emptyset)/_s$.*

Proof. If $\llbracket C_1 \rrbracket(b)/_s \not\sqsubseteq \llbracket C_2 \rrbracket(b)/_s$, choose $C = MAKE_b$.

For the proof of our main result we will need to encode a context that simulates the environment of an arbitrary trace α . To that end we define the following program.

Definition 15. *Given $\alpha = (s_0, s'_0) \cdots (s_n, s'_n)$, define the command $SIMULATE_\alpha$ as*

$$\begin{aligned} SIMULATE_\alpha = & \text{await } IS_{s_0} \text{ then skip;} \\ & \text{await } IS_{s'_0} \text{ then } MAKE_{s_1}; \\ & \text{await } IS_{s'_1} \text{ then } MAKE_{s_2}; \\ & \dots \\ & \text{await } IS_{s'_{n-1}} \text{ then } MAKE_{s_n} \end{aligned}$$

Intuitively, $\llbracket SIMULATE_\alpha \rrbracket$ is given by the closure of the single trace that is “complementary” to α . Formally,

$$\llbracket SIMULATE_\alpha \rrbracket = \lambda b \in \{\emptyset\} . \{ \langle (s'_0, s_1) \cdot (s'_1, s_2) \cdots (s'_{n-1}, s_n), \emptyset \rangle \}^\dagger$$

Lemma 16. *Given α as in Theorem 15, letting $\{flag, finish\}$ be disjoint from $FV(C) \cup \text{dom}(b) \cup \bigcup_i (\text{dom}(s_i) \cup \text{dom}(s'_i))$, and considering the command context,*

$$\mathbb{C}[-] = flag := 0; finish := 0; \left(\begin{array}{c} MAKE_{b_0}; \\ [-] \end{array} \parallel SIMULATE_\alpha \right)$$

we obtain $\langle \alpha, b \rangle \in \llbracket MAKE_{b_0}; C \rrbracket(\emptyset) \iff \langle \alpha_0 \cdot (s_0, s'_n) \cdot \alpha_1, \emptyset \rangle \in \llbracket C[C] \rrbracket(\emptyset)$, where $\langle \alpha', \emptyset \rangle \in \llbracket SIMULATE_\alpha \rrbracket(\emptyset)$, $(s_0, s'_n) \in (\alpha \parallel \alpha')^\ddagger$, $\langle \alpha_0, \emptyset \rangle \in \text{upd}(flag := 0, finish := 0)$, and $\langle \alpha_1, \emptyset \rangle \in \text{upd}(b)$.

This lemma characterizes the IF traces where the final state before flushing the final buffer b is s'_n , the first state is s_0 and the trace terminates by flushing the buffer b . The variables *flag* and *finish* play essentially no role in this lemma and are included only to accommodate the use-case later.

The proof follows Brookes [1996]. For the forward direction, if $\langle \alpha, b \rangle \in \llbracket C \rrbracket(\emptyset)$, the IF trace $\langle (s_0, s'_0) \cdot (s'_0, s_1) \cdot (s_1, s'_1) \cdot (s'_1, s_2) \cdots (s'_{n-1}, s_n) \cdot (s_n, s'_n), b \rangle$ is in $\llbracket \mathbb{C}[C] \rrbracket(\emptyset)$ by interleaving. Thus, by mumbling closure, $\langle (s_0, s'_n), b \rangle \in \llbracket \mathbb{C}[C] \rrbracket(\emptyset)$. Conversely, $\langle (s_0, s'_n), b \rangle \in \llbracket \mathbb{C}[C] \rrbracket(\emptyset)$ for some b only if there is some β that can be interleaved with $(s'_0, s_1) \cdot (s'_1, s_2) \cdots (s'_{n-1}, s_n)$ to fill up the gaps between s_i and s'_i for all i . Such a trace yields α by stuttering and mumbling.

A significant difference from Brookes [1996] is that we need to check that the final buffers – since they are part of the trace semantics – coincide. To that end, we define the following program CHECK_b that “observes” all the updates of buffer b as they are performed one by one into the memory.

