

Context-Bounded Analysis for POWER

Parosh Aziz Abdulla¹, Mohamed Faouzi Atig¹, Ahmed Bouajjani², and Tuan Phong Ngo¹

¹ Uppsala University, Uppsala, Sweden
{parosh, mohamed_faouzi.atig, tuan-phong.ngo}@it.uu.se
² IRIF, Université Paris Diderot, Paris, France
abou@irif.fr

Abstract. We propose an under-approximate reachability analysis algorithm for programs running under the POWER memory model, in the spirit of the work on context-bounded analysis initiated by Qadeer et al. in 2005 for detecting bugs in concurrent programs (supposed to be running under the classical SC model). To that end, we first introduce a new notion of context-bounding that is suitable for reasoning about computations under POWER, which generalizes the one defined by Atig et al. in 2011 for the TSO memory model. Then, we provide a polynomial size reduction of the context-bounded state reachability problem under POWER to the same problem under SC: Given an input concurrent program \mathcal{P} , our method produces a concurrent program \mathcal{P}' such that, for a fixed number of context switches, running \mathcal{P}' under SC yields the same set of reachable states as running \mathcal{P} under POWER. The generated program \mathcal{P}' contains the same number of processes as \mathcal{P} , and operates on the same data domain. By leveraging the standard model checker CBMC, we have implemented a prototype tool and applied it on a set of benchmarks, showing the feasibility of our approach.

1 Introduction

For performance reasons, modern multi-processors may reorder memory access operations. This is due to complex buffering and caching mechanisms that make the response memory queries (load operations) faster, and allow to speed up computations by parallelizing as much as possible independent operations and computation flows. Therefore, operations may not be visible to all processors at the same time, and they are not necessarily seen in the same order by different processors (when they concern different addresses/variables). The only model where all operations are visible immediately to all processors is the Sequential Consistency (SC) model [23] which corresponds to the standard interleaving semantics where the program order between operations of a same processor is preserved. Modern architectures adopt weaker models (in the sense that they allow more behaviours) due to the relaxation in various ways of the program order. Examples of such weak models are TSO adopted in Intel x86 machines for instance, POWER adopted in PowerPC machines, or the model adopted in ARM machines.

Apprehending the effects of all the relaxations allowed in such models is extremely hard. For instance, while TSO allows reordering stores past loads (of different addresses/variables) reflecting the use of store buffers, a model such as POWER allows reordering of all kind of store and load operations under conditions that can be quite

subtle to understand. A lot of work has been devoted to the definition of formal models that capture accurately the program semantics corresponding to models such as TSO and POWER [30, 27, 25, 29, 9]. Still, programming against weak memory models is a hard and error prone task. Therefore, developing formal verification approaches under weak memory models is of paramount importance. In particular, it is crucial in this context to have efficient algorithms for automatic bug detection. This paper addresses precisely this issue and presents an algorithmic approach for checking state reachability in concurrent programs running on the POWER semantics as defined in [18, 29] (which provides some fixes of the operational model of Power presented in [29]).

The verification of concurrent programs under weak memory models is known to be complex. Indeed, encoding the buffering and storage mechanisms used in these models leads in general to complex, infinite-state formal operational models involving unbounded data structures like FIFO queues (or more generally unbounded partial order constraints). For the case of TSO, efficient, yet precise encodings of the effects of its storage mechanism have been designed recently [4, 3]. It is not clear how to define such precise and practical encodings for POWER.

In this paper, we consider an alternative approach. We investigate the issue of defining approximate analysis. Our approach consists in introducing a parametric under-approximation schema in the spirit of context-bounding [28, 26, 22, 21, 10]. Context-bounding has been proposed in [28] as a suitable approach for efficient bug detection in multithreaded programs. Indeed, for concurrent programs, a bounding concept that provides both good coverage and scalability must be based on aspects related to the interactions between concurrent components. It has been shown experimentally that concurrency bug show up in general after a small number of context switches [26].

