
TSO GAMES - ON THE DECIDABILITY OF SAFETY GAMES UNDER THE TOTAL STORE ORDER SEMANTICS

STEPHAN SPENGLER ^a AND SANCHARI SIL ^b

^a Uppsala University, Uppsala, Sweden
e-mail address: stephan.spengler@it.uu.se

^b Chennai Mathematical Institute, Chennai, India
e-mail address: sanchari@cmi.ac.in

ABSTRACT. We consider an extension of the classical Total Store Order (TSO) semantics by expanding it to turn-based 2-player safety games. During her turn, a player can select any of the communicating processes and perform its next transition. We consider different formulations of the safety game problem depending on whether one player or both of them transfer messages from the process buffers to the shared memory. We give the complete decidability picture for all the possible alternatives.

1. INTRODUCTION

Most modern architectures, such as Intel x86 [Int12], SPARC [spa94], IBM's POWER [IBM21], and ARM [ARM14], implement several relaxations and optimisations that reduce the latency of memory accesses. This has the effect of breaking the Sequential Consistency (SC) assumption [Lam79]. SC is the classical strong semantics for concurrent programs that interleaves the parallel executions of processes while maintaining the order in which instructions were issued. Programmers usually assume that the execution of programs follows the SC model. However, this is not true when we consider concurrent programs running on modern architectures. In fact, even simple programs such as mutual exclusion and producer-consumer protocols, that are correct under SC, may exhibit erroneous behaviours. This is mainly due to the relaxation of the execution order of the instructions. For instance, a standard relaxation is to allow the reordering of reads and writes of the same process if the reads have been issued after the writes and they concern different memory locations. This relaxation can be implemented using an unbounded perfect FIFO queue/buffer between each process and the memory. These buffers are used to store delayed writes. The corresponding model is called Total Store Ordering (TSO) and corresponds to the formalisation of SPARC and Intel x86 [OSS09, SSO⁺10].

A process of a concurrent program is modelled as a finite-state transition system. Each transition is labelled with an instruction that describes how it interacts with the shared memory. For instance, the process can read a value from or write a value to a shared variable. The data domain of these instructions is assumed to be finite. More complex instructions like comparisons or arithmetic operations, and infinite data structures managed locally by the process are abstracted through the nondeterminism of the process. An example of a concurrent program is given in Figure 1, which shows the Dekker protocol for mutual




FIGURE 1. A concurrent program $\mathcal{P} = \langle \text{Proc}^1, \text{Proc}^2 \rangle$ modelling (a simplified version of) the Dekker protocol.

exclusion. Under SC semantics, it is not possible for both processes to simultaneously reach the critical sections q_3 and r_3 , respectively.

In TSO, an unbounded buffer is associated with each process. When a process executes a write operation, this write is appended to the end of the buffer of that process. A pending write operation on the variable x at the head of a buffer can be deleted in a non-deterministic manner. This updates the value of the shared variable x in the memory. To perform a read operation on a variable x , the process first checks its buffer for a pending write operation on the variable x . If such a write exists, then the process reads the value written by the newest pending write operation on x . Otherwise, the process fetches the value of the variable x from the memory.

Under TSO, both processes of the Dekker protocol in Figure 1 can reach the critical sections q_3 and r_3 simultaneously, thus violating mutual exclusion. This can happen since after executing $\text{wr}(x, 1)$, the operation only resides in the buffer of Proc¹ and is invisible to Proc². The converse holds for the write instruction of Proc². Thus, both processes are able to read the initial value 0 from the variable y or x , respectively.

In general, verification of programs running under TSO is challenging due to the unboundedness of the buffers. In fact, the induced state space of a program under TSO may be infinite even if the program itself is a finite-state system.

The reachability problem for programs under TSO checks whether a given program state is reachable during program execution. It is also called safety problem, in case the target state is considered to be a bad state. It has been shown decidable using different alternative semantics for TSO (e.g. [ABBM10, AAC⁺12, AABN18]). These alternative semantics, such as the load-buffer semantics in contrast to the traditional store-buffer semantics, provide different formalisations of TSO. While they differ in how they model specific aspects of memory behaviour, they are equivalent with respect to the reachability of program states.

Furthermore, it has been shown in [ABBM10] that lossy channel systems (see e.g., [AJ96, FS01, AJ94, Sch02]) can be simulated by programs running under TSO. This entails that the reachability problem for programs under TSO is non-primitive recursive and that the repeated reachability problem is undecidable. This is an immediate consequence of the fact that the reachability problem for lossy channel is non-primitive recursive [Sch02] and that the repeated reachability problem is undecidable [AJ94]. The termination problem for programs running under TSO has been shown to be decidable in [Ati20] using the framework of well-structured transition systems [Fin87, FS01, AJ96].

The authors of [BDM13, BMM11] consider the robustness problem for programs running under TSO. This problem consists in checking whether, for any given TSO execution, there is an equivalent SC execution of the same program. Two executions are declared equivalent by the robustness criterion if they agree on (1) the order in which instructions are executed

within the same process (i.e., program order), (2) the write instruction from which each read instruction fetches its value (i.e., read-from relation), and (3) the order in which write instruction on the same variable are committed to memory (i.e., store ordering). The problem of checking whether a program is robust has been shown to be PSPACE-complete in [BDM13]. A variant of the robustness problem which is called persistence, declares that two runs are equivalent if (1) they have the same program order and (2) all write instructions reach the memory in the same order. Checking the persistency of a program under TSO has been shown to be PSPACE-complete in [AAP15]. Observe that the persistency and robustness problems are stronger than the safety problem (i.e., if a program is safe under SC and robust/persistent, then it is also safe under TSO).

Due to the non-determinism of the buffer updates, the buffers associated with each process under TSO appear to exhibit a lossy behaviour. Previously, games on lossy channel systems (and more general on monotonic systems) were studied in [ABd08]. Unfortunately these results are not applicable / transferable to programs under TSO whose induced transition systems are not monotone [ABBM10].

In this paper, we consider a natural continuation of the works on both the study of the decidability/complexity of the formal verification of programs under TSO and the study of games on concurrent systems. This is further motivated by the fact that formal games provide a framework to reason about a system's behaviour, which can be leveraged in control model checking, for example in controller synthesis problems. In particular, games can model the adversarial nature of certain problems where one component tries to achieve a goal while another tries to prevent it. This is especially relevant in the context of concurrent programs running under relaxed memory models like TSO, where the non-deterministic behaviour of memory can be seen as an adversary to the correctness of the program.

One specific problem that can be modeled by our approach is the controller synthesis problem (e.g. [PR89, KV97, KV00]). In this context, the goal is to design a controller (player A) that ensures a system (player B) behaves correctly despite the uncertainties introduced by the memory model. The game framework allows us to formally define the objectives of the controller and the system, and to analyse whether a winning strategy exists for the controller. If such a strategy exists, it corresponds to a correct implementation of the controller that guarantees the desired behaviour of the system [BS04, AVW03].

Another problem that can be modeled by our approach is the verification of safety properties in concurrent programs. By framing the verification problem as a game, we can leverage game-theoretic techniques to determine whether a program can reach an undesirable state (a losing condition for player A). This approach provides a clear and formal method to analyse the safety of concurrent programs under TSO and other relaxed memory models.

Additionally, our approach lays the foundation for future research. For instance, by adding weights to the instructions/transitions, we can model fence insertion problems [AAP15]. This could be done by giving a player the choice to take a memory fence instruction or not, but we assign a penalty cost to each memory fence transition. An optimal, that is lowest-cost, winning strategy corresponds to the minimal amount of memory fences that need to be inserted into the program to ensure memory consistency. This is crucial for optimising the performance of concurrent programs under relaxed memory models.

In more detail, we consider (safety) games played on the transition systems induced by programs running under TSO. Given a program under TSO, we construct a game in which two players, A and B, take turns executing instructions. Player B aims to reach a given set of final configurations, while player A tries to prevent this, making it a reachability

| | | Player A: | | | |
|-----------|--------|-----------|--------|---------|--------|
| | | always | before | after | never |
| Player B: | always | I (d) | | | |
| | before | | II (d) | | |
| | after | | | III (u) | |
| | never | | | | IV (d) |

FIGURE 2. Groups of TSO games, where players A and B are allowed to update the buffer: always, before their own move, after their own move, or never. The games in group I (light grey), II (white) and IV (dark grey) are decidable (d), the games in group III (medium grey) are undecidable (u).

game from the perspective of player B. The turn structure determines which player executes the next instruction, but it does not specify how memory updates occur. To address this, we allow players to update memory by removing pending writes from the buffer between instruction executions.

In some variants, only player A (considered to be the *good* player) has control over memory updates. These correspond to the problem of synthesizing an update controller that ensures correct execution despite an adversarial opponent. In other variants, only player B (the *bad* player) has control over updates, modelling the problem of synthesising a process or program under a non-cooperative update mechanism. In both cases, the precise timing of updates remains unclear — whether they occur before, after, or both before and after a player's turn. Additionally, there may be scenarios where it is beneficial to allow both players to update memory, depending on the problem setting.

Considering all these possibilities results in 16 different TSO game variants. While not all are equally relevant in practice, we analyse them for completeness, as they may provide insights for future research and yet unknown applications. We divide these 16 games into four different groups, depending on their decidability results.

- Group I (7 games) can be reduced to TSO games with 2-bounded buffers.
- Group II (1 game) can be reduced to TSO games with bounded buffers.
- Group III (7 games) can simulate perfect channel systems.
- Group IV (1 game) can be reduced to a finite game without buffers.

This classification is shown in Figure 2. Of these four groups, only Group III is undecidable, the others each reduce to a finite game and are thus decidable. Coincidentally, the undecidable group contains all of the most interesting variants, where exactly one of the two players can update memory.

Finally, we establish the exact computational complexity for the decidable games. In fact, we show that the problem is EXPTIME-complete. We prove EXPTIME-hardness by a reduction from the problem of acceptance of a word by a polynomially bounded alternating Turing machine [CS76, CKS81]. This result holds even in the case for games played on a single-process concurrent program following SC semantics. To prove EXPTIME-membership, we show that it is possible to compute the winning regions for the players in exponential time. These results are surprising given the non-primitive recursive complexity of the reachability problem for programs under TSO and the undecidability of the repeated reachability problem.

This paper is an extended and revised version of the conference paper [SS23]. We present the formal proof of the lower complexity bound for the decidable games. These now also include the newly defined SC games, which are games played on programs following SC semantics. We utilise the framework of bisimulations to lift complexity results from SC games to TSO games, and to streamline the EXPTIME-completeness proof of group IV. Furthermore, the appendix presents a formal proof of the undecidability result for the games in group III. We generally revised all sections of this paper. It contains clarifications, more elaborate explanations, and additional examples and figures.

Related Works. In addition to the related work mentioned in the introduction on the decidability / complexity of the verification problems of programs running under TSO, there have been some works on parameterized verification of programs running under TSO. The problem consists in verifying a concurrent program regardless of the number of involved processes (which are identical finite-state systems). The parameterised reachability problem of programs running under TSO has been shown to be decidable in [AABN16, AABN18]. While this problem for concurrent programs performing only read and writing operations (no atomic read-write instructions) is PSPACE-complete [AAR20]. This result has been recently extended to processes manipulating abstract data types over infinite domains [AAF⁺23]. Checking the robustness of a parameterised concurrent system is decidable and EXPSPACE-hard [BDM13].

As far as we know this is the first work that considers the game problem for programs running under TSO. The proofs and techniques used in this paper are different from the ones used to prove decidability / complexity results for the verification of programs under TSO except the undecidability result which uses some ideas from the reduction from the reachability problem for lossy channel systems to its corresponding problem for programs under TSO [ABBM10]. However, our undecidability proof requires us to implement a protocol that detects lossiness of messages in order to turn the lossy channel system into a perfect one (which is the most intricate part of the proof).

2. PRELIMINARIES

2.1. Transition Systems. A *(labelled) transition system* is a triple $\langle C, L, \rightarrow \rangle$, where C is a set of *configurations*, L is a set of *labels*, and $\rightarrow \subseteq C \times L \times C$ is a *transition relation*. We usually write $c_1 \xrightarrow{\text{label}} c_2$ if $\langle c_1, \text{label}, c_2 \rangle \in \rightarrow$. Furthermore, we write $c_1 \rightarrow c_2$ if there exists some *label* such that $c_1 \xrightarrow{\text{label}} c_2$. A *run* π of \mathcal{T} is a sequence of transitions $c_0 \xrightarrow{\text{label}_1} c_1 \xrightarrow{\text{label}_2} c_2 \dots \xrightarrow{\text{label}_n} c_n$. It is also written as $c_0 \xrightarrow{\pi} c_n$. A configuration c' is *reachable* from a configuration c , if there exists a run from c to c' .

For a configuration c , we define $\text{Pre}(c) := \{c' \mid c' \rightarrow c\}$ and $\text{Post}(c) := \{c' \mid c \rightarrow c'\}$. We extend these notions to sets of configurations C' with $\text{Pre}(C') := \bigcup_{c \in C'} \text{Pre}(c)$ and $\text{Post}(C') := \bigcup_{c \in C'} \text{Post}(c)$.

An *unlabelled transition system* is a transition system without labels. Formally, it is defined as a TS with a singleton label set. In this case, we omit the labels.

2.2. Alternating Turing Machines. Let $\mathcal{A} = \langle \Sigma, Q, q_0, q_F, Q_\exists, Q_\forall, \delta \rangle$ be an alternating Turing Machine, where Σ is a finite alphabet, Q is a finite set of states partitioned into existential states Q_\exists and universal states Q_\forall , $q_0 \in Q$ is the *initial state*, $q_F \in Q_\forall$ is the *accepting state*, and $\delta \subseteq Q \times \Sigma \times Q \times \Sigma \times \{L, R\}$ is a transition relation. The alphabet contains the blank symbol \sqcup . Let $\langle q, \sigma, q', \sigma', D \rangle \in \delta$ be a transition. If machine \mathcal{A} is in state q and its head reads letter $\sigma \in \Sigma$, then it replaces the contents of the current cell with the letter σ' , moves the head in direction D ($D = L$ for left and $D = R$ for right) and changes control to state q' . We assume that q_F has no outgoing transitions.

A configuration of \mathcal{A} is a triple $c = \langle q, i, \omega \rangle$, where $q \in Q$, $i \in \mathbb{Z}$ and $\omega \in \Sigma^\mathbb{Z}$. It represents the situation where \mathcal{A} is in state q , the tape contains the (in both directions infinite) word ω and the head is at the i -th position of the tape. A *successor* of c is a configuration c' that can be obtained from c by executing a transition as described above.

Let $C_0 := \{ \langle q_F, i, \omega \rangle \mid i \in \mathbb{Z}, \omega \in \Sigma^\mathbb{Z} \}$ and for all $k > 0$:

$$C_k := \{ c = \langle q, i, \omega \rangle \mid (q \in Q_\exists \wedge \text{Post}(c) \cap C_{k-1} \neq \emptyset) \vee (q \in Q_\forall \wedge \text{Post}(c) \subseteq C_{k-1}) \}$$

We call $C_\infty := \bigcup_{k=0}^\infty C_k$ the set of *accepting configurations* of \mathcal{A} . Intuitively, a configuration $c = \langle q, i, \omega \rangle$ is accepting if and only if $q \in Q_\forall$ and all successors of c are accepting or $q \in Q_\exists$ and at least one successor of c is accepting.

We say that \mathcal{A} accepts a word $w \in \Sigma^n$, if $\langle q_0, 1, \omega(w) \rangle$ is an accepting configuration, where $\omega(w)$ contains w on positions 1 through n and the blank \sqcup on all other positions. The **word acceptance problem** of ATM is, given an alternating Turing machine \mathcal{A} and a word $w \in \Sigma^n$, to decide whether \mathcal{A} accepts w . The word acceptance problem for ATMs that only use space polynomial in the length of the input word is EXPTIME-complete [CS76, CKS81].


2.3. Perfect Channel Systems. Given a finite set of messages M , define the set of channel operations $\text{Op} := \{!m, ?m \mid m \in M\} \cup \{\text{skip}\}$. A *perfect channel system* (PCS) is a triple $\mathcal{L} = \langle S, M, \delta \rangle$, where S is a finite set of states, M is the set of messages, and $\delta \subseteq S \times \text{Op} \times S$ is a transition relation. We write $s_1 \xrightarrow{\text{op}} s_2$ if $\langle s_1, \text{op}, s_2 \rangle \in \delta$.

Intuitively, a PCS models a finite state automaton that is augmented by a *perfect* (i.e. non-lossy) FIFO buffer, called *channel*. During a *send operation* $!m$, the channel system appends m to the tail of the channel. A transition $?m$ is called *receive operation*. It is only enabled if the channel is not empty and m is its oldest message. When the channel system performs this operation, it removes m from the head of the channel. Lastly, a *skip* operation just changes the state, but does not modify the buffer.

The formal semantics of \mathcal{L} are defined by a transition system $\mathcal{T}_\mathcal{L} = \langle C_\mathcal{L}, L_\mathcal{L}, \rightarrow_\mathcal{L} \rangle$, where $C_\mathcal{L} := S \times M^*$, $L_\mathcal{L} := \text{Op}$ and the transition relation $\rightarrow_\mathcal{L}$ is the smallest relation given by:

- If $s_1 \xrightarrow{!m} s_2$ and $w \in M^*$, then $\langle s_1, w \rangle \xrightarrow{!m}_\mathcal{L} \langle s_2, m \cdot w \rangle$.
- If $s_1 \xrightarrow{?m} s_2$ and $w \in M^*$, then $\langle s_1, w \cdot m \rangle \xrightarrow{?m}_\mathcal{L} \langle s_2, w \rangle$.
- If $s_1 \xrightarrow{\text{skip}} s_2$ and $w \in M^*$, then $\langle s_1, w \rangle \xrightarrow{\text{skip}}_\mathcal{L} \langle s_2, w \rangle$.

A state $s_F \in S$ is *reachable* from a configuration $c_0 \in C_\mathcal{L}$, if there exists a configuration $c_F = \langle s_F, w_F \rangle$ such that c_F is reachable from c_0 in $\mathcal{T}_\mathcal{L}$. The **state reachability problem** of PCS is, given a perfect channel system \mathcal{L} , an initial configuration $c_0 \in C_\mathcal{L}$ and a final state $s_F \in S$, to decide whether s_F is reachable from c_0 in $\mathcal{T}_\mathcal{L}$. It is undecidable [BZ83].

FIGURE 3. Concurrent program $\mathcal{P} = \langle \text{Proc}^1, \text{Proc}^2 \rangle$

3. CONCURRENT PROGRAMS

3.1. Syntax. Let Dom be a finite data domain and Vars be a finite set of shared variables over Dom . We define the *instruction set*

$$\begin{aligned} \text{Instrs} := & \{ \text{rd}(x, d), \text{wr}(x, d) \mid x \in \text{Vars}, d \in \text{Dom} \} \\ & \cup \{ \text{arw}(x, d, d') \mid x \in \text{Vars}, d, d' \in \text{Dom} \} \\ & \cup \{ \text{skip}, \text{mf} \} \end{aligned}$$

which are called *read*, *write*, *atomic read-write*, *skip* and *memory fence*, respectively. A process is represented by a finite state labelled transition system. It is given as the triple $\text{Proc} = \langle Q, \text{Instrs}, \delta \rangle$, where Q is a finite set of *local states* and $\delta \subseteq Q \times \text{Instrs} \times Q$ is the transition relation. As with transition systems, we write $q_1 \xrightarrow{\text{instr}} q_2$ if $\langle q_1, \text{instr}, q_2 \rangle \in \delta$ and $q_1 \rightarrow q_2$ if there exists some instr such that $q_1 \xrightarrow{\text{instr}} q_2$.

A *concurrent program* is a tuple of processes $\mathcal{P} = \langle \text{Proc}^\iota \rangle_{\iota \in \mathcal{I}}$, where \mathcal{I} is a finite set of process identifiers. For each $\iota \in \mathcal{I}$ we have $\text{Proc}^\iota = \langle Q^\iota, \text{Instrs}, \delta^\iota \rangle$. A *global state* of \mathcal{P} is a function $\mathcal{S} : \mathcal{I} \rightarrow \bigcup_{\iota \in \mathcal{I}} Q^\iota$ that maps each process to its local state, i.e $\mathcal{S}(\iota) \in Q^\iota$. Figure 3 shows a simple example of a concurrent program \mathcal{P} consisting of two processes Proc^1 and Proc^2 .

3.2. SC Semantics. Under SC (Sequential Consistency) semantics, all processes of a concurrent program interact with the shared memory directly, and all operations appear to be executed in some sequential order that is consistent with the program order of each individual process.

Formally, an SC *configuration* is a tuple $c = \langle \mathcal{S}, \mathcal{M} \rangle$, where $\mathcal{S} : \mathcal{I} \rightarrow \bigcup_{\iota \in \mathcal{I}} Q^\iota$ is a global state of \mathcal{P} and $\mathcal{M} : \text{Vars} \rightarrow \text{Dom}$ represents the memory state of each shared variable. Given a configuration c , we write $\mathcal{S}(c)$ and $\mathcal{M}(c)$ for the global program state and memory state of c .