Definition 17. For any buffer $b = [x_1 := v_1, \dots, x_n := v_n]$ and memories s and $\bar{s}$, define:

$$\begin{aligned} \text{CHECK}_{b,s,\bar{s}} = & \text{await IS}_s \text{ then MAKE}_{\bar{s}}; \\ & \text{await IS}_{\bar{s}[x_1 := v_1]} \text{ then MAKE}_{\bar{s}}; \\ & \dots \\ & \text{await IS}_{\bar{s}[x_n := v_n]} \text{ then MAKE}_{\bar{s}} \end{aligned}$$

Informally, the program CHECK_b starts by replacing the state s for a state $\bar{s}$. In our use case, $\bar{s}$ maps every variable to values that do not appear in the trace generating the state s . The await commands are intended to observe each update from the buffer b of another thread. Upon observing each update in state $\bar{s}$ that state is reinitialized to observe the following buffer update.

Lemma 18. Let $\{flag, finish\}$ be disjoint from $\text{FV}(C) \cup \text{dom}(b) \cup_i (\text{dom}(s_i) \cup \text{dom}(s'_i))$. Let $\bar{s}$ be any memory such that the range of $\bar{s}$ is disjoint from the range of s and b . Considering the command context

$$\begin{aligned} \mathbb{C}[-] = & flag := 0; finish := 0; \\ & \left(\begin{array}{ll} \text{MAKE}_{b_0}; & D; \\ [-]; & \text{await } 1 \text{ then } flag := 1; \\ \text{if } flag \text{ then } finish := 1 & \parallel \text{await } 1 \text{ then } flag := 0; \\ & \text{CHECK}_{b,s,\bar{s}}; \\ & \text{await IS}_{\bar{s}[finish := 1]} \text{ then skip} \end{array} \right) \end{aligned}$$

there exist α_0 and α_1 such that $\langle \alpha_0, b \rangle \in \llbracket C \rrbracket(b_0)$ and $\langle \alpha_1, \varepsilon \rangle \in \llbracket D \rrbracket(\varepsilon)$ with $(s_0, s) \in (\alpha_0 \parallel \alpha_1)^\ddagger$ if and only if $\text{IO}[\llbracket \mathbb{C}[C] \rrbracket(\emptyset)] \neq \emptyset$.

Proof (SKETCH). Let $\langle \alpha_0, b \rangle \in \llbracket C \rrbracket(b_0)$ and $\langle \alpha_1, \emptyset \rangle \in \llbracket D \rrbracket$ such that $(s_0, s) \in (\alpha_0 \parallel \alpha_1)^\ddagger$. Consider the execution given by the following interleaving:

- obviously we start by executing the initial assignments of *flag* and *finish*, which are updated before spawning the new threads,
- C, D execute with an appropriate interleaving to yield the shared memory s and a buffer b' for the thread on the left of the parallel component and an empty buffer for the thread on the right, where $b' \triangleright b$,

- we then execute the first await on the right hand of the parallel composition to set $flag$ in shared memory,
- the left thread of the parallel composition observes the update on $flag$ and sets $finish$ and this update is added to the buffer of the left hand thread,
- the await on the right hand thread executes unsetting $flag$ in shared memory,
- $CHECK_{b,s,\bar{s}}$ terminates successfully since the individual awaits can be interleaved with the propagation of buffer updates from b into the shared memory,
- the update to $finish$ moves into shared memory from the buffer of left thread. Since b was exhausted in the previous step, there is no change in shared memory on $dom(\bar{s})$,
- the final await in the right thread terminates successfully because $finish$ is set and the state remains at $\bar{s}[finish := 1]$.