In the context of weak memory models, context-bounded analysis has been extended in [10] to the case of programs running on TSO. The work we present here aims at extending this approach to the case of POWER. This extension is actually very challenging due to the complexity of POWER that requires developing new techniques that are radically different from, and by far much more involved than the ones used for the case of TSO. First, we introduce a new concept of bounding that is suitable for POWER. Intuitively, the architecture of POWER is somehow similar to a distributed system with a replicated memory, where each processor has his own replica, and where operations are propagated between replicas according to some specific protocol. Our bounding concept is based on this architecture. We consider that a computation is divided in a sequence of “contexts”, where a context is a computation segment for which there is precisely one active processor such that all actions within that context are either operations issued by that processor, or propagation actions performed by its storage subsystem. Then, we consider for the analysis only computations that have a number of contexts that is less or equal than some given bound. Notice that while we bound the number of contexts in a computation, we do not put any bound on the length of the contexts, nor on the size of the storage systems.

We prove that for every bound \mathbb{K} , and for every concurrent program $Prog$, it is possible to construct, using code-to-code translation, another concurrent program $Prog^\bullet$ such that for every \mathbb{K} -bounded computation π in $Prog$ under the POWER semantics there is a corresponding \mathbb{K} -bounded computation π^\bullet of $Prog^\bullet$ under the SC seman-

tics that reaches the same set of states and vice-versa. Thus, the context-bounded state reachability problem for *Prog* can be reduced to the context-bounded state reachability problem for *Prog*[•] under SC. We show that the program *Prog*[•] has the same number of processes as *Prog*, and only $O(|\mathcal{P}||\mathcal{X}|\mathbb{K} + |\mathcal{R}|)$ additional shared variables and local registers compared to *Prog*, where $|\mathcal{P}|$ is the number of processes, $|\mathcal{X}|$ is the number of shared variables and $|\mathcal{R}|$ is the number of local registers in *Prog*. Furthermore, the obtained program has the same type of data structures and variables as the original one. As a consequence, we obtain for instance that for finite-data programs, the context-bounded analysis of programs under POWER is decidable. Moreover, our code-to-code translation allows to leverage existing verification tools for concurrent programs to safety verification under POWER.

To show the practicability of our approach, we have implemented our reduction, and we have used cbmc version 5.1 [15] as the backend tool for solving SC reachability queries. We have carried out several experiments showing the efficiency of our approach. Our experimental results confirm the assumption that concurrency bugs manifest themselves within small bounds of context switches. They also confirm that our approach based on context-bounding is more efficient and scalable than approaches based on bounding sizes of computations and/or of storage systems.

Related work. There is a lot of work on automatic program verification under weak memory models, based on precise, under-approximate, and abstract analyses, e.g., [24, 19, 20, 10, 31, 32, 16, 4, 6, 13, 14, 11, 12, 33, 2, 34, 17, 8]. While most of the existing work concern TSO, only few work address the safety verification problem under POWER [5, 8, 31, 7, 9], while [18] addresses the different issue of checking robustness against POWER, i.e., whether a program has the same (trace) semantics for both POWER and SC.

The work in [7] extends the cbmc framework by taking into account weak memory models including TSO and POWER. While their approach uses reductions to SC analysis, it is conceptually and technically totally different from ours. The work in [8] develops a verification techniques combining partial order techniques with bounded model checking, that is applicable to various weak memory models including again TSO and POWER. However, these techniques are not anymore supported by the latest version of cbmc. The work in [5] develops stateless model-checking techniques under POWER. In Section 4, we compare the performances of our approach with those of [7] and [5]. The tool herd [9] operates on small litmus tests under various memory models. Our tool can handle in an efficient and precise way such litmus tests.

Recently, Tomasco et al. presented a new verification approach [31], based on code-to-code translation, for programs running under TSO and PSO. They also discuss the extension of their approach to POWER however their tool can not handle examples under POWER. Our approach and the one adopted in [31] are orthogonal since we are using different bounding parameters: In this paper, we are bounding the number of contexts while Tomasco et al. [31] are bounding the number of writes.