The semantics of a concurrent program running under SC is defined by a transition system $\mathcal{T}_{\mathcal{P}}^{\text{SC}} = \langle \mathcal{C}_{\mathcal{P}}^{\text{SC}}, \mathcal{L}_{\mathcal{P}}^{\text{SC}}, \rightarrow_{\mathcal{P}}^{\text{SC}} \rangle$, where $\mathcal{C}_{\mathcal{P}}^{\text{SC}}$ is the set of all SC configurations of \mathcal{P} and $\mathcal{L}_{\mathcal{P}}^{\text{SC}} := \{ \text{instr}_\iota \mid \text{instr} \in \text{Instrs}, \iota \in \mathcal{I} \}$ is the set of labels. The transition relation \rightarrow (we usually drop the indices \mathcal{P}^{SC} whenever they are clear from context) is given by the following rules (see Figure 4):

$$\begin{array}{lcl}
\text{read} & \frac{q \xrightarrow{\text{rd}(x,d)} q' \quad \mathcal{S}(\iota)=q \quad \mathcal{M}(x)=d}{\langle \mathcal{S}, \mathcal{M} \rangle \xrightarrow{\text{rd}(x,d)_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{M} \rangle} \\
\text{write} & \frac{q \xrightarrow{\text{wr}(x,d)} q' \quad \mathcal{S}(\iota)=q}{\langle \mathcal{S}, \mathcal{M} \rangle \xrightarrow{\text{wr}(x,d)_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{M}[x \leftarrow d] \rangle} \\
\text{arw} & \frac{q \xrightarrow{\text{arw}(x,d,d')} q' \quad \mathcal{S}(\iota)=q \quad \mathcal{M}(x)=d}{\langle \mathcal{S}, \mathcal{M} \rangle \xrightarrow{\text{arw}(x,d,d')_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{M}[x \leftarrow d'] \rangle} \\
\text{skip} & \frac{q \xrightarrow{\text{skip}} q' \quad \mathcal{S}(\iota)=q}{\langle \mathcal{S}, \mathcal{M} \rangle \xrightarrow{\text{skip}_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{M} \rangle}
\end{array}$$

FIGURE 4. SC semantics

- **Read:** If $\mathcal{S}(\iota) = q$, $\mathcal{M}(x) = d$ and $q \xrightarrow{\text{rd}(x,d)} q'$, then $\langle \mathcal{S}, \mathcal{M} \rangle \xrightarrow{\text{rd}(x,d)_\iota} \langle \mathcal{S}', \mathcal{M} \rangle$, where $\mathcal{S}'(\iota) = q'$ and $\mathcal{S}'(\iota') = \mathcal{S}(\iota')$ for all $\iota' \neq \iota$.
- **Write:** If $\mathcal{S}(\iota) = q$ and $q \xrightarrow{\text{wr}(x,d)} q'$, then $\langle \mathcal{S}, \mathcal{M} \rangle \xrightarrow{\text{wr}(x,d)_\iota} \langle \mathcal{S}', \mathcal{M}' \rangle$, where $\mathcal{S}'(\iota) = q'$, $\mathcal{S}'(\iota') = \mathcal{S}(\iota')$ for all $\iota' \neq \iota$, $\mathcal{M}'(x) = d$ and $\mathcal{M}'(x') = \mathcal{M}(x')$ for all $x' \neq x$.
- **ARW:** If $\mathcal{S}(\iota) = q$, $\mathcal{M}(x) = d$ and $q \xrightarrow{\text{arw}(x,d,d')} q'$, then $\langle \mathcal{S}, \mathcal{M} \rangle \xrightarrow{\text{arw}(x,d,d')_\iota} \langle \mathcal{S}', \mathcal{M}' \rangle$, where $\mathcal{S}'(\iota) = q'$, $\mathcal{S}'(\iota') = \mathcal{S}(\iota')$ for all $\iota' \neq \iota$, $\mathcal{M}'(x) = d$ and $\mathcal{M}'(x') = \mathcal{M}(x')$ for all $x' \neq x$.
- **Skip:** If $\mathcal{S}(\iota) = q$ and $q \xrightarrow{\text{skip}} q'$, then $\langle \mathcal{S}, \mathcal{M} \rangle \xrightarrow{\text{skip}_\iota} \langle \mathcal{S}', \mathcal{M} \rangle$, where $\mathcal{S}'(\iota) = q'$ and $\mathcal{S}'(\iota') = \mathcal{S}(\iota')$ for all $\iota' \neq \iota$.

The memory fence operation **mf** is not used in programs running under SC semantics. If needed, it can be formally defined as following the same semantics as **skip**.

A global state \mathcal{S}_F is *reachable* from an initial configuration c_0 , if there is a configuration c_F with $\mathcal{S}(c_F) = \mathcal{S}_F$ such that c_F is reachable from c_0 in $\mathcal{T}_P^{\text{SC}}$. The **state reachability problem** of SC is, given a program \mathcal{P} , an initial configuration c_0 and a final global state \mathcal{S}_F , to decide whether \mathcal{S}_F is reachable from c_0 in $\mathcal{T}_P^{\text{SC}}$.

3.3. TSO Semantics. Under TSO semantics, the processes of a concurrent program do not interact with the shared memory directly, but indirectly through FIFO *store buffers* instead. When performing a *write* instruction $\text{wr}(x, d)$, the process adds a new message $\langle x, d \rangle$ to the tail of its own store buffer. A *read* instruction $\text{rd}(x, d)$ works differently depending on the current buffer content of the process. If its buffer contains a write message on variable x , the value d must correspond to the value of the most recent such message. Otherwise, the value is read directly from memory. The *atomic read-write* instruction $\text{arw}(x, d, d')$ is only enabled if the buffer of the process is empty and the value of x in the memory is d . When executed, this instruction directly changes the value of x in the memory to d' . A *skip* instruction only changes the local state of the process. The *memory fence* instruction is disabled, i.e. it cannot be executed, unless the buffer of the process is empty. Additionally, at any point during the execution, the process can *update* the write message at the head of its buffer to the memory. For example, if the oldest message in the buffer is $\langle x, d \rangle$, it will be removed from the buffer and the memory value of variable x will be updated to contain the value d . This happens in a non-deterministic manner.

$$\begin{array}{lcl}
\text{read-own-write} & \frac{q \xrightarrow{\text{rd}(x,d)} q' \quad \mathcal{S}(\iota)=q \quad \mathcal{B}(\iota)|_{\{x\} \times \text{Dom}} = \langle x, d \rangle \cdot w}{\langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle \xrightarrow{\text{rd}(x,d)_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{B}, \mathcal{M} \rangle} \\
\text{read-from-memory} & \frac{q \xrightarrow{\text{rd}(x,d)} q' \quad \mathcal{S}(\iota)=q \quad \mathcal{B}(\iota)|_{\{x\} \times \text{Dom}} = \varepsilon \quad \mathcal{M}(x)=d}{\langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle \xrightarrow{\text{rd}(x,d)_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{B}, \mathcal{M} \rangle} \\
\text{write} & \frac{q \xrightarrow{\text{wr}(x,d)} q' \quad \mathcal{S}(\iota)=q}{\langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle \xrightarrow{\text{wr}(x,d)_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{B}[\iota \leftarrow \langle x, d \rangle \cdot \mathcal{B}(\iota)], \mathcal{M} \rangle} \\
\text{arw} & \frac{q \xrightarrow{\text{arw}(x,d,d')} q' \quad \mathcal{S}(\iota)=q \quad \mathcal{B}(\iota)=\varepsilon \quad \mathcal{M}(x)=d}{\langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle \xrightarrow{\text{arw}(x,d,d')_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{B}, \mathcal{M}[x \leftarrow d'] \rangle} \\
\text{skip} & \frac{q \xrightarrow{\text{skip}} q' \quad \mathcal{S}(\iota)=q}{\langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle \xrightarrow{\text{skip}_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{B}, \mathcal{M} \rangle} \\
\text{memory-fence} & \frac{q \xrightarrow{\text{mf}} q' \quad \mathcal{S}(\iota)=q \quad \mathcal{B}(\iota)=\varepsilon}{\langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle \xrightarrow{\text{mf}_\iota} \langle \mathcal{S}[\iota \leftarrow q'], \mathcal{B}, \mathcal{M} \rangle} \\
\text{update} & \frac{\mathcal{B}(\iota)=w \cdot \langle x, d \rangle}{\langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle \xrightarrow{\text{up}_\iota} \langle \mathcal{S}, \mathcal{B}[\iota \leftarrow w], \mathcal{M}[x \leftarrow d] \rangle}
\end{array}$$

FIGURE 5. TSO semantics

Formally, we introduce a TSO *configuration* as a tuple $c = \langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle$, where:

- $\mathcal{S} : \mathcal{I} \rightarrow \bigcup_{\iota \in \mathcal{I}} \mathcal{Q}^\iota$ is a global state of \mathcal{P} .
- $\mathcal{B} : \mathcal{I} \rightarrow (\text{Vars} \times \text{Dom})^*$ represents the buffer state of each process.
- $\mathcal{M} : \text{Vars} \rightarrow \text{Dom}$ represents the memory state of each shared variable.

Given a configuration c , we write $\mathcal{S}(c)$, $\mathcal{B}(c)$ and $\mathcal{M}(c)$ for the global program state, buffer state and memory state of c . The semantics of a concurrent program running under TSO is defined by a transition system $\mathcal{T}_P^{\text{TSO}} = \langle \mathcal{C}_P^{\text{TSO}}, \mathcal{L}_P^{\text{TSO}}, \rightarrow_P^{\text{TSO}} \rangle$, where $\mathcal{C}_P^{\text{TSO}}$ is the set of all possible TSO configurations and $\mathcal{L}_P^{\text{TSO}} := \{\text{instr}_\iota \mid \text{instr} \in \text{Instrs}, \iota \in \mathcal{I}\} \cup \{\text{up}_\iota \mid \iota \in \mathcal{I}\}$ is the set of labels. The transition relation \rightarrow (i.e. $\rightarrow_P^{\text{TSO}}$) is given by the rules in Figure 5, where we use $\mathcal{B}(\iota)|_{\{x\} \times \text{Dom}}$ to denote the restriction of $\mathcal{B}(\iota)$ to write messages on the variable x .

A global state \mathcal{S}_F is *reachable* from an initial configuration c_0 , if there is a configuration c_F with $\mathcal{S}(c_F) = \mathcal{S}_F$ such that c_F is reachable from c_0 in \mathcal{T}_P . The **state reachability problem** of TSO is, given a program \mathcal{P} , an initial configuration c_0 and a final global state \mathcal{S}_F , to decide whether \mathcal{S}_F is reachable from c_0 in \mathcal{T}_P .

We define up^* to be the transitive closure of $\{\text{up}_\iota \mid \iota \in \mathcal{I}\}$, i.e. $c_1 \xrightarrow{\text{up}^*}_P c_2$ if and only if c_2 can be obtained from c_1 by some amount of buffer updates.

4. GAMES

4.1. Definitions. A (*safety*) *game* is an unlabelled transition system, in which two players A and B take turns making a *move* in the transition system, i.e. changing the state of the game from one configuration to an adjacent one. The set of configurations is partitioned

into two sets C_A and C_B representing the configurations of player A and B, respectively. Whenever the game is in a configuration that belongs to player A, it is her turn to decide which move to make. This continues until the game reaches a configuration that belongs to player B, at which point the control is passed to her instead. The goal of player B is to reach a given set of final configurations, while player A tries to avoid this. Thus, it can also be seen as a *reachability* game with respect to player B.

Formally, a game is defined as a tuple $\mathcal{G} = \langle C, C_A, C_B, \rightarrow, C_F \rangle$, where C is the set of configurations, C_A and C_B form a partition of C , the transition relation is $\rightarrow \subseteq C \times C$, and $C_F \subseteq C_A$ is a set of *final configurations*. Furthermore, we require that \mathcal{G} is deadlock-free, i.e. $\text{Post}(c) \neq \emptyset$ for all $c \in C$.

A *play* P of \mathcal{G} is an infinite sequence c_0, c_1, \dots such that $c_i \rightarrow c_{i+1}$ for all $i \in \mathbb{N}$. In the context of safety games, P is *winning* for player B if there is $i \in \mathbb{N}$ such that $c_i \in C_F$. Otherwise, it is *winning* for player A. This means that player B tries to force the play into C_F , while player A tries to avoid this.

A *strategy* of player A is a partial function $\sigma_A : C^* \rightarrow C_B$, such that $\sigma_A(c_0, \dots, c_n)$ is defined if and only if c_0, \dots, c_n is a prefix of a play, $c_n \in C_A$ and $\sigma_A(c_0, \dots, c_n) \in \text{Post}(c_n)$. A strategy σ_A is called *positional*, if it only depends on c_n , i.e. if $\sigma_A(c_0, \dots, c_n) = \sigma_A(c_n)$ for all (c_0, \dots, c_n) on which σ_A is defined. Thus, a positional strategy is usually given as a total function $\sigma_A : C_A \rightarrow C$. Given two games \mathcal{G} and \mathcal{G}' and a strategy σ_A for \mathcal{G} , an *extension* of σ_A to \mathcal{G}' is a strategy σ'_A of \mathcal{G}' that is also an extension of σ_A to the configuration set of \mathcal{G}' in the mathematical sense, i.e. $\sigma_A(c_0, \dots, c_n) = \sigma'_A(c_0, \dots, c_n)$ for all (c_0, \dots, c_n) on which σ_A is defined. Conversely, σ_A is called the *restriction* of σ'_A to \mathcal{G} . For player B, strategies are defined accordingly.

Two strategies σ_A and σ_B together with an initial configuration c_0 induce a play $P(c_0, \sigma_A, \sigma_B) = c_0, c_1, \dots$ such that $c_{i+1} = \sigma_A(c_0, \dots, c_i)$ for all $c_i \in C_A$ and $c_{i+1} = \sigma_B(c_0, \dots, c_i)$ for all $c_i \in C_B$. A strategy σ_A is *winning* from a configuration c , if for *all* strategies σ_B it holds that $P(\sigma_A, \sigma_B, c)$ is a winning play for player A. A configuration c is *winning* for player A if she has a strategy that is winning from c . Equivalent notions exist for player B. The **safety problem** for a game \mathcal{G} and a configuration c is to decide whether c is winning for player A.

Lemma 4.1 (Proposition 2.21 in [Maz01]). *In safety games, every configuration is winning for exactly one player. A player with a winning strategy also has a positional winning strategy.*

Since we only consider safety games in this paper, strategies will be considered to be positional unless explicitly stated otherwise. Furthermore, Lemma 4.1 implies the following:

- $c_A \in C_A$ is winning for player A \iff there is $c \in \text{Post}(c_A)$ that is winning for player A.
- $c_B \in C_B$ is winning for player A \iff all $c \in \text{Post}(c_B)$ are winning for player A.

A *finite game* is a game with a finite set of configurations. It is rather intuitive that the safety problem is decidable for finite games, e.g. by applying a backward induction algorithm. In particular, the winning configurations for each player are computable in linear time:

Lemma 4.2 (Chapter 2 in [GTW02]). *Computing the set of winning configurations for a finite game with n configurations and m transitions is in $\mathcal{O}(n + m)$.*




FIGURE 6. Configurations in a bisimulation




FIGURE 7. Configurations and strategies in Lemma 4.3

Bisimulations. A *bisimulation* (also called *zig-zag relation*) between two games $\mathcal{G} = \langle C^{\mathcal{G}}, C_A^{\mathcal{G}}, C_B^{\mathcal{G}}, \rightarrow, C_F^{\mathcal{G}} \rangle$ and $\mathcal{H} = \langle C^{\mathcal{H}}, C_A^{\mathcal{H}}, C_B^{\mathcal{H}}, \rightarrow, C_F^{\mathcal{H}} \rangle$ is a relation $Z \subseteq C^{\mathcal{G}} \times C^{\mathcal{H}}$ such that for all pairs of related configurations $(c_1, c_2) \in Z$ it holds that (cf. Figure 6):

- (*zig*) for each transition $c_1 \rightarrow c_3$ there is a transition $c_2 \rightarrow c_4$ such that $(c_3, c_4) \in Z$.
- (*zag*) for each transition $c_2 \rightarrow c_4$ there is a transition $c_1 \rightarrow c_3$ such that $(c_3, c_4) \in Z$.
- c_1 and c_2 are owned by the same player: $c_1 \in C_A^{\mathcal{G}} \iff c_2 \in C_A^{\mathcal{H}}$
- c_1 and c_2 agree on being a final configuration: $c_1 \in C_F^{\mathcal{G}} \iff c_2 \in C_F^{\mathcal{H}}$

We say that two related configurations c_1 and c_2 are *bisimilar* and write $c_1 \approx c_2$. We call \mathcal{G} and \mathcal{H} *bisimilar* if there is a bisimulation between them. It is common knowledge in game theory that winning strategies are preserved under bisimulations:

Lemma 4.3. *Given two bisimilar configurations $c_0^{\mathcal{G}} \in C^{\mathcal{G}}$ and $c_0^{\mathcal{H}} \in C^{\mathcal{H}}$, it holds that $c_0^{\mathcal{G}}$ is winning for player A if and only if $c_0^{\mathcal{H}}$ is winning for player A.*

Proof. Suppose that $c_0^{\mathcal{G}}$ is winning for player A with (positional) strategy $\sigma_A^{\mathcal{G}}$ and consider the case $c_0^{\mathcal{G}} \in C_A^{\mathcal{G}}$. Let $\sigma_B^{\mathcal{H}}$ be an arbitrary strategy for player A in \mathcal{H} . In the following, we will describe two (non-positional) winning strategies $\sigma_A^{\mathcal{H}}$ and $\sigma_B^{\mathcal{G}}$.

For $n \in \mathbb{N}$, define recursively (see Figure 7):

- $c_{2n+1}^{\mathcal{G}} := \sigma_A^{\mathcal{G}}(c_{2n}^{\mathcal{G}})$
- $c_{2n+1}^{\mathcal{H}}$ such that $c_{2n}^{\mathcal{H}} \rightarrow c_{2n+1}^{\mathcal{H}}$ and $c_{2n+1}^{\mathcal{G}} \approx c_{2n+1}^{\mathcal{H}}$, which exists by the zig property of Z .
- $\sigma_A^{\mathcal{H}}(c_0^{\mathcal{H}}, \dots, c_{2n}^{\mathcal{H}}) := c_{2n+1}^{\mathcal{H}}$
- $c_{2n+2}^{\mathcal{H}} := \sigma_B^{\mathcal{H}}(c_{2n+1}^{\mathcal{H}})$
- $c_{2n+2}^{\mathcal{G}}$ such that $c_{2n+1}^{\mathcal{G}} \rightarrow c_{2n+2}^{\mathcal{G}}$ and $c_{2n+2}^{\mathcal{G}} \approx c_{2n+2}^{\mathcal{H}}$, which exists by the zag property of Z .
- $\sigma_B^{\mathcal{G}}(c_0^{\mathcal{G}}, \dots, c_{2n+1}^{\mathcal{G}}) := c_{2n+2}^{\mathcal{G}}$

Since $\sigma_A^{\mathcal{G}}$ is a winning strategy for player A in \mathcal{G} , the sequence $c_0^{\mathcal{G}}, c_1^{\mathcal{G}}, \dots$ is a winning play for player A, i.e. it does not visit the set of final configurations $C_F^{\mathcal{G}}$. Thus, $c_0^{\mathcal{H}}, c_1^{\mathcal{H}}, \dots$

is also a winning play for player A, because for all $n \in \mathbb{N}$, $c_n^G \approx c_n^H$ and $c_n^G \notin C_F^G$ implies $c_n^H \notin C_F^H$. Since σ_B^H was chosen arbitrarily, this shows that σ_A^H is a winning strategy for player A in \mathcal{H} . Note that by Lemma 4.1, player A could also choose a positional strategy instead.

The case where $c_0^G \in C_B^G$ is similar, but with the recursive definition shifted by one, i.e. $c_{2n+1}^G := \sigma_B^G(c_{2n}^G)$ and so on. \square

4.2. SC Games. We want to define an *SC game* as a safety game where the underlying transition system follows Sequential Consistency semantics. While the set of game configurations, the transition relation of the game and the set of its final configurations could be taken directly from the transition system \mathcal{T}_P^{SC} , it is not straightforward to decide how the configuration set should be partitioned into C_A and C_B . One idea that comes to mind would be to partition the *local states* of each process. But whose turn is it if Proc^1 is in a state belonging to player A but Proc^2 is in a state belonging to player B? Which player is allowed to move first? The standard approach to solve this problem is to consider the product of the concurrent processes instead, i.e. to partition the set of *global states*. Another common approach is to switch to turn-based games, where the players alternate in making a move in the game.

For this present work, we chose the latter approach, which means that the players take turns executing exactly one instruction from an arbitrary process. This is because we can argue that turn-based games do not pose a significant restriction compared to games with state ownership, but instead can emulate most of their behaviour (see below).