□

Lemma 19. $C_1 \not\sqsubseteq C_2 \Rightarrow C_1 \not\leq_{IO} C_2$

Proof (SKETCH). We have to construct a command context to distinguish the IO behavior of C_1, C_2 . By Theorem 14, we can assume that $\llbracket MAKE_{b_0}; C_1 \rrbracket(\emptyset) \not\sqsubseteq \llbracket MAKE_{b_0}; C_2 \rrbracket(\emptyset)$. Now let $\langle \alpha, b \rangle \in \llbracket C_1 \rrbracket(\emptyset) \setminus \llbracket C_2 \rrbracket(\emptyset)$. Consider the program context

$$\mathbb{C}[-] = flag := 0; finish := 0; \left(\begin{array}{c} MAKE_{b_0}; \\ [-]; \\ \text{if } flag \text{ then } finish := 1 \end{array} \parallel \begin{array}{c} SIMULATE_{\alpha}; \\ \text{await } 1 \text{ then } flag := 1; \\ \text{await } 1 \text{ then } flag := 0; \\ CHECK_{b,s,\bar{s}}; \\ \text{await } IS_{\bar{s}[finish:=1]} \text{ then skip} \end{array} \right)$$

where $flag, finish, \bar{s}, s$ and b satisfy the naming constraints of Lemmas 18 and 16. Since $\langle \alpha, b \rangle \in \llbracket C_1 \rrbracket(b_0)$, we use the forward direction of Lemmas 16 and 18 to deduce that $IO[\llbracket C[C_1] \rrbracket](\emptyset) \neq \emptyset$. Let $IO[\llbracket C[C_2] \rrbracket](\emptyset) \neq \emptyset$. Then there are α_0 and α' with $\langle \alpha_0, b \rangle \in \llbracket C_2 \rrbracket(b_0)$ and $\langle \alpha', \emptyset \rangle \in \llbracket SIMULATE_{\alpha} \rrbracket$ such that $(s_0, s) \in \{\alpha_0 \parallel \alpha'\}^{\ddagger}$. So by Theorem 18 $\langle \alpha, b \rangle \in \llbracket C_2 \rrbracket(b_0)$, which is a contradiction. □

Combining Theorem 12 and Theorem 19, we deduce that the denotational semantics $\llbracket \cdot \rrbracket$ is inequationally fully abstract.

Theorem 20 (Full Abstraction). *For any commands C and D we have*

$$C \sqsubseteq D \iff C \leq_{IO} D$$

A simple corollary of the proof of Theorem 19 is that it suffices to consider simple sequential contexts to prove inter-substitutivity of programs. For a given sequential command D and a given b , consider:

$$\mathbb{C}_D^b[-] = flag := 0; finish := 0; \left(\begin{array}{c} MAKE_b; \\ [-]; \\ \text{if } flag \text{ then } finish := 1 \end{array} \parallel D \right)$$

where *flag* and *finish* satisfy the naming constraints of Lemmas 18 and 16. Then:

$$C_1 \sqsubseteq C_2 \iff \forall D, b. (\text{IO}[\llbracket C_D^b[C_1] \rrbracket] \neq \emptyset \Rightarrow \text{IO}[\llbracket C_D^b[C_2] \rrbracket] \neq \emptyset)$$

This validates the folklore analysis of TSO programs using only sequential testers in parallel.

5 Examples and Laws

We examine some laws of parallel programming under a TSO memory model, and consider some standard TSO examples from the perspective of the denotational semantics introduced in Section 3.

Laws of parallel programming. Most of the laws inherited from Brookes [1996] hold in our setting.

$$\text{skip}; C \equiv C \equiv C; \text{skip} \quad (1)$$

$$(C_1; C_2); C_3 \equiv C_1; (C_2; C_3) \quad (2)$$

$$C_1 \parallel C_2 \equiv C_2 \parallel C_1 \quad (3)$$

$$(C_1 \parallel C_2) \parallel C_3 \equiv C_1 \parallel (C_2 \parallel C_3) \quad (4)$$

$$(\text{if } E \text{ then } C_0 \text{ else } C_1); C \equiv \text{if } E \text{ then } C_0; C \text{ else } C_1; C \quad (5)$$

$$\text{while } E \text{ do } C \equiv \text{if } E \text{ then } (C; \text{while } E \text{ do } C) \text{ else skip} \quad (6)$$

In (1) and (2) we see that sequential composition is associative with unit skip. Laws (3) and (4) say that parallel composition is commutative and associative. However, skip is not a unit for parallel composition in general, since parallel composition requires flushing the buffers before spawning the threads and when synchronizing them at the end. Instead what holds is:

$$\text{skip} \parallel C \equiv (\text{mfence}; C; \text{mfence})$$

Law (5) implies that sequential composition distributes into conditionals, and finally law (6) is the usual unrolling law for while loops. Also, The usual laws for local variables hold. If x is not free in C then:

$$\text{local } x \text{ in } C \equiv C$$

$$\text{local } x \text{ in } C; D \equiv C; \text{local } x \text{ in } D$$

$$\text{local } x \text{ in } (C \parallel D) \equiv C \parallel \text{local } x \text{ in } D$$

Thread inlining. Thread inlining is always sound in Brookes [1996], where for example the following rule holds

$$x := y; C \sqsubseteq x := y; \parallel C$$

In our setting however, this equation holds only if C does not read reference x . In the case where C reads x , C in the left hand side can potentially access newer local updates that are not available globally. In this case, a mfence is needed to validate the equation:

$$x := y; \text{mfence}; C \sqsubseteq x := y \parallel C$$

$$\begin{array}{c}
\left[\begin{array}{l} \text{local } r_0, r_1 \text{ in} \\ x := 1; \\ r_0 := x; \\ r_1 := y \end{array} \right] \parallel \left[\begin{array}{l} \text{local } r_2, r_3 \text{ in} \\ y := 1; \\ r_2 := y; \\ r_3 := x \end{array} \right] \quad [x := 1] \parallel [y := 1] \parallel \left[\begin{array}{l} \text{local } r_0, r_1 \text{ in} \\ r_0 := x; \\ r_1 := y \end{array} \right] \parallel \left[\begin{array}{l} \text{local } r_2, r_3 \text{ in} \\ r_2 := y; \\ r_3 := x \end{array} \right] \\
\text{Possible: } r_0 = r_2 = 1 \ \& \ r_1 = r_3 = 0 \qquad \text{Impossible: } r_0 = r_2 = 0 \ \& \ r_1 = r_3 = 1 \\
\text{(a) Buffer Forwarding} \qquad \qquad \qquad \text{(b) IRIW}
\end{array}$$

Fig. 4. TSO Examples

Commutation of independent statements. The TSO memory model permits reads to move ahead of previous writes on independent references. This is generally seen with the example below. Using the denotational semantics, we are able to prove the inequality, and moreover the denotations imply the existence of counterexamples to show that the inequality cannot be strengthened to an equality. Thus we get:

$$\left[\begin{array}{l} \text{local } r \text{ in} \\ r := y; \\ x := 1; \\ z := r; \end{array} \right] \sqsubseteq \left[\begin{array}{l} \text{local } r \text{ in} \\ x := 1; \\ r := y; \\ z := r; \end{array} \right] \quad \& \quad \left[\begin{array}{l} \text{local } r \text{ in} \\ x := 1; \\ r := y; \\ z := r; \end{array} \right] \not\sqsubseteq \left[\begin{array}{l} \text{local } r \text{ in} \\ r := y; \\ x := 1; \\ z := r; \end{array} \right]$$

In general, TSO does not permit writes of independent references or reads of independent reference to commute. However, a special case of this latter class of transformation can be modeled by the capability of reading one threads own writes (as shown in the example of Figure 4a). Notice in particular that the example in Figure 4a is a case of inlining of the standard IRIW example (shown in Figure 4b), which provides evidence of our previous claim that inlining is not a legal TSO transformation in general. Our denotational semantics is able to explain this relaxed behavior by means of the inequalities below. In particular, the one on the right can be proved using the inequality discussed above and the one on the left.