In the following formal definition of an SC game, we will create two copies of each SC configuration, annotated with *A* and *B*, respectively. The transition relation of the game will be restricted to transitions between configurations of different players. The use of two identical sets of configurations, one for each player, may seem unnecessary at first glance. However, this distinction is crucial for the definition of TSO games in the next section, where the available instructions to execute differ between the two players A and B. Furthermore, it allows us to describe turn-based games using the framework of games with configuration ownership, as introduced in the beginning of this section.

A program $\mathcal{P} = \langle \text{Proc}^\iota \rangle_{\iota \in \mathcal{I}}$ and a set of final local states $Q_F^P \subseteq Q^P$ induce a safety game $\mathcal{G}^{SC}(\mathcal{P}, Q_F^P) = \langle C, C_A, C_B, \rightarrow, C_F \rangle$ as follows. The sets C_A and C_B are copies of the set C^P of TSO configurations, annotated by *A* and *B*, respectively:

$$C_A := \{c_A \mid c \in C_P^{SC}\} \quad \text{and} \quad C_B := \{c_B \mid c \in C_P^{SC}\}$$

The set of final configurations is defined as:

$$C_F := \{ \langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle_A \in C_A \mid \exists \iota \in \mathcal{I} : \mathcal{S}(\iota) \in Q_F^P \}$$

That is, it is the set of all configurations where at least one process is in a final state. As described previously, the transition relation is defined by the program's instructions, in a way that the two players take turns alternatingly to execute these instructions: For each transition $c \xrightarrow{\text{instr}_\iota} c'$ where $c, c' \in C_P^{SC}$, $\iota \in \mathcal{I}$ and $\text{instr} \in \text{Instrs}$, it holds that $c_A \rightarrow c'_B$ and $c_B \rightarrow c'_A$. Figure 8 shows the transitions relation of the SC game induced by the program from Figure 3.

Note that in the analysis of SC games (and later of TSO games), we often talk about the program instruction that gives rise to a transition of \mathcal{G}^{SC} instead of talking about the

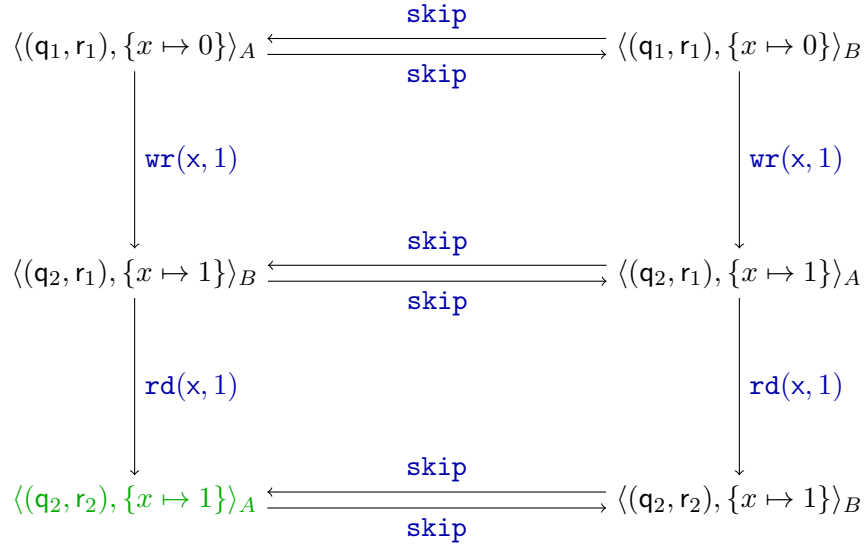

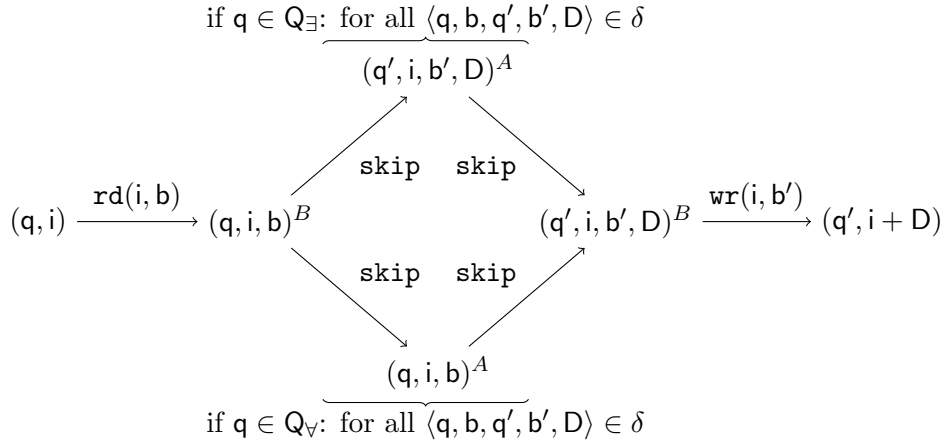


FIGURE 8. The transition relation of the SC game $\mathcal{G}^{\text{SC}}(\mathcal{P}, \{r_2\})$ induced by the program \mathcal{P} from Figure 3. Note that only configurations reachable from $\langle (q_1, r_1), \{x \mapsto 0\} \rangle_A$ are shown. The labels in blue are not formally part of the game definition, but are included to indicate which instruction of \mathcal{P} gives rise to the transition. The configuration in green is the final state induced by the set of final local states $Q_F^{\mathcal{P}} := \{r_2\}$.

transition explicitly. For example, in a situation where the game is in some configuration c_A with $\mathcal{S}(c)(\iota) = q$, we might say that player A executes instruction $q \xrightarrow{\text{instr}} q'$ and mean that player A moves from c_A to c'_B , where c' is the unique configuration obtained from executing instr at c . Similarly, for better readability we might drop the index A or B from a game configuration if it is clear from the context which player's turn it is.

State ownership. As mentioned above, turn-based games can emulate games with state ownership. The following example illustrates this informal idea. Consider an instruction $q_1 \xrightarrow{\text{instr}} q_2$ in some process, where q_1 is supposed to belong to player B. We modify the process by adding some new states and transitions, as shown in Figure 9. Note that through a suitable definition of the set of final configurations C_F we can ensure that reaching q_F means that player B wins the game. Suppose that the game is in a configuration where the process is in q_1 , it is player B's turn and she wants to execute the instruction instr . She can do so by taking the transition from q_1 to q'_2 . Afterwards, it is the turn of player A. She is forced to respond by taking the skip instruction from q'_2 to q_2 . Otherwise, player B could move to q_F in her turn and win immediately. Now consider the same situation as before but it is the turn of player A instead. If she would execute the instruction instr and move to q'_2 , then player B could respond with moving to q_F and again win immediately. Therefore, player A is prevented from executing any instruction starting in q_1 .

Complexity. In the following, we consider the safety problem for SC games. Unexpectedly, this problem is EXPTIME-hard even for games played on a single process. We prove this

FIGURE 9. The gadget for a transition $q_1 \xrightarrow{\text{instr}} q_2$ FIGURE 10. Construction of a program \mathcal{P} that simulates an ATM \mathcal{A} . The states and transitions shown are added to Proc for all $q \in Q$ and $-p(n) \leq i \leq p(n)$. We set $L := -1$ and $R := 1$ in the expression $i + D$.

by reducing the *word acceptance problem* of polynomial-space bounded *alternating Turing machines* (ATM) to the safety problem of a single-process SC game.

Consider an ATM $\mathcal{A} = \langle \Sigma, Q, q_0, q_F, Q_{\exists}, Q_{\forall}, \delta \rangle$ that is space-bounded by some polynomial $p(n)$. Given a word $w \in \Sigma^n$, we construct a concurrent program \mathcal{P} with a single process that simulates \mathcal{A} . The key idea is to store the state and head position of the ATM in the local states of the process, and use a set of variables to save the word on the working tape. Based on the alternations of the Turing machine, either player A or player B decides which transition the program will simulate. More precisely, we let player B simulate \mathcal{A} at the existential states while player A has control over the simulation at the universal states. Player B will win the SC game induced by \mathcal{P} if and only if \mathcal{A} accepts w .

$\mathcal{P} = \langle \text{Proc} \rangle$ uses the variables $\text{Vars} := \{i \mid -p(n) \leq i \leq p(n)\} \subset \mathbb{Z}$ over the domain $\text{Dom} := \Sigma$ to store the content of the tape. The control state of \mathcal{A} , the head position i and all necessary intermediate values are stored in the local states of Proc . The construction of Proc is shown in Figure 10. Note that e.g. $(q', i, b', D)^A$ is a *state* of Proc and not a *configuration* of \mathcal{P} or even \mathcal{G}^{SC} .

Define the set of local final states $Q_F^{\mathcal{P}} := \{(q_F, i) \mid -p(n) \leq i \leq p(n)\}$ and the initial memory state \mathcal{M}_0 , where $\mathcal{M}_0(i) = w(i)$ for $1 \leq i \leq n$ and $\mathcal{M}_0(i) = \perp$ otherwise.

Theorem 4.4. \mathcal{A} accepts $w \in \Sigma^n$ if and only if player B wins the game $\mathcal{G}^{\text{SC}}(\mathcal{P}, Q_F^{\mathcal{P}})$ from the initial configuration $\langle (q_0, 1), \mathcal{M}_0 \rangle_A$.

Proof. Recall the definitions of C_∞ and C_k from subsection 2.2 and suppose that \mathcal{A} accepts w . We argue that player B can force the simulation into q_F and describe a strategy for her to do so. The main idea is that if the simulation is in a configuration $c \in C_k$ ($k > 0$), then player B can force the simulation into another configuration $c' \in C_{k-1}$.

Let the game be in configuration $\langle (q, i), \omega \rangle_A$ which simulates the \mathcal{A} -configuration $\langle q, i, \omega \rangle$. Note that (q, i) is the local state of **Proc** and $\omega : \text{Vars} \rightarrow \Sigma$ is its memory state. (Formally, $\omega : \mathbb{Z} \rightarrow \Sigma$, but since $\text{Vars} \subset \mathbb{Z}$, this is only a slight abuse of notation. Since $\omega(i) = \perp$ for all $i \notin \text{Vars}$, $\omega : \text{Vars} \rightarrow \Sigma$ captures the same information as $\omega : \mathbb{Z} \rightarrow \Sigma$.) Let $b = \omega(i)$ be the symbol under the head of the Turing machine and $\langle q, i, \omega \rangle \in C_k$ for some $k > 0$. Player A has to take the only enabled transition, which is $\text{rd}(i, b)$ leading to the local state $(q, i, b)^B$. Now, the possible choices for player B depend on whether q is an existential or universal state.

If $q \in Q_\exists$, then player B chooses an \mathcal{A} -transition $\langle q, b, q', b', D \rangle \in \delta$ with $\langle q', i + D, \omega[i \leftarrow b'] \rangle \in C_{k-1}$, which exists by definition of C_k . Player B then moves to the local state $(q', i, b', D)^A$ and player A has to respond with $(q', i, b', D)^B$.

Otherwise, if $q \in Q_\forall$, then player B must take the transition to the state $(q, i, b)^A$. From there, player A moves to some state $(q', i, b', D)^B$. This implies that there is an \mathcal{A} -transition $\langle q, b, q', b', D \rangle \in \delta$. From the construction of C_k , it follows that $\langle q', i + D, \omega[i \leftarrow b'] \rangle \in C_{k-1}$.

In both cases, the game is now in some configuration $\langle (q', i, b', D)^B, \omega \rangle_B$ with $\langle q', i + D, \omega[i \leftarrow b'] \rangle_A \in C_{k-1}$. Player B takes the transition $\text{wr}(i, b')$ and the game is in configuration $\langle (q', i + D), \omega[i \leftarrow b'] \rangle_A$. It follows by simple induction over k that player B can force the game into simulating an \mathcal{A} -configuration of C_0 .

For the other direction, assume that \mathcal{A} does not accept w . This time we describe a strategy for player A to win the game. Similar to the previous argumentation, we can show by induction that player A can simulate a sequence of configurations that are *not* accepting. The difference is that if $q \in Q_\exists$, then player B can never choose any transition that leads to an accepting configuration, while in the case of $q \in Q_\forall$, player A can always avoid accepting configurations. This follows from the fact that the simulated \mathcal{A} -configurations $\langle q, i, \omega \rangle$ are not elements of any of the C_k . \square

Theorem 4.5. *The safety problem for SC games is EXPTIME-complete.*

Proof. The word acceptance problem of a linearly bounded ATM \mathcal{A} is EXPTIME-hard [CS76, CKS81]. In Theorem 4.4, it is reduced to the safety problem for the SC game induced by a concurrent program \mathcal{P} . This program has $\mathcal{O}(|Q| \cdot |\Sigma| \cdot p(n))$ local states, which is polynomial in the size of \mathcal{A} . It follows that the safety problem for TSO games is EXPTIME-hard. Membership follows directly from Lemma 4.2 and the fact that for any \mathcal{P} , the game $\mathcal{G}^{\text{SC}}(\mathcal{P}, Q_F^{\mathcal{P}})$ has $\mathcal{O}(\prod_{\iota \in \mathcal{I}} |Q^\iota| \cdot |\text{Dom}|^{|\text{Vars}|})$ many configurations. \square

4.3. TSO Games. Given a $\mathcal{P} = \langle \text{Proc}^\iota \rangle_{\iota \in \mathcal{I}}$ and a set of final local states $Q_F^{\mathcal{P}} \subseteq Q^{\mathcal{P}}$, the TSO game $\mathcal{G}^{\text{TSO}}(\mathcal{P}, Q_F^{\mathcal{P}})$ is defined similar to the SC game $\mathcal{G}^{\text{SC}}(\mathcal{P}, Q_F^{\mathcal{P}})$. The only addition is that the game needs to be able to handle buffer updates. Therefore, we allow one or both of the players to perform buffer updates, either before or after their turn. Which player is allowed to do so and when exactly she is allowed to do so depends on the concrete instantiation of the TSO game. The complete transition rules of the TSO game are:

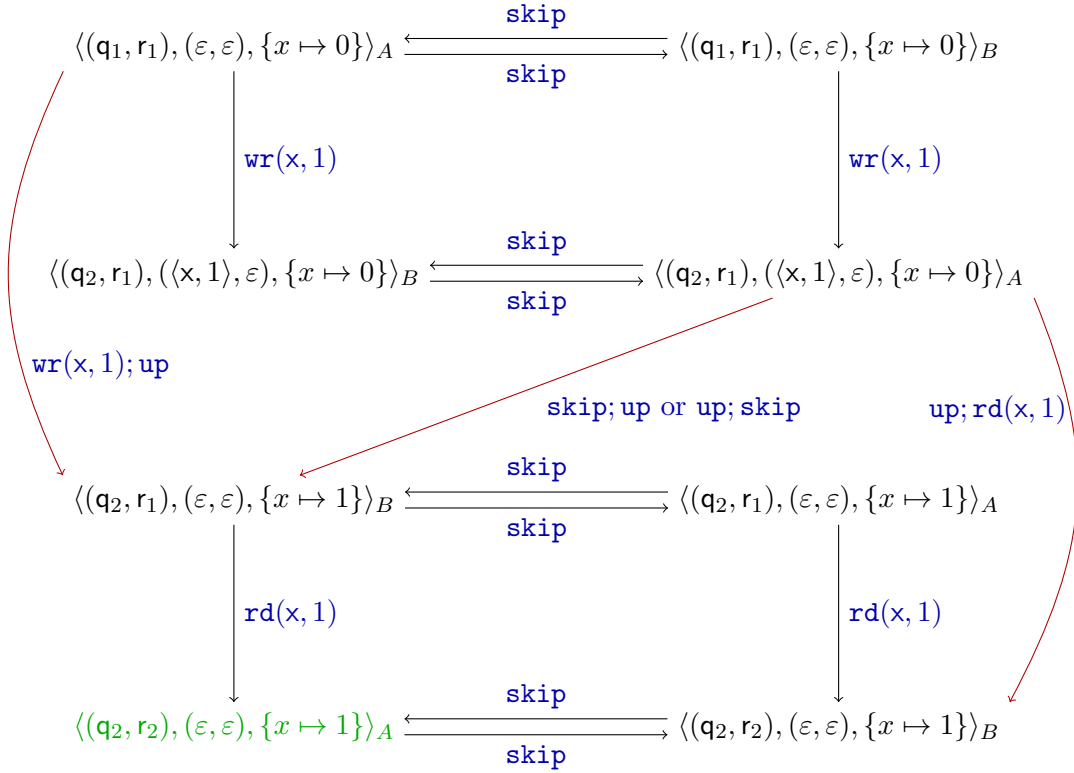


FIGURE 11. The transition relation of the TSO game $\mathcal{G}^{\text{TSO}}(\mathcal{P}, \{r_2\})$ induced by the program \mathcal{P} from Figure 3, in the case where player A is allowed to update before and after her turn, but player B is not allowed to update buffer messages. Note that only configurations reachable from $\langle (q_1, r_1), (\epsilon, \epsilon), \{x \mapsto 0\} \rangle_A$ are shown. The labels in blue are not formally part of the game definition, but are included to indicate which instruction of \mathcal{P} gives rise to the transition. The transitions in red are due to updates; they correspond to both an instruction and an update operation. The configuration in green is the final state induced by the set of final local states $Q_F^{\mathcal{P}} := \{r_2\}$.

- For each transition $c \xrightarrow{\text{instr}_\iota} c'$ where $c, c' \in \mathcal{C}^{\mathcal{P}}$, $\iota \in \mathcal{I}$ and $\text{instr} \in \text{Instrs}$, it holds that $c_A \rightarrow c'_A$ and $c_B \rightarrow c'_B$. This is the same as for SC and means that each player can execute any TSO instruction, but they take turns alternatingly.
- *If player A can update before her own turn:* For each transition $c_A \rightarrow c'_A$ introduced by any of the previous rules, it holds that $\tilde{c}_A \rightarrow c'_A$ for all \tilde{c} with $\tilde{c} \xrightarrow{\text{up}^*} c$.
- *If player A can update after her own turn:* For each transition $c_A \rightarrow c'_A$ introduced by any of the previous rules, it holds that $c_A \rightarrow \tilde{c}'_A$ for all \tilde{c}' with $c' \xrightarrow{\text{up}^*} \tilde{c}'$.
- The update rules for player B are defined in a similar manner.

From this definition, we obtain 16 different variants of TSO games, which differ in whether each of the players can update *never*, *before* her turn, *after* her turn, or *always* (before and after her turn). Figure 11 shows the transition relation of the TSO game induced

by the program from Figure 3, in the case where player A is allowed to update buffer messages both before and after her turn, but player B is not allowed to do so.

We group games with similar decidability and complexity results together. An overview of these four groups is presented in Figure 2. Each group is described in detail in the following sections.

But first, we use the complexity results for SC games to obtain a lower bound for all groups of TSO games. Interestingly, we can argue that the exact type of TSO game is irrelevant. Note that the program \mathcal{P} from Theorem 4.4 uses only one process and no memory fences or atomic read-writes. Consider the view the single process **Proc** has on a variable of the concurrent system. Independently of the buffer updates, the process will always read the last value that it has written: Either the last write is still in the buffer and will be read from there, or the process reads directly from the memory, which must contain the last value written by the process. Thus, any sequence of instructions executed in the SC game is also enabled in the TSO game. This is formalised through a bimsimulation between the games $\mathcal{G}^{\text{TSO}}(\mathcal{P}, \mathcal{Q}_F^{\mathcal{P}}) = \langle \mathcal{C}^{\text{TSO}}, \mathcal{C}_A^{\text{TSO}}, \mathcal{C}_B^{\text{TSO}}, \rightarrow, \mathcal{C}_F^{\text{TSO}} \rangle$ and $\mathcal{G}^{\text{SC}}(\mathcal{P}, \mathcal{Q}_F^{\mathcal{P}}) = \langle \mathcal{C}^{\text{SC}}, \mathcal{C}_A^{\text{SC}}, \mathcal{C}_B^{\text{SC}}, \rightarrow, \mathcal{C}_F^{\text{SC}} \rangle$.

For a TSO configuration c , let \bar{c} be the configuration obtained from c by updating all buffer messages to the memory, i.e. $c \xrightarrow{\text{up}^*} \bar{c}$ and $\mathcal{B}(\bar{c}) = \langle \varepsilon \rangle$. Note that \bar{c} is unique since \mathcal{P} has only one process and thus there is only one way to update all messages. We extend this notation to game configurations in the obvious way.

Lemma 4.6. *The relation*

$$\mathbf{Z} := \{(c, \langle \mathcal{S}(c), \mathcal{M}(\bar{c}) \rangle) \mid c \in \mathcal{C}^{\text{TSO}}\} \subset (\mathcal{C}^{\text{TSO}} \times \mathcal{C}^{\text{SC}})$$

is a bisimulation between $\mathcal{G}^{\text{TSO}}(\mathcal{P}, \mathcal{Q}_F^{\mathcal{P}})$ and $\mathcal{G}^{\text{SC}}(\mathcal{P}, \mathcal{Q}_F^{\mathcal{P}})$.

Proof. We need to show that for all $(c_1, c_2) \in \mathbf{Z}$:

- For all $c_1 \rightarrow c_3$ in \mathcal{G}^{TSO} , there is $c_2 \rightarrow c_4$ in \mathcal{G}^{SC} with $c_3 \approx c_4$.
- For all $c_2 \rightarrow c_4$ in \mathcal{G}^{SC} , there is $c_1 \rightarrow c_3$ in \mathcal{G}^{TSO} with $c_3 \approx c_4$.
- $c_1 \in \mathcal{C}_A^{\text{TSO}}$ if and only if $c_2 \in \mathcal{C}_A^{\text{SC}}$.
- $c_1 \in \mathcal{C}_F^{\text{TSO}}$ if and only if $c_2 \in \mathcal{C}_F^{\text{SC}}$.

For the first property, consider a transition $c_1 \rightarrow c_3$ in \mathcal{G}^{TSO} . It is due to some instruction $\mathcal{S}(c_1) \xrightarrow{\text{instr}} \mathcal{S}(c_3)$ in **Proc** of \mathcal{P} . If $\text{instr} = \text{rd}(x, d)$ for some variable x and value d , then d must be the value of the last x -message in the buffer of c_1 , or the buffer does not contain such a message and d is the value of x in the memory. In both cases, $\mathcal{M}(\bar{c}_1)(x) = d$. Thus, $\text{rd}(x, d)$ is enabled at c_2 , since $\mathcal{S}(c_1) = \mathcal{S}(c_2)$ and $\mathcal{M}(\bar{c}_1) = \mathcal{M}(c_2)$. Otherwise, if instr is an instruction other than rd , it is always enabled at $\mathcal{S}(c_2)$. Note that $\text{instr} \neq \text{mf}$ and $\text{instr} \neq \text{arw}(x, d, d')$ since \mathcal{P} does not use memory fences.

Let c_4 be the unique SC configuration obtained from c_2 after executing the program instruction $\mathcal{S}(c_1) \xrightarrow{\text{instr}} \mathcal{S}(c_3)$. If $\text{instr} = \text{wr}(x, d)$ for some variable x and value d , then $\mathcal{M}(\bar{c}_3) = \mathcal{M}(\bar{c}_1)[x \leftarrow d]$ and $\mathcal{M}(c_4) = \mathcal{M}(c_2)[x \leftarrow d]$. This holds even if $c_1 \rightarrow c_3$ includes buffer message updates of any kind. Otherwise, if instr is an instruction other than wr , it simply holds that $\mathcal{M}(\bar{c}_3) = \mathcal{M}(\bar{c}_1)$ and $\mathcal{M}(c_4) = \mathcal{M}(c_2)$. Since $\mathcal{M}(\bar{c}_1) = \mathcal{M}(c_2)$, we have $\mathcal{M}(\bar{c}_3) = \mathcal{M}(c_4)$. Using $\mathcal{S}(c_3) = \mathcal{S}(c_4)$ we conclude that $c_3 \approx c_4$.

The second property is proven analogously. The third and fourth properties are trivially fulfilled by the definition of \mathbf{Z} . \square




FIGURE 12. Commutative diagram of strategies in games of group I.

Corollary 4.7. *The reachability problem for TSO games is EXPTIME-hard.*

Proof. This follows directly from Lemma 4.3, Theorem 4.4, Lemma 4.6 and the fact that the word acceptance problem of linearly bounded ATMs is EXPTIME-hard. Note that Lemma 4.3 is applicable since $c \approx c'$ implies $\mathcal{S}(c) \in \mathcal{Q}_F^{\mathcal{P}} \iff \mathcal{S}(c') \in \mathcal{Q}_F^{\mathcal{P}}$. \square

5. GROUP I

All TSO games in this group have the following in common: There is one player that can update messages *after* her turn, and the other player can update messages *before* her turn. Both players might be allowed to do more than that, but fortunately we do not need to differentiate between those cases. In the following, we call the player that updates after her turn *player X*, and the other one *player Y*. Although the definition of safety games seems to be of asymmetric nature (player B tries to *reach* a final configuration, while player A tries to *avoid* them), the proof does not rely on the exact identity of player X and Y.

In this section, given a configuration c , we write \bar{c} to denote a configuration obtained from c after updating all messages to the memory. More formally, $c \xrightarrow{\text{up}^*} \bar{c}$ and all buffers of \bar{c} are empty. Note that if the buffers of multiple processes contain messages at configuration c , then \bar{c} is not unique: Although the global state $\mathcal{S}(\bar{c})$ and the buffer content $\mathcal{B}(\bar{c})$ are uniquely defined, the memory $\mathcal{M}(\bar{c})$ may depend on the order in which the buffer messages have been updated to the memory.

Let $\mathcal{G} = \langle \mathcal{C}, \mathcal{C}_A, \mathcal{C}_B, \rightarrow, \mathcal{C}_F \rangle$ be a TSO game as described above, currently in some configuration $c_0 \in \mathcal{C}$. We first consider the situation where player X has a winning strategy σ_X from c_0 . Let σ_Y be an arbitrary strategy for player Y and define two more strategies $\bar{\sigma}_X : c \mapsto \sigma_X(\bar{c})$ and $\bar{\sigma}_Y : c \mapsto \sigma_Y(\bar{c})$. That is, they act like σ_X and σ_Y , respectively, with the addition that $\bar{\sigma}_X$ empties the buffer *after* each turn and $\bar{\sigma}_Y$ empties the buffer *before* each turn. From the definitions it follows directly that $\bar{\sigma}_Y(\sigma_X(c)) = \sigma_Y(\bar{\sigma}_X(c))$ for all $c \in \mathcal{C}_X$. An example can be seen in Figure 12.

We argue that $\bar{\sigma}_X$ is a winning strategy for player X. The intuition behind this is as follows: Using the notation of Figure 12, if a configuration c'' is reachable from \bar{c}' , then it is also reachable from c' , since player Y can empty all buffers at the start of her turn and then proceed as if she started in \bar{c}' . On the other hand, there might be configurations reachable from c' but not \bar{c}' , for example a read transition might get disabled by one of the buffer updates. Thus, player X never gets a disadvantage by emptying the buffers.

Claim 5.1. $\bar{\sigma}_X$ is a winning strategy from c_0 .

Proof. **Case $c_0 \in C_X$:** Since $\bar{\sigma}_Y(\sigma_X(c)) = \sigma_Y(\bar{\sigma}_X(c))$ for all $c \in C_X$, the plays $P_1 := P(c_0, \sigma_X, \bar{\sigma}_Y)$ and $P_2 := P(c_0, \bar{\sigma}_X, \sigma_Y)$ agree on every second configuration, i.e. the configurations in C_X . Moreover, the configurations in between (after an odd number of steps) at least share the same global state, i.e. $\mathcal{S}(\sigma_X(c)) = \mathcal{S}(\bar{\sigma}_X(c))$. In particular, the sequence of visited global TSO states is the same in both plays. Since σ_X is a winning strategy from c_0 , it means that P_1 is winning for player X. This means that P_2 is also winning, because for both players, a winning play is clearly determined by the sequence of visited global TSO states. Because we chose σ_Y arbitrarily, it follows that $\bar{\sigma}_X$ is a winning strategy.

Case $c_0 \in C_Y$: For the other case, we consider the configurations in $\text{Post}(c_0)$ instead. We observe that σ_X must be a winning strategy for all $c \in \text{Post}(c_0)$. We apply the first case of this proof to each of these configurations and obtain that $\bar{\sigma}_X$ is a winning strategy for all of them. It follows that $\bar{\sigma}_X$ is a winning strategy for c_0 . \square

Suppose that player X plays her modified strategy as described above. We observe that after at most two steps, every play induced by her strategy and an arbitrary strategy of the opposing player only visits configurations with at most one message in the buffers: Player X will empty all buffers at the end of each of her turns and player Y can only add at most one message to the buffers in between. Hence, they can play on a finite set of configurations instead.

To show this, we construct a finite game $\mathcal{G}' = \langle C', C'_A, C'_B, \rightarrow', C'_F \rangle$ as follows. C'_Y contains all configurations of C_Y that have at most one buffer message, i.e.:

$$C'_Y := \{ \langle \mathcal{S}, \mathcal{B}, \mathcal{M} \rangle_Y \in C_Y \mid \sum_{\iota \in \mathcal{I}} |\mathcal{B}(\iota)| \leq 1 \}$$

If $c_0 \in C_Y$, we also add it to C'_Y , otherwise to C'_X . Lastly, we add $\text{Post}(C'_Y)$ to C'_X , where Post is with respect to \mathcal{G} . \rightarrow' is defined as the restriction of \rightarrow to configurations of \mathcal{G}' , and $C'_F = C_F \cap C'_A$. Note that C'_X also contains configurations with two messages. This is needed to account for the case that player Y has a winning strategy, which is handled later in this proof. Now, let $\bar{\sigma}'_X$ be the restriction of $\bar{\sigma}_X$ to C'_X (in the mathematical sense, i.e. $\bar{\sigma}'_X : C'_X \rightarrow C_Y$ and $\bar{\sigma}_X(c) = \bar{\sigma}'_X(c)$ for all $c \in C'_X$).

Claim 5.2. $\bar{\sigma}'_X$ is a winning strategy for c_0 in \mathcal{G}' .

Proof. Looking at the definitions, we confirm that $\bar{\sigma}'_X$ actually is a valid strategy for \mathcal{G}' , i.e. $\bar{\sigma}'_X(c) \in C'_Y$, for all $c \in C'_X$, since $\bar{\sigma}'_X(c)$ has empty buffers. (This makes $\bar{\sigma}'_X$ the restriction of $\bar{\sigma}_X$ to \mathcal{G}' .) Consider a strategy σ'_Y for player Y in \mathcal{G}' and an arbitrary extension σ_Y to \mathcal{G} . Because $\bar{\sigma}'_X$ and $\bar{\sigma}_X$ agree on C'_X and $\bar{\sigma}'_Y$ and $\bar{\sigma}_Y$ agree on C'_Y , $P := P(c_0, \bar{\sigma}'_X, \bar{\sigma}_Y)$ and $P' := P(c_0, \bar{\sigma}'_X, \bar{\sigma}_Y)$ are in fact the exact same play. Since $\bar{\sigma}_X$ is a winning strategy, P is a winning play, and thus also P' . Here, note that \mathcal{G} and \mathcal{G}' agree on the final configurations within C' . Since σ'_Y was arbitrary, it follows that $\bar{\sigma}'_X$ is a winning strategy from c_0 in \mathcal{G}' . \square

What is left to show is that a winning strategy for \mathcal{G}' induces a winning strategy for \mathcal{G} . Suppose σ'_X is a winning strategy for player X in game \mathcal{G}' for the configuration c_0 . Let σ_X be an arbitrary extension of σ'_X to \mathcal{G} .

Claim 5.3. σ_X is a winning strategy for c_0 in \mathcal{G} .

Proof. Let σ_Y be a strategy of player Y in \mathcal{G} and σ'_Y the restriction of σ_Y to C'_Y (again, in the mathematical sense). Since the outgoing transitions of every $c \in C'_Y$ are the same in both \mathcal{G} and \mathcal{G}' , σ'_Y is a strategy for \mathcal{G}' (and the restriction of σ_Y to \mathcal{G}'). Furthermore,

starting from c_0 , we see that σ_X and σ_Y induce the exact same play in \mathcal{G} as σ'_X and σ'_Y in \mathcal{G}' . Since the former play is winning, so must be the latter one. \square

Now, we quickly cover the situation where it is player Y that has a winning strategy. We follow the same arguments as previously, with minor changes. This time, assume σ_Y to be a winning strategy and let σ_X be arbitrary. Define $\bar{\sigma}_X$ and $\bar{\sigma}_Y$ as above. Following the beginning of the proof of Claim 5.1, we can conclude that the sequence of visited global TSO states is the same in both play P_1 and P_2 . For the remainder of the proof, we swap the roles of X and Y and obtain that $\bar{\sigma}_Y$ is a winning strategy.

Let $\bar{\sigma}'_Y$ be the restriction of $\bar{\sigma}_Y$ to C'_Y . Since $\bar{\sigma}'_Y(C'_Y) = \bar{\sigma}_Y(C'_Y) \subseteq \text{Post}(C'_Y) \subseteq C'_X$, it follows that $\bar{\sigma}'_Y$ is a strategy of \mathcal{G}' (Post is again with respect to \mathcal{G}). Consider a strategy σ'_X for player X in \mathcal{G}' and an arbitrary extension σ_X to \mathcal{G} . Similar as in Claim 5.2, we see that $P(c_0, \bar{\sigma}'_X, \bar{\sigma}_Y) = P(c_0, \bar{\sigma}'_X, \bar{\sigma}_Y)$ and conclude that $\bar{\sigma}'_Y$ is a winning strategy.

The other direction follows from the proof of Claim 5.3, with the roles of X and Y swapped.

Theorem 5.4. *The safety problem for games of group I is EXPTIME-complete.*

Proof. By Claim 5.1 and Claim 5.2, if a configuration c_0 is winning for player X in \mathcal{G} , then it is also winning in \mathcal{G}' . The reverse holds by Claim 5.3. The equivalent statement for player Y follows from results outlined above. Thus, the safety problem for \mathcal{G} is equivalent to the safety problem for \mathcal{G}' . \mathcal{G}' is finite and has exponentially many configurations. EXPTIME-completeness follows immediately from Lemma 4.2 (membership) and Corollary 4.7 (hardness). \square

Remark 5.5. In the game where both players are allowed to update the buffer at any time, we can show an interesting conclusion. By Claim 5.1 and the equivalent statement for the second player, we can restrict both players to strategies that empty the buffer after each turn. Thus, the game is played only on configurations with empty buffer, except for the initial configuration which might contain some buffer messages. This implies that the TSO program that is described by the game implicitly follows SC semantics.

6. GROUP II

This group contains TSO games where both players are allowed to update the buffer *only* before their own move. Let player X be the player that has a winning strategy and player Y her opponent. Note that this differs from the previous section, in which the players X and Y were defined based on their updating capabilities.

Similar to the argumentation for Group I, we want to show that player X also has a winning strategy where she empties the buffer in each move. But, in contrast to before, this time there is an exception: Since the player has to update the buffer *before* her move, by updating a memory variable she might disable a read transition that she intended to execute. Thus, we do not require her to empty the buffer in that case.

Formally, let $\mathcal{G} = \langle C, C_A, C_B, \rightarrow, C_F \rangle$ be a TSO game where both players are allowed to perform buffer updates exactly before their own moves. Suppose σ_X is a winning strategy for player X and some configuration c_0 . We construct another strategy $\bar{\sigma}_X$ for player X. Let $c \in C_X$, $c' = \sigma_X(c)$ and \bar{c} as in the previous section, i.e. a (non-unique) configuration such that $c \xrightarrow{\text{up}^*} \bar{c}$ and the buffers of \bar{c} are empty. Suppose that $c \xrightarrow{\text{instr}_\ell} c'$, where instr_ℓ is not a




FIGURE 13. Commutative diagram of strategies in games of group II, in the case where $\text{instr}_\iota \neq \text{rd}(x, d)$ and $\text{instr}_\iota \neq \text{arw}(x, d, d')$.

read or atomic read-write instruction. Then, starting from c , updating all buffer messages does not change that the transition from $\mathcal{S}(c)(\iota)$ to $\mathcal{S}(c')(\iota)$ is enabled. Thus, instr_ι can also be executed from \bar{c} . We call the resulting configuration \tilde{c}' and observe that $\bar{c} \rightarrow \tilde{c}'$ and $c' \xrightarrow{\text{up}^*} \tilde{c}'$. We define $\bar{\sigma}_X(c) := \tilde{c}'$. This can be seen in Figure 13. Note that \tilde{c}' may have at most one message in its buffers. In the other case, where there is no transition from c to c' other than read or atomic read-write instructions, we define $\bar{\sigma}_X(c) := \sigma_X(c) = c'$ and observe that c' cannot have more buffer messages than c .

Claim 6.1. $\bar{\sigma}_X$ is a winning strategy for c_0 .

Proof. First, suppose that $c_0 \in C_X$ and let σ_Y be an arbitrary strategy of player Y. We define another (non-positional) strategy $\bar{\sigma}_Y$, that depends on the last two configurations, by $\bar{\sigma}_Y(c, c') := \sigma_Y(\bar{\sigma}_X(c))$. We observe that for all $c \in C_X$, it holds that $\bar{\sigma}_Y(c, \sigma_X(c)) = \sigma_Y(\bar{\sigma}_X(c))$. It follows that the play P_1 induced by σ_X and $\bar{\sigma}_Y$ and the play P_2 induced by $\bar{\sigma}_X$ and σ_Y agree on every second configuration, i.e. the configurations in C_X . In particular, the sequence of visited global TSO configurations is the same in both plays. Since σ_X is winning, it means that P_1 is winning for player X and thus also P_2 is winning. Because we chose σ_Y arbitrarily, it follows that $\bar{\sigma}_X$ is a winning strategy.

Otherwise, if $c_0 \in C_Y$, we consider the successors of c_0 instead. We note that $\bar{\sigma}_X$ must also be a winning strategy for each $c \in \text{Post}(c_0)$. But then, we can apply the previous arguments to each of those configurations and conclude that $\bar{\sigma}_X$ is a winning strategy for all of them. Thus, it is also a winning strategy for c_0 . \square

We conclude that if player X has a winning strategy σ_X , then she also has a winning strategy $\bar{\sigma}_X$ where she empties the buffers before every turn in which she does not perform a read operation. By symmetry, the same holds true for player Y. Thus, we can limit our analysis to this type of strategies. We see that the number of messages in the buffers is bounded: Suppose that the game is in configuration $c \in C_X$. Then, $\bar{\sigma}_X$ either empties the buffer and adds at most one new message, or it performs a transition due to a read instruction, which does not increase the size of the buffers. The analogous argumentation holds for player Y. Hence, we can reduce the game to a game on bounded buffers, which is finite state and thus decidable.

Given the configuration c_0 as above, we construct a finite game $\mathcal{G}' = \langle C', C'_X, C'_Y, \rightarrow', C'_F \rangle$ as follows. The set C'_X contains all configurations from C_X which have at most as many buffer messages than c_0 (or at most one message, if c_0 has empty buffers):

$$C'_X := \{c \in C_X \mid |\mathcal{B}(c)| \leq \max\{1, |\mathcal{B}(c_0)|\}\} \quad \text{where} \quad |\mathcal{B}| = \sum_{\iota \in \mathcal{I}} |\mathcal{B}(\iota)|$$

The set C'_Y is defined accordingly. Note that both sets are finite. Lastly, \rightarrow' is defined as the restriction of \rightarrow to configurations of \mathcal{G}' , and $C'_F := C_F \cap C'_A$. We define $\bar{\sigma}'_X$ to be the restriction of $\bar{\sigma}_X$ to C'_X . Since $\bar{\sigma}'_X(c) \in C'_Y$ for all $c \in C'_X$, $\bar{\sigma}'_X$ is indeed a valid strategy for \mathcal{G}' . In particular, it is the restriction of $\bar{\sigma}_X$ to \mathcal{G}' .

Claim 6.2. $\bar{\sigma}'_X$ is a winning strategy for c_0 in \mathcal{G}' .

Proof. First, consider the case where $c_0 \in C_X$. Let σ'_Y be a strategy for player Y in \mathcal{G}' and let σ_Y be an arbitrary extension of σ'_Y to \mathcal{G} . The play P induced by $\bar{\sigma}_X$ and σ_Y in \mathcal{G} is the same as the play P' induced by $\bar{\sigma}'_X$ and σ'_Y in \mathcal{G}' . Since $\bar{\sigma}_X$ is a winning strategy, P is a winning play. It follows that P' must also be a winning strategy. Since σ'_Y was arbitrary, it follows that $\bar{\sigma}'_X$ is a winning strategy and c_0 is winning in \mathcal{G}' . \square

Theorem 6.3. *The safety problem for games of group II is EXPTIME-complete.*

Proof. By Claim 6.1 and Claim 6.2, if a configuration c_0 is winning for player A in game \mathcal{G} , then it is also winning in \mathcal{G}' . The same holds true for player B. Thus, the safety problem for \mathcal{G} is equivalent to the safety problem for \mathcal{G}' . Similar to the games of group I, \mathcal{G}' is finite and has exponentially many configurations. By Lemma 4.2 and Corollary 4.7, we can again conclude that the safety problem is EXPTIME-complete. \square

7. GROUP III

This group consists of all games where exactly one player has control over the buffer updates, and additionally the game where both players are allowed to update buffer messages *after* their own move. Intuitively, all of them have in common that the TSO program can attribute a buffer update to one specific player. If only one player can update messages, this is clear. In the other game, the first player who observes that a buffer message has reached the memory is not the one who has performed the buffer update. Thus, the program is able to punish misbehaviour, i.e. not following protocols or losing messages.