$$\left[\begin{array}{l} \text{local } r \text{ in} \\ x := 1; \\ r := 1; \\ z := r \end{array} \right] \sqsubseteq \left[\begin{array}{l} \text{local } r \text{ in} \\ x := 1; \\ r := x; \\ z := r \end{array} \right] \qquad \left[\begin{array}{l} \text{local } r_1, r_2 \text{ in} \\ r_2 := y; \\ x := 1; \\ r_1 := 1; \\ z_0 := r_1; z_1 := r_2 \end{array} \right] \sqsubseteq \left[\begin{array}{l} \text{local } r_1, r_2 \text{ in} \\ x := 1; \\ r_1 := x; \\ r_2 := y; \\ z_0 := r_1; z_1 := r_2 \end{array} \right]$$

6 Conclusion

We describe how to modify the Brookes semantics for a shared variable parallel programming language Brookes [1996] to address the TSO relaxed memory model. We view our results as the foundations towards two developments: (a) separation logics for relaxed memory models, and (b) refinement theory for relaxed memory models.

References

- Adve, S.V., Boehm, H.-J.: Memory models: a case for rethinking parallel languages and hardware. *Commun. ACM* 53, 90–101 (2010) ISSN 0001-0782
- Adve, S.V., Gharachorloo, K.: Shared memory consistency models: A tutorial. *Computer* 29(12), 66–76 (1996)
- Adve, S.V., Hill, M.D.: Weak ordering - a new definition. In: ISCA, pp. 2–14 (1990)
- Boudol, G., Petri, G.: Relaxed memory models: an operational approach. In: POPL, pp. 392–403 (2009)
- Brookes, S.: A semantics for concurrent separation logic. *Theor. Comput. Sci.* 375(1-3), 227–270 (2007)
- Brookes, S.: Full abstraction for a shared-variable parallel language. *Inf. Comput.* 127(2), 145–163 (1996)
- Jagadeesan, R., Pitcher, C., Riely, J.: Generative Operational Semantics for Relaxed Memory Models. In: Gordon, A.D. (ed.) ESOP 2010. LNCS, vol. 6012, pp. 307–326. Springer, Heidelberg (2010)
- Lamport, L.: How to make a multiprocessor computer that correctly executes multiprocess program. *IEEE Trans. Comput.* 28(9), 690–691 (1979)
- Manson, J., Pugh, W., Adve, S.V.: The java memory model. In: POPL, pp. 378–391 (2005)
- O’Hearn, P.W.: Resources, concurrency, and local reasoning. *Theor. Comput. Sci.* 375(1-3), 271–307 (2007)
- Owens, S., Sarkar, S., Sewell, P.: A Better x86 Memory Model: x86-TSO. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 391–407. Springer, Heidelberg (2009)
- Parkinson, M.J., Bornat, R., O’Hearn, P.W.: Modular verification of a non-blocking stack. In: POPL, pp. 297–302 (2007)
- Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: LICS, pp. 55–74 (2002)
- Sevcík, J., Vafeiadis, V., Nardelli, F.Z., Jagannathan, S., Sewell, P.: Relaxed-memory concurrency and verified compilation. In: POPL, pp. 43–54 (2011)
- Sewell, P., Sarkar, S., Owens, S., Nardelli, F.Z., Myreen, M.O.: x86-TSO: a rigorous and usable programmer’s model for x86 multiprocessors. *Commun. ACM* 53(7), 89–97 (2010)
- Inc. CORPORATE. SPARC. The SPARC Architecture Manual (version 9). Prentice-Hall, Inc., Upper Saddle River (1994)
- Turon, A.J., Wand, M.: A separation logic for refining concurrent objects. *SIGPLAN Not.* 46, 247–258 (2011) ISSN 0362-1340
- Vafeiadis, V., Parkinson, M.J.: A Marriage of Rely/Guarantee and Separation Logic. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703, pp. 256–271. Springer, Heidelberg (2007)