We will show that the safety problem is undecidable for this group of games. To accomplish that, we reduce the state reachability problem of PCS to the safety problem of each game. Since the former problem is undecidable, so is the latter.

The case where player A is allowed to perform buffer updates at any time is called the *A-TSO game*. It is explained in detail in the following. The other cases work similar, but require slightly different program constructions. They are presented in the appendix.

Consider the A-TSO game, i.e. the case where player A can update messages at any time, but player B can never do so. Given a PCS $\mathcal{L} = \langle S, M, \delta \rangle$ and a final state $s_F \in S$, we construct a TSO program \mathcal{P} that simulates \mathcal{L} . We design the program such that s_F is reachable in \mathcal{L} if and only if player B wins the safety game induced by \mathcal{P} . Thus, the construction gives her the initiative to decide which transitions of \mathcal{L} will be simulated. Meanwhile, the task of player A is to take care of the buffer updates.

\mathcal{P} consists of three processes Proc^1 , Proc^2 and Proc^3 , that operate on the variables $\{x_{\text{wr}}, x_{\text{rd}}, y\}$ over the domain $M \uplus \{0, 1, \perp\}$. The first process simulates the control flow and the message channel of the PCS \mathcal{L} . The second process provides a mean to read from the channel. The only task of the third process is to prevent deadlocks, or rather to make any deadlocked player lose. Proc^3 achieves this with four states: the initial state, an intermediate state, and one winning state for each player, respectively. If one of the players cannot move

in both Proc^1 and Proc^2 , they have to take a transition in Proc^3 . From the initial state of this process, there exists only one outgoing transition, which is to the intermediate state. From there, the other player can move to her respective winning state and the process will only self-loop from then on. For player A, her state is winning because she can refuse to update any messages, which will ensure that player B keeps being deadlocked in Proc^1 and Proc^2 . For player B, her state simply is contained in $Q_F^{\mathcal{P}}$. In the following, we will mostly omit Proc^3 from the analysis and just assume that both players avoid reaching a configuration where they cannot take any transition in either Proc^1 or Proc^2 .

As mentioned above, we will construct Proc^1 and Proc^2 to simulate the perfect channel system in a way that gives player B the control about which channel operation will be simulated. To achieve this, each channel operation will need an even number of transitions to be simulated in \mathcal{P} . Since player B starts the game, this means that after every fully completed simulation step, it is again her turn and she can initiate another simulation step as she pleases. Furthermore, during the simulation of a skip or send operation, we want to prevent player A from executing Proc^2 , since this process is only needed for the receive operation. Suppose that we want to block player A from taking a transition $q \xrightarrow{\text{instr}} q'$. We add a new transition $q' \xrightarrow{\text{skip}} q_F$, where $q_F \in \mathcal{S}_F^{\mathcal{P}}$. Hence, reaching q' is immediately losing for player A, since player B can respond by moving to q_F .

Next, we will describe how Proc^1 and Proc^2 simulate the perfect channel system \mathcal{L} . For each transition in \mathcal{L} , we construct a sequence of transitions in Proc^1 that simulates both the state change and the channel behaviour of the \mathcal{L} -transition. To achieve this, Proc^1 uses its buffer to store the messages of the PCS's channel. In particular, to simulate a send operation $!m$, Proc^1 adds the message $\langle x_{\text{wr}}, m \rangle$ to its buffer. For receive operations, Proc^1 cannot read its own oldest buffer message, since it is overshadowed by the more recent messages. Thus, the program uses Proc^2 to read the message from memory and copies it to the variable x_{rd} , where it can be read by Proc^1 . We call the combination of reading a message m from x_{wr} and writing it to x_{rd} the *rotation* of m .

While this is sufficient to simulate all behaviours of the PCS, it also allows for additional behaviour that is not captured by \mathcal{L} . More precisely, we need to ensure that each channel message is received *once and only once*. Equivalently, we need to prevent the *loss* and *duplication* of messages. This can happen due to multiple reasons.

The first phenomenon that allows the loss of messages is the seeming lossiness of the TSO buffer. Although it is not strictly lossy, it can appear so: Consider an execution of \mathcal{P} that simulates two send operations $!m_1$ and $!m_2$, i.e. Proc^1 adds $\langle x_{\text{wr}}, m_1 \rangle$ and $\langle x_{\text{wr}}, m_2 \rangle$ to its buffer. Assume that player A decides to update both messages to the memory, without Proc^2 performing a message rotation in between. The first message m_1 is overwritten by the second message m_2 and is lost beyond recovery.

To prevent this, we extend the construction of Proc^1 such that it inserts an auxiliary message $\langle y, 1 \rangle$ into its buffer after the simulation of each send operation. After a message rotation, that is, after Proc^2 copied a message from x_{wr} to x_{rd} , the process then resets the value of x_{wr} to its initial value \perp . Next, the process checks that y contains the value 0, which indicates that only one message was updated to the memory. Now, player A is allowed to update exactly one $\langle y, 1 \rangle$ buffer message, after which Proc^2 resets y to 0. To ensure that player A has actually updated only one message in this step, Proc^2 then checks that x_{wr} is still empty. Since player A is exclusively responsible for buffer updates, Proc^2 deadlocks her whenever one of these checks fails.

In the next scenario, we discover a different way of message loss. Consider again an execution of \mathcal{P} that simulates two send operations $!m_1$ and $!m_2$. Assume Player A updates m_1 to the memory and Proc^2 performs a message rotation. Immediately afterwards, the same happens to m_2 , without Proc^1 simulating a receive operation in between. Again, m_1 is overwritten by m_2 before being received, thus it is lost.

Player A is prevented from losing a message in this way by disallowing her to perform a complete message rotation (including the update of one $\langle y, 1 \rangle$ -message and the reset of the variables) entirely on her own. More precisely, we add a winning transition for player B to Proc^2 that she can take if and only if player A is the one initiating the update of $\langle y, 1 \rangle$. On the other hand, player A can prevent player B from performing two rotations right after each other by refusing to update the next buffer message until Proc^1 initiates the simulation of a receive operation.

Lastly, we investigate message duplication. This occurs if Proc^1 simulates two receive operations without Proc^2 performing a message rotation in between. In this case, the most recently rotated message is received twice.

The program prevents this by blocking Proc^1 from progressing after a receive operation until Proc^2 has finished a full rotation. In detail, at the very end of the message rotation and $\langle y, 1 \rangle$ -update, Proc^2 reset the value of x_{rd} to its initial value \perp . After simulating a receive operation, Proc^1 is blocked until it can read this value from memory.

This concludes the mechanisms implemented to ensure that each channel message is received *once and only once*. Thus, we have constructed an A-TSO game that simulates a perfect channel system. We summarise our results in the following theorem. The formal proof can be found in Appendix A.

Theorem 7.1. *The safety problem for the A-TSO game is undecidable.*

8. GROUP IV

In TSO games where no player is allowed to perform any buffer updates, there is no communication between the processes at all. A read operation of a process Proc^l on a variable x either reads the initial value from the shared memory, or the value of the last write of Proc^l on x from the buffer, if such a write operation has happened.

Thus, we are only interested in the transitions that are enabled for each process, but we do not need to care about the actual buffer content. In particular, the information that we need to capture from the buffers and the memory is the values that each process can read from the variables, and whether a process can execute memory fence and atomic read-write instructions or not. Together with the global state of the current configuration, this completely determines the enabled transitions in the system.

We call this concept the *view* of the processes on the concurrent system and define it formally as a tuple $v = \langle \mathcal{S}, \mathcal{V}, \mathcal{F} \rangle$, where:

- $\mathcal{S} : \mathcal{I} \rightarrow \bigcup_{\iota \in \mathcal{I}} Q^\iota$ is a global state of \mathcal{P} .
- $\mathcal{V} : \mathcal{I} \times \text{Vars} \rightarrow \text{Dom}$ defines which value each process reads from a variable.
- $\mathcal{F} : \mathcal{I} \rightarrow \{\text{true}, \text{false}\}$ represents the possibility to perform a memory fence instruction.

Given a view $v = \langle \mathcal{S}, \mathcal{V}, \mathcal{F} \rangle$, we write $\mathcal{S}(v)$, $\mathcal{V}(v)$ and $\mathcal{F}(v)$ for the global program state \mathcal{S} , the value state \mathcal{V} and the fence state \mathcal{F} of v .

The view of a configuration c is denoted by $v(c)$ and defined in the following way. First, $\mathcal{S}(v(c)) = \mathcal{S}(c)$. For all $\iota \in \mathcal{I}$ and $x \in \text{Vars}$, if $\mathcal{B}(c)(\iota)|_{\{x\} \times \text{Dom}} = \langle x, d \rangle \cdot w$, then

$\mathcal{V}(\mathbf{v}(\mathbf{c}))(\iota, \mathbf{x}) = \mathbf{d}$. Otherwise, $\mathcal{V}(\mathbf{v}(\mathbf{c}))(\iota, \mathbf{x}) = \mathcal{M}(\mathbf{c})(\mathbf{x})$. Lastly, $\mathcal{F}(\mathbf{v}(\mathbf{c}))(\iota) = \text{true}$ if and only if $\mathcal{B}(\mathbf{c})(\iota) = \varepsilon$. We extend the notation to sets of configurations in the usual way, i.e. $\mathbf{v}(\mathbf{C}') := \{\mathbf{v}(\mathbf{c}) \mid \mathbf{c} \in \mathbf{C}'\}$, and to game configurations by $\mathbf{v}(\mathbf{c}_A) := \mathbf{v}(\mathbf{c})_A$ and $\mathbf{v}(\mathbf{c}_B) := \mathbf{v}(\mathbf{c})_B$.

If $\mathbf{v}(\mathbf{c}) = \mathbf{v}(\mathbf{c}')$ for some $\mathbf{c}, \mathbf{c}' \in \mathbf{C}_{\mathcal{P}}$, a local process of \mathcal{P} cannot differentiate between \mathbf{c} and \mathbf{c}' in the sense that the enabled transitions in both configurations are the same. We formalise this idea by defining a bisimulation between \mathcal{G} and a game $\mathcal{H} = \langle \mathbf{V}, \mathbf{V}_A, \mathbf{V}_B, \rightarrow, \mathbf{V}_F \rangle$ played on the views of \mathcal{G} . Define:

$$\mathbf{V}_A := \{\mathbf{v}(\mathbf{c})_A \mid \mathbf{c}_A \in \mathbf{C}_A\} \quad \mathbf{V}_B := \{\mathbf{v}(\mathbf{c})_B \mid \mathbf{c}_B \in \mathbf{C}_B\} \quad \mathbf{V}_F := \{\mathbf{v}(\mathbf{c})_A \mid \mathbf{c}_A \in \mathbf{C}_F\}$$

Lastly, we have $\mathbf{v}(\mathbf{c}) \rightarrow \mathbf{v}(\mathbf{c}')$ whenever $\mathbf{c} \rightarrow \mathbf{c}'$.

Lemma 8.1. *The relation $\mathbf{Z} := \{(\mathbf{c}, \mathbf{v}(\mathbf{c})) \mid \mathbf{c} \in \mathbf{C}\}$ is a bisimulation between \mathcal{G} and \mathcal{H} .*

Proof. We need to show that for all $(\mathbf{c}_1, \mathbf{v}_2 := \mathbf{v}(\mathbf{c}_1)) \in \mathbf{Z}$:

- For all $\mathbf{c}_1 \rightarrow \mathbf{c}_3$ in \mathcal{G} , there is $\mathbf{v}_2 \rightarrow \mathbf{v}_4$ in \mathcal{H} with $\mathbf{c}_3 \approx \mathbf{v}_4$.
- For all $\mathbf{v}_2 \rightarrow \mathbf{v}_4$ in \mathcal{H} , there is $\mathbf{c}_1 \rightarrow \mathbf{c}_3$ in \mathcal{G} with $\mathbf{c}_3 \approx \mathbf{v}_4$.
- $\mathbf{c}_1 \in \mathbf{C}_A^{\mathcal{G}}$ if and only if $\mathbf{v}_2 \in \mathbf{C}_A^{\mathcal{H}}$.
- $\mathbf{c}_1 \in \mathbf{C}_F^{\mathcal{G}}$ if and only if $\mathbf{v}_2 \in \mathbf{C}_F^{\mathcal{H}}$.

The first property is trivially fulfilled since $\mathbf{v}(\mathbf{c}_1) \rightarrow \mathbf{v}(\mathbf{c}_3)$ for any $\mathbf{c}_1 \rightarrow \mathbf{c}_3$ by definition of \mathcal{H} .

For the second property, suppose that $\mathbf{v}_2 \rightarrow \mathbf{v}_4$ is due to some transition $\mathbf{c}_2 \xrightarrow{\text{instr}_\iota} \mathbf{c}_4$. We first show that instr_ι is enabled at \mathbf{c}_1 . Since $\mathbf{v}(\mathbf{c}_1) = \mathbf{v}(\mathbf{c}_2)$, it holds that $\mathcal{S}(\mathbf{c}_1) = \mathcal{S}(\mathbf{c}_2)$. Furthermore, if $\text{instr}_\iota = \text{rd}(\mathbf{x}, \mathbf{d})_\iota$, then $\mathcal{V}(\mathbf{v}(\mathbf{c}_1))(\iota, \mathbf{x}) = \mathcal{V}(\mathbf{v}(\mathbf{c}_2))(\iota, \mathbf{x}) = \mathbf{d}$. Also, if $\text{instr}_\iota = \text{mf}_\iota$, then $\mathcal{B}(\mathbf{c}_2)(\iota) = \varepsilon$ and since $\mathcal{F}(\mathbf{v}(\mathbf{c}_1))(\iota) = \mathcal{F}(\mathbf{v}(\mathbf{c}_2))(\iota) = \text{true}$ it follows that $\mathcal{B}(\mathbf{c}_1)(\iota) = \varepsilon$. Similarly, if $\text{instr}_\iota = \text{arw}(\mathbf{x}, \mathbf{d}, \mathbf{d}')_\iota$, then $\mathcal{V}(\mathbf{v}(\mathbf{c}_1))(\iota, \mathbf{x}) = \mathcal{V}(\mathbf{v}(\mathbf{c}_2))(\iota, \mathbf{x}) = \mathbf{d}$ and $\mathcal{B}(\mathbf{c}_1)(\iota) = \mathcal{B}(\mathbf{c}_2)(\iota) = \varepsilon$. From these considerations and the definition of the TSO semantics (see Figure 5), it follows that instr_ι is indeed enabled at \mathbf{c}_1 .

Let \mathbf{c}_3 be the configuration obtained after performing instr_ι , i.e. $\mathbf{c}_1 \xrightarrow{\text{instr}_\iota} \mathbf{c}_3$. It holds that $\mathcal{S}(\mathbf{c}_3) = \mathcal{S}(\mathbf{c}_4) = \mathcal{S}(\mathbf{c}_1)[\iota \leftarrow \mathcal{S}(\mathbf{c}_4)(\iota)]$. If $\text{instr}_\iota = \text{wr}(\mathbf{x}, \mathbf{d})_\iota$, then $\mathcal{V}(\mathbf{v}(\mathbf{c}_3)) = \mathcal{V}(\mathbf{v}(\mathbf{c}_4)) = \mathcal{V}(\mathbf{v}(\mathbf{c}_1))[(\iota, \mathbf{x}) \leftarrow \mathbf{d}]$ and $\mathcal{F}(\mathbf{v}(\mathbf{c}_3)) = \mathcal{F}(\mathbf{v}(\mathbf{c}_4)) = \mathcal{F}(\mathbf{v}(\mathbf{c}_1))[\iota \leftarrow \text{false}]$. Similarly, if $\text{instr}_\iota = \text{arw}(\mathbf{x}, \mathbf{d}, \mathbf{d}')_\iota$, then $\mathcal{V}(\mathbf{v}(\mathbf{c}_3)) = \mathcal{V}(\mathbf{v}(\mathbf{c}_4)) = \mathcal{V}(\mathbf{v}(\mathbf{c}_1))[(\iota, \mathbf{x}) \leftarrow \mathbf{d}']$ but $\mathcal{F}(\mathbf{v}(\mathbf{c}_3)) = \mathcal{F}(\mathbf{v}(\mathbf{c}_4)) = \mathcal{F}(\mathbf{v}(\mathbf{c}_1))$. Otherwise, $\mathcal{V}(\mathbf{v}(\mathbf{c}_3)) = \mathcal{V}(\mathbf{v}(\mathbf{c}_4)) = \mathcal{V}(\mathbf{v}(\mathbf{c}_1))$ and $\mathcal{F}(\mathbf{v}(\mathbf{c}_3)) = \mathcal{F}(\mathbf{v}(\mathbf{c}_4)) = \mathcal{F}(\mathbf{v}(\mathbf{c}_1))$. In all cases it follows that $\mathbf{v}(\mathbf{c}_3) = \mathbf{v}(\mathbf{c}_4) = \mathbf{v}_4$.

The third and fourth properties are trivially fulfilled by the definition of \mathbf{Z} and \mathcal{H} . \square

Theorem 8.2. *The safety problem for games in group IV is EXPTIME-complete.*

Proof. Apply Lemma 4.3 to \mathbf{Z} and obtain that the safety problem for \mathcal{G} is equivalent to the safety problem of \mathcal{H} . Since there exist only exponentially many views, EXPTIME-completeness follows from Lemma 4.2 and Corollary 4.7. \square

9. CONCLUSION AND FUTURE WORK

In this work we have addressed for the first time the game problem for programs running under weak memory models in general and TSO in particular. Surprisingly, our results show that depending on when the updates take place, the problem can turn out to be undecidable or decidable. In fact, there is a subtle difference between the decidable (group I, II and IV) and undecidable (group III) TSO games. For the former games, when a player is taking a turn, the system does not know who was responsible for the last update. But for the latter

games, the last update can be attributed to a specific player. Another surprising finding is the complexity of the game problem for the groups I, II and IV which is EXPTIME-complete in contrast with the non-primitive recursive complexity of the reachability problem for programs running under TSO and the undecidability of the repeated reachability problem.

In future work, the games where exactly one player has control over the buffer seem to be the most natural ones to expand on. In particular, the A-TSO game (where player A can update before and after her move) and the B-TSO game (same, but for player B). On the other hand, the games of groups I, II and IV seem to be degenerate cases and therefore rather uninteresting. In particular, we have shown that they are not more powerful than games on programs that follow SC semantics.

Another direction for future work is considering other memory models, such as the partial store ordering semantics, the release-acquire semantics, and the ARM semantics. It is also interesting to define stochastic games for programs running under TSO as extension of the probabilistic TSO semantics [AAA⁺22].

REFERENCES

- [AAA⁺22] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Raj Aryan Agarwal, Adwait Godbole, and Shankara Narayanan Krishna. Probabilistic total store ordering. In Ilya Sergey, editor, *Programming Languages and Systems - 31st European Symposium on Programming, ESOP 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings*, volume 13240 of *Lecture Notes in Computer Science*, pages 317–345. Springer, 2022. doi:10.1007/978-3-030-99336-8_12.
- [AABN16] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Ahmed Bouajjani, and Tuan Phong Ngo. The benefits of duality in verifying concurrent programs under TSO. In Josée Desharnais and Radha Jagadeesan, editors, *27th International Conference on Concurrency Theory, CONCUR 2016, August 23-26, 2016, Québec City, Canada*, volume 59 of *LIPIcs*, pages 5:1–5:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.CONCUR.2016.5.
- [AABN18] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Ahmed Bouajjani, and Tuan Phong Ngo. A load-buffer semantics for total store ordering. *Log. Methods Comput. Sci.*, 14(1), 2018. doi:10.23638/LMCS-14(1:9)2018.
- [AAC⁺12] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Yu-Fang Chen, Carl Leonardsson, and Ahmed Rezzine. Counter-example guided fence insertion under TSO. In Cormac Flanagan and Barbara König, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 18th International Conference, TACAS 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012. Proceedings*, volume 7214 of *Lecture Notes in Computer Science*, pages 204–219. Springer, 2012. doi:10.1007/978-3-642-28756-5_15.
- [AAF⁺23] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Florian Furbach, Adwait Amit Godbole, Yacoub G. Hendi, Shankara Narayanan Krishna, and Stephan Spengler. Parameterized verification under TSO with data types. In Sriram Sankaranarayanan and Natasha Sharygina, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 29th International Conference, TACAS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22-27, 2023, Proceedings, Part I*, volume 13993 of *Lecture Notes in Computer Science*, pages 588–606. Springer, 2023. doi:10.1007/978-3-031-30823-9_30.
- [AAP15] Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Ngo Tuan Phong. The best of both worlds: Trading efficiency and optimality in fence insertion for TSO. In Jan Vitek, editor, *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9032 of *Lecture Notes in Computer Science*, pages 308–332. Springer, 2015. doi:10.1007/978-3-662-46669-8_13.

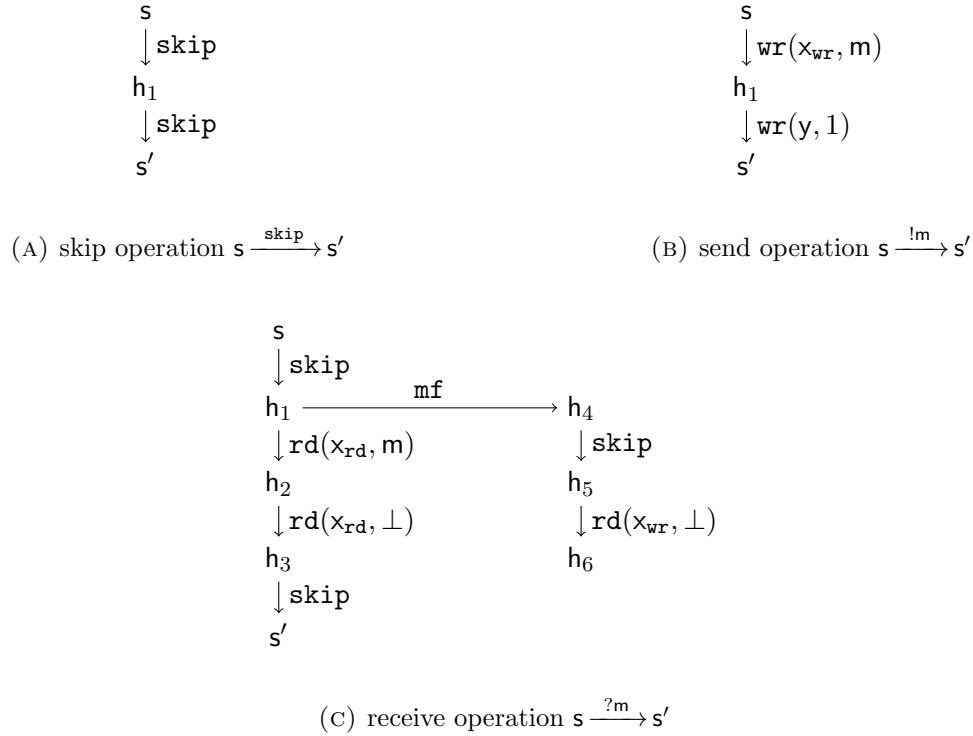
- [AAR20] Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Rojin Rezvan. Parameterized verification under TSO is pspace-complete. *Proc. ACM Program. Lang.*, 4(POPL):26:1–26:29, 2020. doi:10.1145/3371094.
- [ABBM10] Mohamed Faouzi Atig, Ahmed Bouajjani, Sebastian Burckhardt, and Madanlal Musuvathi. On the verification problem for weak memory models. In Manuel V. Hermenegildo and Jens Palsberg, editors, *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*, pages 7–18. ACM, 2010. doi:10.1145/1706299.1706303.
- [ABd08] Parosh Aziz Abdulla, Ahmed Bouajjani, and Julien d’Orso. Monotonic and downward closed games. *J. Log. Comput.*, 18(1):153–169, 2008. doi:10.1093/logcom/exm062.
- [AJ94] Parosh Aziz Abdulla and Bengt Jonsson. Undecidable verification problems for programs with unreliable channels. In Serge Abiteboul and Eli Shamir, editors, *Automata, Languages and Programming, 21st International Colloquium, ICALP94, Jerusalem, Israel, July 11-14, 1994, Proceedings*, volume 820 of *Lecture Notes in Computer Science*, pages 316–327. Springer, 1994. doi:10.1007/3-540-58201-0_78.
- [AJ96] Parosh Aziz Abdulla and Bengt Jonsson. Verifying programs with unreliable channels. *Inf. Comput.*, 127(2):91–101, 1996. doi:10.1006/inco.1996.0053.
- [ARM14] ARM. *ARM Architecture Reference Manual, ARmv7-A and ARmv7-R edition*, 2014. URL: <https://developer.arm.com/documentation/ddi0406/latest/>.
- [Ati20] Mohamed Faouzi Atig. What is decidable under the TSO memory model? *ACM SIGLOG News*, 7(4):4–19, 2020. doi:10.1145/3458593.3458595.
- [AVW03] A. Arnold, A. Vincent, and I. Walukiewicz. Games for synthesis of controllers with partial observation. *Theoretical Computer Science*, 303(1):7–34, 2003. Logic and Complexity in Computer Science. URL: <https://www.sciencedirect.com/science/article/pii/S0304397502004425>, doi:10.1016/S0304-3975(02)00442-5.
- [BDM13] Ahmed Bouajjani, Egor Derevenetc, and Roland Meyer. Checking and enforcing robustness against TSO. In Matthias Felleisen and Philippa Gardner, editors, *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*, volume 7792 of *Lecture Notes in Computer Science*, pages 533–553. Springer, 2013. doi:10.1007/978-3-642-37036-6_29.
- [BMM11] Ahmed Bouajjani, Roland Meyer, and Eike Möhlmann. Deciding robustness against total store ordering. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part II*, volume 6756 of *Lecture Notes in Computer Science*, pages 428–440. Springer, 2011. doi:10.1007/978-3-642-22012-8_34.
- [BS04] R.-J. Back and C.C. Seceleanu. Contracts and games in controller synthesis for discrete systems. In *Proceedings. 11th IEEE International Conference and Workshop on the Engineering of Computer-Based Systems, 2004.*, pages 307–314, 2004. doi:10.1109/ECBS.2004.1316713.
- [BZ83] Daniel Brand and Pitro Zafiropulo. On communicating finite-state machines. *J. ACM*, 30(2):323–342, 1983. doi:10.1145/322374.322380.
- [CKS81] Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981. doi:10.1145/322234.322243.
- [CS76] Ashok K. Chandra and Larry J. Stockmeyer. Alternation. In *17th Annual Symposium on Foundations of Computer Science, Houston, Texas, USA, 25-27 October 1976*, pages 98–108. IEEE Computer Society, 1976. doi:10.1109/SFCS.1976.4.
- [Fin87] Alain Finkel. A generalization of the procedure of karp and miller to well structured transition systems. In Thomas Ottmann, editor, *Automata, Languages and Programming, 14th International Colloquium, ICALP87, Karlsruhe, Germany, July 13-17, 1987, Proceedings*, volume 267 of *Lecture Notes in Computer Science*, pages 499–508. Springer, 1987. doi:10.1007/3-540-18088-5_43.
- [FS01] Alain Finkel and Philippe Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1-2):63–92, 2001. doi:10.1016/S0304-3975(00)00102-X.
- [GTW02] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002. doi:10.1007/3-540-36387-4.

- [IBM21] IBM. *Power ISA, Version 3.1b*, 2021. URL: https://files.openpower.foundation/s/dAYSdGzTfW4j2r2/download/OPF_PowerISA_v3.1B.pdf.
- [Int12] Intel Corporation. *Intel 64 and IA-32 Architectures Software Developers Manual*, 2012. URL: <https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html>.
- [KV97] Oded Kupferman and Mosab Y. Vardi. Synthesis with incomplete information. In *Proceedings of the 2nd International Conference on Temporal Logic*, volume 171 of *Lecture Notes in Computer Science*, pages 91–106. Springer, 1997. doi:10.1007/978-94-015-9586-5_6.
- [KV00] Orna Kupferman and Moshe Y. Vardi. μ -calculus synthesis. In Mogens Nielsen and Branislav Rovan, editors, *Mathematical Foundations of Computer Science 2000, 25th International Symposium, MFCS 2000, Bratislava, Slovakia, August 28 - September 1, 2000, Proceedings*, volume 1893 of *Lecture Notes in Computer Science*, pages 497–507. Springer, 2000. doi:10.1007/3-540-44612-5_45.
- [Lam79] Leslie Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691, 1979. doi:10.1109/TC.1979.1675439.
- [Maz01] René Mazala. Infinite games. In Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors, *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*, pages 23–42. Springer, 2001. doi:10.1007/3-540-36387-4_2.
- [OSS09] Scott Owens, Susmit Sarkar, and Peter Sewell. A better x86 memory model: x86-tso. In Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel, editors, *Theorem Proving in Higher Order Logics, 22nd International Conference, TPHOLs 2009, Munich, Germany, August 17-20, 2009. Proceedings*, volume 5674 of *Lecture Notes in Computer Science*, pages 391–407. Springer, 2009. doi:10.1007/978-3-642-03359-9_27.
- [PR89] Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, January 11-13, 1989*, pages 179–190. ACM Press, 1989. doi:10.1145/75277.75293.
- [Sch02] Philippe Schnoebelen. Verifying lossy channel systems has nonprimitive recursive complexity. *Inf. Process. Lett.*, 83(5):251–261, 2002. doi:10.1016/S0020-0190(01)00337-4.
- [spa94] SPARC International, Inc. *SPARC Architecture Manual Version 9*, 1994. URL: <https://sparc.org/wp-content/uploads/2014/01/SPARCV9.pdf.gz>.
- [SS23] Stephan Spengler and Sanchari Sil. TSO games - on the decidability of safety games under the total store order semantics. In Antonis Achilleos and Dario Della Monica, editors, *Proceedings of the Fourteenth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2023, Udine, Italy, 18-20th September 2023*, volume 390 of *EPTCS*, pages 82–98, 2023. doi:10.4204/EPTCS.390.6.
- [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, 2010. doi:10.1145/1785414.1785443.

APPENDIX A. UNDECIDABILITY OF THE A-TSO GAME

In this section, we will formally prove the correctness of the reduction presented in section 7. We begin with the construction of the TSO program $\mathcal{P} = \langle \text{Proc}^1, \text{Proc}^2, \text{Proc}^3 \rangle$. For Proc^1 , we start by adding the states of \mathcal{S} to \mathcal{Q}^1 . Then, for each transition of \mathcal{L} , we add some auxiliary states h_1, h_2, \dots and transitions between them to simulate the PCS behaviour. Figure 14 shows this construction for each of the transitions $s \xrightarrow{\text{skip}} s'$, $s \xrightarrow{!m} s'$ and $s \xrightarrow{?m} s'$. Note that all the auxiliary states of Proc^1 (labelled as h_i) are supposed to be distinct for each transition.

Additionally, Figure 15 and Figure 16 show the process definitions of Proc^2 and Proc^3 , respectively. The set of final states is $\mathcal{Q}_F^{\mathcal{P}} := \mathcal{S}_F \cup \{q_F, r_F\}$, where $\mathcal{S}_F \subset \mathcal{Q}^1$, $q_F \in \mathcal{Q}^2$ and $r_F \in \mathcal{Q}^3$.

FIGURE 14. Proc^1 of the A-TSO reduction from PCS



Theorem A.1. *Consider the A-TSO game, where player A can update the buffer before and after her turn, but player B can never update. The set of final states S_F of \mathcal{L} is reachable from $s_0 \in S$ if and only if player B wins the game $\mathcal{G}^{\text{TSO}}(\mathcal{P}, Q_F^{\mathcal{P}})$ starting from the configuration $c_0 := \langle \langle s_0, q_1, r_1 \rangle, \langle \varepsilon, \varepsilon, \varepsilon \rangle, \{x_{\text{wr}} \mapsto \perp, x_{\text{rd}} \mapsto \perp, y \mapsto 0\} \rangle_B \in C_B$.*

Note that the definition of Proc^3 secures deadlock-freedom of \mathcal{P} . Furthermore, if one of the players has no enabled transitions in either Proc^1 or Proc^2 in her turn, she will lose the game: If player A has no enabled transition, she needs to move from r_1 to r_2 in Proc^3 . Player B can then respond with $r_2 \rightarrow r_3 \in Q_F^{\mathcal{P}}$, which player A cannot escape from. Similarly, if player B needs to move to r_2 , player A can respond with r_4 . Since player A controls the buffer updates, she can decide to not perform any buffer updates, which maintains the situation that player B can only move in Proc^3 . Thus, the game will loop in $r_4 \rightarrow r_4$ forever, which means that player A wins.

In the following, we say that a player is *deadlocked*, if it is her turn and there is no enabled outgoing transition in either Proc^1 or Proc^2 . Naturally, both players avoid being deadlocked, since it means that they lose the game, respectively.

A.1. From \mathcal{L} -Reachability to a Winning Strategy in A-TSO. Suppose S_F is reachable from s_0 . Let $c_0^{\mathcal{L}} \xrightarrow{\text{op}_1} c_1^{\mathcal{L}} \xrightarrow{\text{op}_2} \dots \xrightarrow{\text{op}_n} c_n^{\mathcal{L}}$ be a run with $c_k^{\mathcal{L}} = \langle s_k, w_k \rangle$ for all $k = 1, \dots, n$ and $s_n \in S_F$. Here, w_k is a word over M consisting of the letters $w_k[1], \dots, w_k[|w_k|]$. Without loss of generality we can assume that this run does not contain any duplicate configurations.

For all $k = 0, \dots, n$, we define a game configuration $c_k = \langle S_k, \mathcal{B}_k, \mathcal{M}_k \rangle_B \in C_B$:

FIGURE 15. Proc^2 of the A-TSO reduction from PCSFIGURE 16. Proc^3 of the A-TSO reduction from PCS

- (1) $\mathcal{S}_k := \langle s_k, q_1, r_1 \rangle$
- (2) $\mathcal{B}_k := \langle \langle y, 1 \rangle, \langle x_{wr}, w_k[1] \rangle, \dots, \langle y, 1 \rangle, \langle x_{wr}, w_k[|w_k|] \rangle, \varepsilon, \varepsilon \rangle$
- (3) $\mathcal{M}_k := \{x_{wr} \mapsto \perp, x_{rd} \mapsto \perp, y \mapsto 0\}$

The basic idea is to develop a strategy for player B that visits all configurations c_1, \dots, c_n . Inspecting the construction of Proc^2 , we observe that there is no way to prevent player A from performing the buffer update and rotation of the oldest channel message whenever she decides to do so. We account for that by introducing additional configurations $c'_k = \langle \mathcal{S}'_k, \mathcal{B}'_k, \mathcal{M}'_k \rangle_B$:

- (1) $\mathcal{S}'_k := \langle s_k, q_3, r_1 \rangle$

- (2) $\mathcal{B}'_k := \langle \langle y, 1 \rangle, \langle x_{\text{wr}}, w_k[1] \rangle, \dots, \langle y, 1 \rangle, \langle x_{\text{wr}}, w_k[|w_k| - 1] \rangle, \langle y, 1 \rangle \rangle, \varepsilon, \varepsilon \rangle$
 (3) $\mathcal{M}_k^{(3)} := \{x_{\text{wr}} \mapsto w_k[|w_k|], x_{\text{rd}} \mapsto w_k[|w_k|], y \mapsto 0\}$

The configurations c'_k refer to the case where player A has updated the oldest message to the memory and also performed the message rotation. Note that if $w_k = \varepsilon$, then c'_k is not defined.

As a first step, consider a situation where it is player A's turn and Proc^2 is in state q_1 . If player A decides to initiate the message rotation by following the transition $q_1 \xrightarrow{\text{rd}(x_{\text{wr}}, m)} q_m$ for some $m \in M$, we always let player B immediately respond with $q_m \xrightarrow{\text{wr}(x_{\text{rd}}, m)} q_3$. On the other hand, if Proc^2 is already in state q_3 , player A cannot take any further transitions in this process, since after $q_3 \xrightarrow{\text{wr}(x_{\text{wr}}, \perp)} q_4$, player B can take the transition $q_4 \xrightarrow{\text{skip}} q_F$ and wins. We will use these observations in the following argumentation.

We show by induction that starting from c_0 , Player B can force a play that either visits one of c_k , c'_k or C_F . For $k = 0$, this clearly holds true. So, suppose that the claim holds for some arbitrary but fixed $0 \leq k < n$. By the induction hypothesis, we can assume that the game is either in state c_k or c'_k .

If $\text{op}_{k+1} = \text{skip}$, then player B takes the transition $s_k \xrightarrow{\text{skip}} h_1$ in Proc^1 . After a possible message rotation in Proc^2 (i.e. one move by player A and B, respectively), player A is left with only one non-losing transition, which is $h_1 \xrightarrow{\text{skip}} s_{k+1}$ in Proc^1 . Assuming that player A can prevent an immediate defeat, we now need to show that the game is either in configuration c_{k+1} or c'_{k+1} .

First, consider the case where Proc^2 is in state q_1 , which implies that the sub-run started at c_k rather than c'_k . If player A had updated at least one message from the buffer to the memory, x_{wr} would now contain some channel message m . Thus, player B could immediately win by taking the transition $q_1 \xrightarrow{\text{rd}(x_{\text{wr}}, m)} q_F$. We conclude that the game must be in configuration c_{k+1} .

In the other case, Proc^2 is in state q_3 . It follows that either the sub-run started in c'_k or player A has updated a message from the buffer and performed the message rotation. Suppose she had additionally updated at least one more message, in particular a message $\langle y, 1 \rangle$. Then, player B is again able to win immediately by following $q_3 \xrightarrow{\text{rd}(y, 1)} q_F$. Thus, the game must be in configuration c'_{k+1} .

Next, we examine the situation where $\text{op}_{k+1} = !m$, which is very similar to the previous one. Player B takes the transition $s_k \xrightarrow{\text{wr}(x_{\text{wr}}, m)} h_1$ and player A eventually responds with $h_1 \xrightarrow{\text{wr}(y, 1)} s_{k+1}$. Using the same argumentation as above, either the game is now in configuration c_{k+1} or c'_{k+1} , or player B can force an immediate win.

Lastly, if $\text{op}_{k+1} = ?m$, player B starts by taking the transition $s_k \xrightarrow{\text{skip}} h_1$ in Proc^1 . For player A, $h_1 \xrightarrow{\text{mf}} h_4$ is losing: Since $\text{op}_{k+1} = ?m$, we know that $w_k \neq \varepsilon$ and after flushing the buffer, the last character of w_k is written to x_{wr} . Thus, after player B responds with $h_4 \xrightarrow{\text{skip}} h_5$, player A is effectively deadlocked, since the outgoing transition in Proc^1 is disabled and all actions in Proc^2 will eventually lead to player B reaching q_F .

We conclude that player A updates and rotates the oldest message if not already done so, and then proceeds with $h_1 \xrightarrow{\text{rd}(x_{\text{rd}}, m)} h_2$ in Proc^1 . The only possible continuation for

both players is now in Proc^2 :

$$q_3 \xrightarrow{\text{wr}(x_{\text{wr}}, \perp)} q_4 \xrightarrow{\text{mf}} q_5 \xrightarrow{\text{skip}} q_6 \xrightarrow{\text{rd}(y, 1)} q_7 \xrightarrow{\text{wr}(y, 0)} q_8 \xrightarrow{\text{mf}} q_9 \xrightarrow{\text{wr}(x_{\text{rd}}, \perp)} q_{10}$$

Because of $q_5 \xrightarrow{\text{rd}(y, 1)} q_F$, player A cannot update the pending $\langle y, 1 \rangle$ -message before reaching state q_6 , but then definitely has to do so to enable the next transition. In q_9 , y is reset to 0 again. Furthermore, the transition $q_9 \xrightarrow{\text{rd}(x_{\text{wr}}, m)} q_F$ ensures that no other message was updated. When q_{10} is reached, player A needs to flush the buffer of Proc^2 again, to enable both $h_2 \xrightarrow{\text{rd}(x_{\text{rd}}, \perp)} h_3$ in Proc^1 and $q_9 \xrightarrow{\text{mf}} q_{10}$ in Proc^2 . She then takes one of those two transitions and player B takes the other one. Eventually (after maybe the update and rotation of the next message), player A will take the transition $h_3 \xrightarrow{\text{skip}} s_{k+1}$ in Proc^1 .

At this point, Proc^1 moved from s_k to s_{k+1} and Proc^2 is back in q_1 or q_3 . Exactly one $\langle y, 1 \rangle$ -message was removed from the buffer and the memory values of y and x_{rd} were reset to 0 and \perp , respectively. Using the same argumentation as previously, we conclude that the game is now either in configuration c_{k+1} or c'_{k+1} , or player B can force an immediate win.

This concludes the proof by induction. In summary, starting from c_0 , player B can force a play that reaches C_B or one of the configurations c_n or c'_n . In both of these configurations, Proc^1 is in the local state s_n , which is a final state of \mathcal{L} . Thus, $c_n, c'_n \in C_B$ and player B wins in any case.

A.2. From a Winning Strategy in A-TSO to \mathcal{L} -Reachability. For the other direction, suppose that player B has a winning strategy. Consider a strategy of player A that avoids reaching both $q_F \in Q^2 \cap Q_F^P$ and $r_F \in Q^3 \cap Q_F^P$ whenever possible. Furthermore, assume that player A prefers to not update any messages and to move in Proc^2 , as long as it does not contradict the previous statement.

Let s_0, \dots, s_n be the sequence of channel states that are visited by Proc^1 during the run induced by these two strategies. Since player B uses a winning strategy, this sequence is finite. We will show by induction that for each $k = 0, \dots, n$, the run contains a game configuration $c_k = \langle \mathcal{S}_k, \mathcal{B}_k, \mathcal{M}_k \rangle_B \in C_B$ such that:

- (1) $\mathcal{S}_k := \langle s_k, q_1, r_1 \rangle$
- (2) $\mathcal{B}_k := \langle \langle \langle y, 1 \rangle, \langle x_{\text{wr}}, w_k[1] \rangle, \dots, \langle y, 1 \rangle, \langle x_{\text{wr}}, w_k[|w_k|] \rangle, \varepsilon, \varepsilon \rangle$, where $w_k \in M^*$
- (3) $\mathcal{M}_k := \{x_{\text{wr}} \mapsto \perp, x_{\text{rd}} \mapsto \perp, y \mapsto 0\}$
- (4) If $k > 0$, there is a label op_k such that $\langle s_{k-1}, w_{k-1} \rangle \xrightarrow{\text{op}_k} \langle s_k, w_k \rangle$.

For $k = 0$, this clearly holds true. So, suppose that the claim holds for some arbitrary but fixed $0 \leq k < n$.

Starting in c_k , player B cannot move in Proc^2 , since $\mathcal{M}(x_{\text{wr}}) = \perp$, $\mathcal{B}(2) = \varepsilon$ and the only outgoing transitions from $\mathcal{S}(2) = q_1$ are of the form $\text{rd}(x_{\text{wr}}, m)$ for $m \in M$. Thus, player B must first move in Proc^1 instead. Looking at the construction of Proc^1 , we see that the sub-run from c_k to (the very next occurrence of) s_{k+1} takes a sequence of transitions that were added to \mathcal{P} due to some unique transition $s_k \xrightarrow{\text{op}_{k+1}} s_{k+1}$ of \mathcal{L} .

If $\text{op}_{k+1} = \text{skip}$, we conclude that player B takes the transition $s_k \xrightarrow{\text{skip}}_P h_1$. Player A has to respond with the only enabled transition in the resulting configuration, which is $h_1 \xrightarrow{\text{skip}}_P s_{k+1}$. The game is now in a configuration with states Q_k , as defined above. Furthermore, we conclude that the buffer must be \mathcal{B}_{k+1} , where $w_{k+1} := w_k$, and the memory

is \mathcal{M}_{k+1} . Finally, we have that $\langle s_k, w_k \rangle \xrightarrow{\text{skip}} \langle s_{k+1}, w_{k+1} \rangle$. In other words, the game is now in position c_k .

If $\text{op}_{k+1} = !m$, we argue in the very same way as in the previous case that the run must have followed the transitions $s_k \xrightarrow{\text{wr}(x_{\text{wr}}, m)} h_1 \xrightarrow{\text{wr}(y, 1)} s_{k+1}$ in Proc^1 . Similarly, the game is now in configuration $c_{k+1} = \langle \mathcal{S}_{k+1}, \mathcal{B}_{k+1}, \mathcal{M}_{k+1} \rangle_B$ with $w_{k+1} := m \bullet w_k$.

Lastly, if $\text{op}_{k+1} = ?m$, then player B takes $s_k \xrightarrow{\text{skip}} h_1$. Assume that $w_k = \varepsilon$, i.e. the buffer of Proc^1 is empty. Then, player A cannot update any message to the memory, which means that the transition $h_1 \xrightarrow{\text{mf}} h_4$ is the only one that is enabled. This leads to the path $h_1 \xrightarrow{\text{mf}} h_4 \xrightarrow{\text{skip}} h_5 \xrightarrow{\text{rd}(x_{\text{wr}}, \perp)} h_6$, which ends in a deadlocked configuration for player B. Since we already know that player B is winning, we conclude that the assumption was wrong and there is at least one pair of messages in the buffer of Proc^1 .

Thus, player A instead performs one buffer update of the oldest message $\langle x_{\text{wr}}, m \rangle$, where m is the last character of w_k . This enables $q_1 \xrightarrow{\text{rd}(x_{\text{wr}}, m)} q_m$ in Proc^2 , which player A executes. Player B responds with $q_m \xrightarrow{\text{wr}(x_{\text{rd}}, m)} q_3$ in Proc^2 . Player A immediately updates this message to the memory, which enables $h_1 \xrightarrow{\text{rd}(x_{\text{rd}}, m)} h_2$ in Proc^1 , which she takes. Note that this is her only choice since $q_3 \xrightarrow{\text{wr}(x_{\text{wr}}, \perp)} q_4$ opens up the possibility for player B to end the game by moving to q_F .

Now, Proc^1 is blocked again and the run continues in Proc^2 . After

$$q_3 \xrightarrow{\text{wr}(x_{\text{wr}}, \perp)} q_4 \xrightarrow{\text{mf}} q_5 \xrightarrow{\text{skip}} q_6 ,$$

player A updates the oldest buffer message of Proc^1 , which is $\langle y, 1 \rangle$. This is needed to enable

$$q_6 \xrightarrow{\text{rd}(y, 1)} q_7 \xrightarrow{\text{wr}(y, 0)} q_8 \xrightarrow{\text{mf}} q_9 \xrightarrow{\text{wr}(x_{\text{rd}}, \perp)} q_{10} .$$

Now, player A empties the buffer of Proc^2 to execute the transition $q_{10} \xrightarrow{\text{mf}} q_1$. This is followed by $h_2 \xrightarrow{\text{rd}(x_{\text{rd}}, \perp)} h_3 \xrightarrow{\text{skip}} s_{k+1}$ in Proc^1 .

We see that Proc^1 is now in state s_{k+1} and Proc^2 is back in state q_1 , i.e. the process state is \mathcal{S}_{k+1} . Since all memory values have also been reset to their default values, we conclude that the buffer state of the current game configuration is \mathcal{B}_{k+1} , where $w_{k+1} \bullet m := w_k$, and the memory is \mathcal{M}_{k+1} . Furthermore, the existence of the transition $s_k \xrightarrow{?m} s_{k+1}$ follows from the construction of \mathcal{P} , and it is enabled at the configuration $\langle s_k, w_k \rangle$ of \mathcal{L} , since m is the last character in w_k .

This concludes the proof by induction. In summary, we showed that \mathcal{L} can reach the state s_n starting from s_0 . What is left to show is that $s_n \in S_F$. Consider the continuation of the run after c_n . If player B starts to simulate a skip or send operation, the run will inadvertently visit another state of S after two moves, which is a contradiction, since we know that s_n is the last such state that is visited. Otherwise, if player B starts to simulate a receive operation, there are again two cases to consider. If the buffer of Proc^1 is empty, we have already seen that player A wins, which we know is not the case. Otherwise, player A can just proceed as usually and continue to simulate the receive operation in Proc^1 and Proc^2 . In any case, player B has no opportunity to force the play into q_F or into a deadlock of player A. In particular, player A can avoid to take $q_3 \xrightarrow{\text{skip}} q_4$, since there are exactly four (i.e. an even number of) moves to take before.

We conclude that the only way that player B wins the game is that $s_n \in Q_F^{\mathcal{P}}$, or $s_n \in S_F$, equivalently. Since $s_k \rightarrow s_{k+1}$ for all $k < n$, it follows that S_F is reachable from s_0 .

A.3. Proof of Theorem 7.1. This follows directly from Theorem A.1 and the undecidability of the state reachability problem for perfect channel systems [BZ83].

A.4. Variants of A-TSO. In the game where player A is restricted to perform updates only before her turn, one can check that the construction still works. However, if player A is only allowed to update after her turn, we need to adjust the construction. This is due to the fact that in the A-TSO game, there exist transitions where she is expected to update before her move.

We can easily fix this problem by adding two auxiliary states just before each of the affected transitions, and connect them with `skip` instructions. This gives player A time to update the buffer after taking the auxiliary transition.

Furthermore, we add a transition from the second auxiliary state to a sink state, that is, a state with just a self-loop. This prevents player B from taking the first transition, since player A can then go to the sink state, which effectively disables the process. It also ensures that whenever player A takes the first auxiliary transition, player B has to respond by taking the second one.

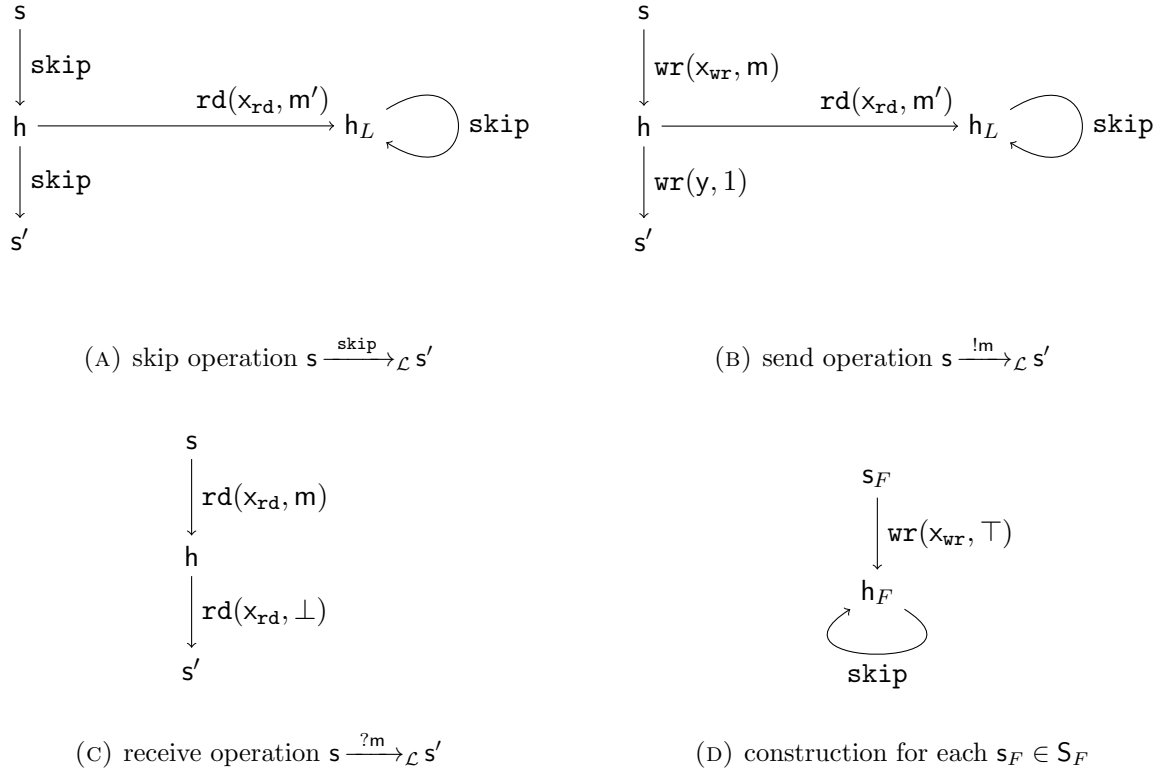
We conclude that the modified construction gives rise to the same behaviour than the construction for A-TSO. The formal proof is along the lines of Theorem A.1 or Theorem B.1.

APPENDIX B. UNDECIDABILITY OF THE B-TSO GAME

In this section we prove the undecidability of the B-TSO game, which is the TSO game where player B has full control over the buffer updates. Similar to what was presented in the A-TSO section, we construct a program consisting of two processes $\mathcal{P} = \langle \text{Proc}^1, \text{Proc}^2 \rangle$. As previously, Proc^1 starts out with all states of S and is then augmented by auxiliary states and transitions to simulate the PCS behaviour. This is sketched in Figure 17. The construction of Proc^2 is shown in Figure 18. Note that each state of Proc^1 labelled as h is supposed to be distinct for each transition. However, all other states with the same label, i.e. h_L , h_F , q_L , q_F and q_1 , are actually representing the same state, respectively. The visual separation was done just for clarity. Note that Proc^2 is deadlock-free, which makes a dedicated Proc^3 superfluous. The set of final states is defined as $Q_F^{\mathcal{P}} := \{q_F\}$.

This construction, which will be described in more detail in the following, also works out for the variants of B-TSO where player B is allowed to update either only before or only after her move. In the former case, this is clear, since the construction never requires her to perform an update between her move and the next one of player A. In the latter case, this is due to the fact that the construction gives her enough opportunities to update during auxiliary `skip` transitions, which would not be needed in the B-TSO game with full buffer control.

In particular, consider the transition $q_1 \xrightarrow{\text{rd}(x_{w\mathbf{r}}, m)} q_m$ in Proc^2 as an example. If player B wants to take this transition while the $\langle x_{w\mathbf{r}}, m \rangle$ message is still in the buffer, she can instead move to h_1 and perform an update after her move. Player A then immediately needs to respond with going back to q_1 , since q_F is winning for player B. For the same reason,

FIGURE 17. Proc^1 of the B-TSO reduction from PCS

player A cannot transition to h_1 herself, which prevents the potential of entering an infinite loop.

Furthermore, consider q_L in Proc^2 . Note that it is immediately losing for player B, since the final state q_F is not reachable from there.

Theorem B.1. *Consider the B-TSO game or one of its variants. The set of final states S_F of \mathcal{L} is reachable from $s_0 \in S$ if and only if player B wins the game $\mathcal{G}^{\text{TSO}}(\mathcal{P}, Q_F^{\mathcal{P}})$ starting from the configuration $c_0 := \langle \langle s_0, q_1 \rangle, \langle \varepsilon, \varepsilon \rangle, \{x_{\text{wr}} \mapsto \perp, x_{\text{rd}} \mapsto \perp, y \mapsto 0\} \rangle_B \in C_B$.*

B.1. From \mathcal{L} -Reachability to a Winning Strategy in B-TSO. Proceeding as in the case of A-TSO games, suppose S_F is reachable from s_0 . Given a run $c_0^{\mathcal{L}}, \dots, c_n^{\mathcal{L}}$ in the PCS, recall the definitions of $c_k^{\mathcal{L}}$ in Appendix A. For all $k = 0, \dots, n$, we define a game configuration $c_k = \langle S_k, \mathcal{B}_k, \mathcal{M}_k \rangle_B$:

- (1) $S_k := \langle s_k, q_1 \rangle$
- (2) $\mathcal{B}_k := \langle \langle y, 1 \rangle, \langle x_{\text{wr}}, w_k[1] \rangle, \dots, \langle y, 1 \rangle, \langle x_{\text{wr}}, w_k[|w_k|] \rangle, \langle y, 1 \rangle, \varepsilon \rangle$
- (3) $\mathcal{M}_k := \{x_{\text{wr}} \mapsto \perp, x_{\text{rd}} \mapsto \perp, y \mapsto 0\}$

We describe a strategy for player B that forces the play into visiting all configurations c_1, \dots, c_n . Since $c_n \in C_F$, this means that the strategy is winning. To achieve this, we show by induction over k that starting from c_0 , Player B can force a play that visits either c_k or C_F . For $k = 0$, this is true since c_0 is the initial configuration. Let the induction hypothesis

be that the claim holds true for some arbitrary but fixed $0 \leq k < n$. For the induction step, we can assume that the game is in configuration c_k .

In the following, we assume that player A avoids immediate defeats. This does not lose generality, since otherwise the game would be in C_F and the induction would still hold.

If $op_{k+1} = \text{skip}$, then player B first takes the transition $s_k \xrightarrow{\text{skip}} h$ in Proc^1 . Since all other outgoing transitions are disabled, player A has to follow up with the transition $h \xrightarrow{\text{skip}} s_{k+1}$. As there is no change in the states of Proc^2 , the buffers, or the memory, we can conclude that the current configuration of the game is c_{k+1} .

Next, consider the case where $op_{k+1} = !m$, which is similar to the previous one. Player B takes $s_k \xrightarrow{wr(x_{wr}, m)} h$ and the only option for player A is $h \xrightarrow{wr(y, 1)} s_{k+1}$. The buffer of Proc^1 has now two additional pending write operations, which correspond to the channel message m , i.e. the last letter of w_{k+1} . Thus, the game is in configuration c_{k+1} .

The last case is where $op_{k+1} = ?m$, which is more complicated. Player B starts by updating the message $\langle x_{wr}, m \rangle$ and taking the transition $q_1 \xrightarrow{rd(x_{wr}, m)} q_m$ in Proc^2 (recall the remark about the B-TSO variants above – she might need to take a detour over h_1). Player A has to respond with $q_m \xrightarrow{wr(x_{rd}, m)} q_3$, since otherwise player B could move to q_F and win. Now, player B is in a similar situation, where she needs to leave q_3 , or else player A can move to q_L which prevents the game from ever reaching q_F . So, player B empties the buffer of Proc^2 (possibly using h_2) and takes the transition $q_3 \xrightarrow{mf} q_4$. Player A has to respond with $q_4 \xrightarrow{wr(x_{wr}, \perp)} q_5$. At this point, player B cannot further proceed in Proc^2 , since player A could respond with $q_6 \xrightarrow{\text{skip}} q_L$ and win. Thus, she has to continue in Proc^1 instead, where the only enabled transition is $s_k \xrightarrow{rd(x_{rd}, m)} h$. Now, Proc^1 is disabled and the play continues in Proc^2 . Eventually, it will go through all states q_6, \dots, q_{14} . Note that player B has to update the messages $\langle x_{wr}, \perp \rangle$, $\langle y, 1 \rangle$, $\langle y, 0 \rangle$ and $\langle x_{rd}, \perp \rangle$ along the way, and has enough opportunities to do so. Finally, player B takes the transition $q_{14} \xrightarrow{mf} q_1$ in Proc^2 and player A is forced to respond with $h \xrightarrow{rd(x_{rd}, \perp)} s_{k+1}$ in Proc^1 . Note that the last transition is enabled since the $\langle x_{rd}, \perp \rangle$ -message has certainly reached the memory.

To summarise the changes: Proc^1 transitioned from s_k to s_{k+1} , while Proc^2 is back in state q_1 . The buffer of Proc^1 has updated its two oldest messages, which correspond to the last letter in w_k , and the buffer of Proc^2 has been emptied. All variables have their default values.

We conclude that the game is in configuration c_{k+1} , which completes the proof by induction. In summary, starting from c_0 , the game either reaches the configuration $c_n \in C_F$, or another configuration in C_F . In the latter case, player B already won, but in the former case, she takes the transition $s_n \xrightarrow{wr(x_{wr}, \top)} h_F$, updates this message, and finally reaches $q_1 \xrightarrow{rd(x_{wr}, \top)} q_F \in Q_F^P$ in her next turn.

B.2. From a Winning Strategy in B-TSO to \mathcal{L} -Reachability. For the other direction, suppose that player B has a winning strategy. Consider a strategy of player A that avoids reaching $q_F \in Q^2 \cap Q_F^P$ whenever possible. Furthermore, to support our argumentation we can assume arbitrary behaviour of player A, since the strategy of player B must be winning in any case.

Let s_0, \dots, s_n be the sequence of channel states that are visited by Proc^1 during the run induced by these two strategies. Since player B uses a winning strategy, this sequence is finite. We will show by induction that for each $k = 0, \dots, n$, the run contains the game configuration $c_k = \langle \mathcal{S}_k, \mathcal{B}_k, \mathcal{M}_k \rangle_B \in C_B$, where:

- (1) $\mathcal{S}_k := \langle s_k, q_1 \rangle$
- (2) $\mathcal{B}_k := \langle \langle y, 1 \rangle, \langle x_{wr}, w_k[1] \rangle, \dots, \langle y, 1 \rangle, \langle x_{wr}, w_k[|w_k|] \rangle, \varepsilon \rangle$, where $w_k \in M^*$
- (3) $\mathcal{M}_k := \{x_{wr} \mapsto \perp, x_{rd} \mapsto \perp, y \mapsto 0\}$
- (4) If $k > 0$, there is a label op_k such that $\langle s_{k-1}, w_{k-1} \rangle \xrightarrow{op_k} \langle s_k, w_k \rangle$.

The induction base case clearly holds true, since c_0 is the initial configuration of the game. So, suppose that the claim holds for some arbitrary but fixed $0 \leq k < n$. By the induction hypothesis, we can assume that the game is in configuration c_k . Consider the different moves player B can make.

First, she may take a transition $s_k \xrightarrow{\text{skip}} h$. She surely does not update any buffer messages during that move, since this would enable player A to take $q_1 \xrightarrow{\text{rd}(x_{wr}, m)} q_L$ and win. So, we assume that player A moves along $h \xrightarrow{\text{skip}} s'$. By the construction of Proc^1 , the path $s_k \xrightarrow{\text{skip}} h \xrightarrow{\text{skip}} s'$ implies the existence of a transition $s_k \xrightarrow{\text{skip}} s'$. We define $s_{k+1} := s'$ and $w_{k+1} := w_k$ and conclude that the game is now in configuration c_{k+1} .

Another option for player A is to take the transition $s_k \xrightarrow{\text{wr}(x_{wr}, m)} h$. The argumentation continues exactly as in the previous paragraph, the only difference is that $s_k \xrightarrow{\text{!m}} s'$ and $w_{k+1} := m \bullet w_k$.

Since all other transitions in Proc^1 are disabled, the last possibility for player B to move is in Proc^2 . She first has to update a $\langle x_{wr}, m \rangle$ -message to the memory (potentially using the loop through h_1), and then she is able to perform $q_1 \xrightarrow{\text{rd}(x_{wr}, m)} q_m$. Assume that player A responds with $q_m \xrightarrow{\text{wr}(x_{rd}, m)} q_3$. Since staying in q_3 is losing for player B, she needs to update the $\langle x_{rd}, m \rangle$ -message to the memory, which enables $q_3 \xrightarrow{\text{skip}} q_4$ in Proc^2 . Player A is now forced to proceed from there and takes $q_4 \xrightarrow{\text{wr}(x_{wr}, \perp)} q_5$.

At this point, player B cannot continue in Proc^2 , since moving to q_6 opens up the opportunity for player A to immediately force her win. However, also in Proc^1 she has limited possibilities. In case she starts to simulate a skip or send operation, player A can react with $h \xrightarrow{\text{rd}(x_{rd}, m)} q_L$ and win. Since player B employs a winning strategy, we conclude that at least one transition $s_k \xrightarrow{\text{rd}(x_{rd}, m)} h$ is enabled, which player B takes. Player A then has to proceed in Proc^2 and the play continues there until reaching q_{14} .

We note the following observations, based on the fact that player A cannot force a win. First, in q_7 , the variable y apparently has value 0. This means that up to this point, no message from the buffer of Proc^1 has been updated (other than the single $\langle x_{wr}, m \rangle$ from before). Second, between q_8 and q_9 , the value of y is 1. We conclude that at least one $\langle y, 1 \rangle$ -message has been updated. Third, in q_{11} , player A reads \perp from x_{wr} . This shows that since the transition $q_6 \xrightarrow{\text{mf}} q_7$, no further $\langle x_{wr}, m' \rangle$ -message has been updated. In summary, we conclude that during the whole process, exactly the two messages $\langle x_{wr}, m \rangle$ and $\langle y, 1 \rangle$ have been updated to the memory.

We continue with the execution from q_{14} . Player B empties the buffer of Proc^2 and takes either $h \xrightarrow{\text{rd}(x_{\text{rd}}, \perp)} s'$ in Proc^1 or $q_{14} \xrightarrow{\text{mf}} q_1$ in Proc^2 . We assume that player A takes the other one, respectively.

We define $s_{k+1} := s'$ and w_{k+1} by $w_{k+1} \bullet m := w_k$. The path $s_k \xrightarrow{\text{rd}(x_{\text{rd}}, m)} h \xrightarrow{\text{rd}(x_{\text{rd}}, \perp)} s_{k+1}$ implies that $\langle s_k, w_k \rangle \xrightarrow{?m} \langle s_{k+1}, w_{k+1} \rangle$. Furthermore, the observations above imply that the game is in configuration c_{k+1} as defined in the beginning. This concludes the proof by induction.

Now, we investigate how the run can continue after c_n . Player B cannot start to simulate a skip or send operation, since we can assume that player A will move to s' , which contradicts the initial assumption that s_n is the last channel state visited in this run. The same holds true for a receive operation: Following the argumentation from above, we have already seen that player A can either force a win, or lead the game into completing the simulation of the receive operation. The only possibility left is that s_n is a final channel state and player B takes the transition $s_n \xrightarrow{\text{wr}(x_{\text{wr}}, \top)} h_F$.

In summary, we constructed a path through \mathcal{L} from $\langle s_0, \varepsilon \rangle$ to $\langle s_n, w_n \rangle$, where $s_n \in S_F$. This shows that the set of final states of \mathcal{L} is reachable.

B.3. Undecidability.

Theorem B.2. *The safety problem for the B-TSO game is undecidable.*

Proof. This follows directly from Theorem B.1 and the undecidability of the state reachability problem for perfect channel systems [BZ83]. \square

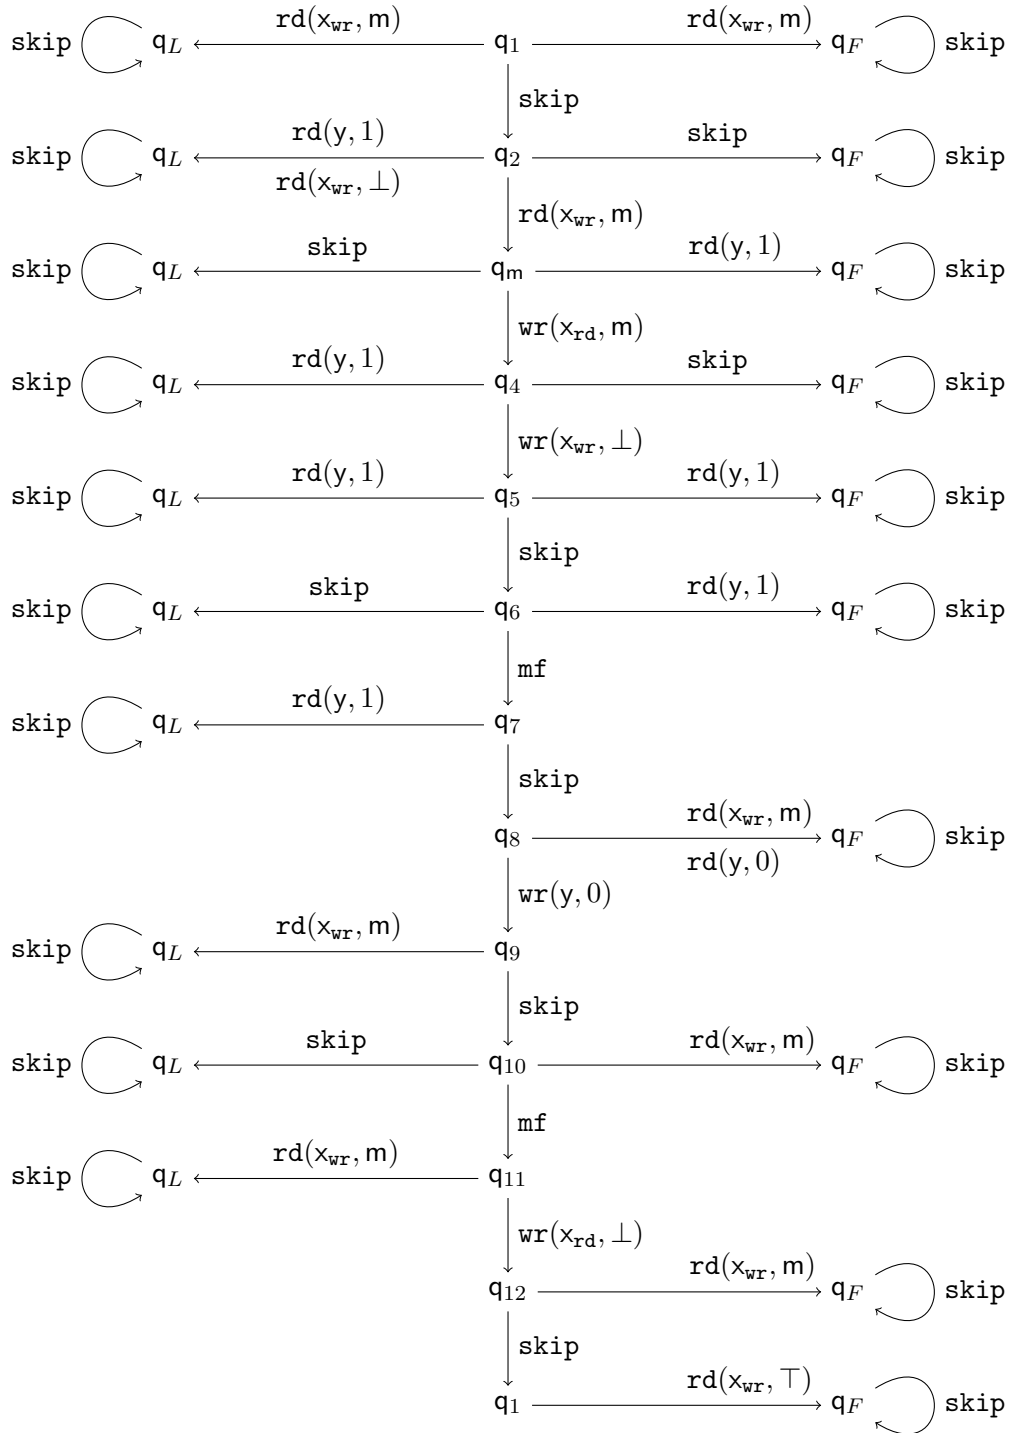
APPENDIX C. UNDECIDABILITY OF THE AB-TSO GAME

In this final section we prove undecidability of the AB-TSO game, which is the TSO game where the players share control over buffer updates in the way that both of them are allowed to update exactly after their own move. The construction will be very similar to the one for the B-TSO game. In particular, Proc^1 is exactly the same and is sketched in Figure 17. The construction for Proc^2 on the other hand is different and shown in Figure 19. Note that this process is deadlock-free by design. As previously, the set of final states is defined as $Q_F^{\mathcal{P}} := \{q_F\}$. Note that q_L in Proc^2 is immediately losing for player B, since the final state q_F is not reachable from there.

Theorem C.1. *Consider the AB-TSO game. The set of final states S_F of \mathcal{L} is reachable from $s_0 \in S$ if and only if player B wins the game $\mathcal{G}^{\text{TSO}}(\mathcal{P}, Q_F^{\mathcal{P}})$ starting from the configuration $c_0 := \langle \langle s_0, q_1 \rangle, \langle \varepsilon, \varepsilon \rangle, \{x_{\text{wr}} \mapsto \perp, x_{\text{rd}} \mapsto \perp, y \mapsto 0\} \rangle_B \in C_B$.*

C.1. From \mathcal{L} -Reachability to a Winning Strategy in AB-TSO. Proceeding as in the case of the B-TSO games, suppose S_F is reachable from s_0 . Given a run $c_0^{\mathcal{L}}, \dots, c_n^{\mathcal{L}}$ in the PCS, recall the definitions of $c_k^{\mathcal{L}}$ in Appendix A and of c_k in Appendix B.

We describe a strategy for player B that forces the play into visiting all configurations c_1, \dots, c_n . Since $c_n \in C_F$, this means that the strategy is winning. To achieve this, we show by induction over k that starting from c_0 , Player B can force a play that visits either c_k or C_F . For $k = 0$, this is true since c_0 is the initial configuration. Let the induction hypothesis

FIGURE 19. Proc^2 of the AB-TSO reduction from PCS

be that the claim holds true for some arbitrary but fixed $0 \leq k < n$. For the induction step, we can assume that the game is in configuration c_k .

In the following, we assume that player A avoids immediate defeats. This does not lose generality, since otherwise the game would be in C_F and the induction would still hold.

If $op_{k+1} = \text{skip}$, then player B first takes the transition $s_k \xrightarrow{\text{skip}} h$ in Proc^1 . For player A, there are only two enabled transitions. After $q_1 \xrightarrow{\text{skip}} q_2$, player B could move to q_F and win. Thus, player A has to follow up with the transition $h \xrightarrow{\text{skip}} s_{k+1}$ instead. Suppose she also updates one or more messages after her move. Then, player B could take $q_1 \xrightarrow{\text{rd}(x_{wr}, m')} q_F$ for some message m' and win. So, we continue assuming that player A does not update anything to the memory. As there is no change in the states of Proc^2 , the buffers or the memory, we can conclude that the current configuration of the game is c_{k+1} .

Next, consider the case where $op_{k+1} = !m$, which is similar to the previous one. Player B takes $s_k \xrightarrow{\text{wr}(x_{wr}, m)} h$ and the only option for player A is $h \xrightarrow{\text{wr}(y, 1)} s_{k+1}$. As above, we argue that she does not perform any buffer updates. The buffer of Proc^1 has now two additional pending write operations, which correspond to the channel message m , i.e. the last letter of w_{k+1} . Thus, the game is in configuration c_{k+1} .

The last case is where $op_{k+1} = ?m$, which is more complicated. Player B starts by taking the transition $q_1 \xrightarrow{\text{skip}} q_2$ in Proc^2 and then updates the message $\langle x_{wr}, m \rangle$ to the memory. Now, player A is forced to leave q_2 , since otherwise player B could move to q_F in her next turn. The only enabled transition from there is $q_2 \xrightarrow{\text{rd}(x_{wr}, m)} q_3$. Again, we see that she cannot update any messages without losing in the very next turn. Player B responds with $q_3 \xrightarrow{\text{wr}(x_{rd}, m)} q_4$, which she immediately updates to the memory. As before, player A must leave this state, the only option for this is to follow $q_4 \xrightarrow{\text{wr}(x_{wr}, \perp)} q_5$. She does so without buffer updates. Now, player B takes the transition $s_k \xrightarrow{\text{rd}(x_{rd}, m)} h$ in Proc^1 . She also empties the buffer of Proc^2 . Player A has to continue in Proc^2 , since all outgoing transitions in Proc^1 are disabled. The only option is to go to q_6 , as usual without performing any updates. Next, player B can perform $q_6 \xrightarrow{\text{mf}} q_7$, since she emptied the buffer one turn earlier. Player A continues to q_8 . If she does not perform any updates, or updates more than one message, i.e. at least one $\langle y, 1 \rangle$ and one $\langle x_{wr}, m \rangle$ buffer message (for some channel message m), then player B can immediately win in her next turn. So, we proceed assuming that player A updates exactly one message $\langle y, 1 \rangle$. Then, player B takes the transition $q_8 \xrightarrow{\text{wr}(y, 0)} q_9$ and immediately updates this message to the memory. This leaves player A with the only possibility of moving to q_{10} . As usual, she cannot update any buffer message without losing. Player B continues with $q_{10} \xrightarrow{\text{mf}} q_{11}$. Player A has to perform $q_{11} \xrightarrow{\text{wr}(x_{rd}, \perp)} q_{12}$ and may update this message, but nothing else. Next, player B moves to q_1 and empties the buffer of Proc^2 . Since player A cannot move to q_2 without losing, she takes the transition $h \xrightarrow{\text{rd}(x_{rd}, \perp)} s_{k+1}$. This is enabled since the buffer of Proc^2 was last emptied.

To summarise the changes: Proc^1 transitioned from s_k to s_{k+1} , while Proc^2 is back in state q_1 . The buffer of Proc^1 has updated its two oldest messages, which correspond to the last letter in w_k , and the buffer of Proc^2 has been emptied. All variables have their default values.

We conclude that the game is in configuration c_{k+1} , which completes the proof by induction. In summary, starting from c_0 , the game either reaches the configuration $c_n \in C_F$, or another configuration in C_F . In the latter case, player B already won, but in the former case, she takes the transition $s_n \xrightarrow{\text{wr}(x_{\text{wr}}, \top)} h_F$, updates this message, and finally reaches $q_1 \xrightarrow{\text{rd}(x_{\text{wr}}, \top)} q_F \in Q_F^{\mathcal{P}}$ in her next turn.

C.2. From a Winning Strategy in AB-TSO to \mathcal{L} -Reachability. For the other direction, suppose that player B has a winning strategy. Consider a strategy of player A that avoids reaching $q_F \in Q^2 \cap Q_F^{\mathcal{P}}$ whenever possible. Furthermore, to support our argumentation we can assume arbitrary behaviour of player A, since the strategy of player B must be winning in any case.

Let s_0, \dots, s_n be the sequence of channel states that are visited by Proc^1 during the run induced by these two strategies. Since player B uses a winning strategy, this sequence is finite. We will show by induction that for each $k = 0, \dots, n$, the run contains the game configuration $c_k = \langle \mathcal{S}_k, \mathcal{B}_k, \mathcal{M}_k \rangle_B \in C_B$, where:

- (1) $\mathcal{S}_k := \langle s_k, q_1 \rangle$
- (2) $\mathcal{B}_k := \langle \langle \langle y, 1 \rangle, \langle x_{\text{wr}}, w_k[1] \rangle, \dots, \langle y, 1 \rangle, \langle x_{\text{wr}}, w_k[|w_k|] \rangle \rangle, \varepsilon \rangle$, where $w_k \in M^*$
- (3) $\mathcal{M}_k := \{x_{\text{wr}} \mapsto \perp, x_{\text{rd}} \mapsto \perp, y \mapsto 0\}$
- (4) If $k > 0$, there is a label op_k such that $\langle s_{k-1}, w_{k-1} \rangle \xrightarrow{\text{op}_k} \langle s_k, w_k \rangle$.

The induction base case clearly holds true, since c_0 is the initial configuration of the game. So, suppose that the claim holds for some arbitrary but fixed $0 \leq k < n$. By the induction hypothesis, we can assume that the game is in configuration c_k . Consider the different moves player B can make.

First, she may take a transition $s_k \xrightarrow{\text{skip}} h$. She surely does not update any buffer messages during that move, since this would enable player A to take $q_1 \xrightarrow{\text{rd}(x_{\text{wr}}, m)} q_L$ and win. So, we assume that player A moves along $h \xrightarrow{\text{skip}} s'$. By the construction of Proc^1 , the path $s_k \xrightarrow{\text{skip}} h \xrightarrow{\text{skip}} s'$ implies the existence of a transition $s_k \xrightarrow{\text{skip}} s'$. We define $s_{k+1} := s'$ and $w_{k+1} := w_k$ and conclude that the game is now in configuration c_{k+1} .

Another option for player A is to take the transition $s_k \xrightarrow{\text{wr}(x_{\text{wr}}, m)} h$. The argumentation continues exactly as in the previous paragraph, the only difference is that $s_k \xrightarrow{!m} s'$ and $w_{k+1} := m \bullet w_k$.

Since all other transitions in Proc^1 are disabled, the last possibility for player B to move is in Proc^2 . She moves from q_1 to q_2 . We can assume that she updates exactly one message from the buffer of Proc^1 , since otherwise player A could win in her next turn. Instead, player A has to follow $q_2 \xrightarrow{\text{rd}(x_{\text{wr}}, m)} q_m$. In the following, we assume that player A performs no buffer updates unless stated otherwise. Player B takes the transition $q_m \xrightarrow{\text{wr}(x_{\text{rd}}, m)} q_4$. She might immediately update this message, but nothing else, following the same argumentation as previously. Player A continues with $q_4 \xrightarrow{\text{wr}(x_{\text{wr}}, \perp)} q_5$, we can assume that she also empties the buffer of Proc^2 . Moving to q_6 is losing for player B, thus we can safely say that she performs $s_k \xrightarrow{\text{rd}(x_{\text{rd}}, m)} h$ instead, which is the only enabled transition in Proc^1 . Player A then has to proceed in Proc^2 and the play continues there until reaching q_{11} .

We note the following observations, based on the fact that player A cannot force a win. First, in q_7 , the variable y apparently has value 0, otherwise player A would win. This means that up to this point, no message from the buffer of Proc^1 has been updated (other than the single $\langle x_{wr}, m \rangle$ from before). Second, after moving to q_8 , player A performs such a buffer update to prevent player B from immediately winning. Third, in q_{11} , the value of x_{wr} in the memory must be \perp , since otherwise player A could win with $q_{11} \xrightarrow{\text{rd}(x_{wr}, m)} q_L$ for some channel message m . This shows that since the transition $q_6 \xrightarrow{mf} q_7$, no further $\langle x_{wr}, m' \rangle$ -message has been updated. In summary, we conclude that during the whole process, exactly the two messages $\langle x_{wr}, m \rangle$ and $\langle y, 1 \rangle$ have been updated to the memory.

We continue with the execution from q_{11} . Player A performs $q_{11} \xrightarrow{wr(x_{rd}, \perp)} q_{12}$ and immediately updates this message to the memory. Player B either takes either $h \xrightarrow{\text{rd}(x_{rd}, \perp)} s'$ in Proc^1 or $q_{11} \xrightarrow{\text{skip}} q_1$ in Proc^2 . We assume that player A takes the other one, respectively.

We define $s_{k+1} := s'$ and w_{k+1} by $w_{k+1} \bullet m := w_k$. The path $s_k \xrightarrow{\text{rd}(x_{rd}, m)} h \xrightarrow{\text{rd}(x_{rd}, \perp)} s_{k+1}$ implies that $\langle s_k, w_k \rangle \xrightarrow{?m} \langle s_{k+1}, w_{k+1} \rangle$. Furthermore, the observations above imply that the game is in configuration c_{k+1} as defined in the beginning.

Now, we investigate how the run can continue after c_n . Player B cannot start to simulate a skip or send operation, since we can assume that player A will move to s' , which contradicts the initial assumption that s_n is the last channel state visited in this run. The same holds true for a receive operation: Following the argumentation from above, we have already seen that player A can either force a win, or lead the game into completing the simulation of the receive operation. The only possibility left is that s_n is a final channel state and player B takes the transition $s_n \xrightarrow{wr(x_{wr}, \top)} h_F$.

In summary, we constructed a path through \mathcal{L} from $\langle s_0, \varepsilon \rangle$ to $\langle s_n, w_n \rangle$, where $s_n \in S_F$. This shows that the set of final states of \mathcal{L} is reachable.

C.3. Undecidability.

Theorem C.2. *The safety problem for the AB-TSO game is undecidable.*

Proof. This follows directly from Theorem C.1 and the undecidability of the state reachability problem for perfect channel systems [BZ83]. \square