2 Concurrent Programs

In this section, we first introduce some notations and definitions. Then, we present the syntax we use for *concurrent programs* and its semantics under *Power* as in [18, 29].

Preliminaries. Consider sets A and B . We use $[A \mapsto B]$ to denote the set of functions from A to B , and write $f : A \mapsto B$ to indicate that $f \in [A \mapsto B]$. We write $f(a) = \perp$ to denote that f is undefined for a . We use $f[a \leftarrow b]$ to denote the function g such that $g(a) = b$ and $g(x) = f(x)$ if $x \neq a$. We will use a function gen which, for a given set A , returns an arbitrary element $\text{gen}(A) \in A$. For integers i, j , we use $[i..j]$ to denote the set $\{i, i+1, \dots, j\}$. We use A^* to denote the set of finite words over A . For words $w_1, w_2 \in A^*$, we use $w_1 \cdot w_2$ to denote the concatenation of w_1 and w_2 .

Syntax. Fig. 1 gives the grammar for a small but general assembly-like language that we use for defining concurrent programs. A program Prog first declares a set \mathcal{X} of (shared) variables followed by the code of a set \mathcal{P} of processes. Each process p has a finite $\mathcal{R}(p)$ of (local) *registers*. We assume w.l.o.g. that the sets of registers of the different processes are disjoint, and define $\mathcal{R} := \cup_p \mathcal{R}(p)$. The code of each process $p \in \mathcal{P}$ starts by declaring a set of registers followed by a sequence of instructions.

For the sake of simplicity, we assume that the data domain of both the shared variables and registers is a single set \mathcal{D} . We assume a special element $0 \in \mathcal{D}$ which is the initial value of each shared variable or register. Each instruction i is of the form $\lambda : \mathfrak{s}$ where λ is a unique label (across all processes) and \mathfrak{s} is a statement. We define $\text{lbl}(i) := \lambda$ and $\text{stmt}(i) := \mathfrak{s}$. We define \mathcal{I}_p to be the set of instructions occurring in p , and define $\mathcal{I} := \cup_{p \in \mathcal{P}} \mathcal{I}_p$. We assume that \mathcal{I}_p contains a designated *initial* instruction i_p^{init} from which p starts its execution. A *read* instruction in a process $p \in \mathcal{P}$ has a statement of the form $\$r \leftarrow x$, where $\$r$ is a register in p and $x \in \mathcal{X}$ is a variable. A *write* instruction has a statement of the form $x \leftarrow \text{exp}$ where $x \in \mathcal{X}$ is a variable and exp is an *expression*. We will assume a set of expressions containing a set of operators applied to constants and registers, but not referring to the content of memory (i.e., the set of variables). Assume, conditional, and iterative instructions (collectively called *aci* instructions) can be explained in a similar manner. The statement `term` will cause the process to terminate its execution. We assume that `term` occurs only once in the code of a process p and that it has the label λ_p^{term} . For an expression exp , we use $\mathcal{R}(\text{exp})$ to denote the set of registers that occur in exp . For a write or an aci instruction i , we define $\mathcal{R}(i) := \mathcal{R}(\text{exp})$ where exp is the expression that occurs in $\text{stmt}(i)$.

For an instruction $i \in \mathcal{I}_p$, we define $\text{next}(i)$ to be the set of instructions that may follow i in a run of a process. Notice that this set contains two elements if i is an aci instruction (in the case of an assume instruction, we assume that if the condition evaluates to *false*, then the process moves to `term`), no element if i is a terminating instruction, and a single element otherwise. We define $\text{Tnext}(i)$ (resp. $\text{Fnext}(i)$) to be

$$\begin{aligned} \text{Prog} &::= \text{var } x^* (\text{proc } p \text{ reg } \$r^* i^*)^* \\ i &::= \lambda : \mathfrak{s} \\ \mathfrak{s} &::= \$r \leftarrow x \mid x \leftarrow \text{exp} \mid \text{assume } \text{exp} \\ &\quad \mid \text{if } \text{exp} \text{ then } i^* \text{ else } i^* \\ &\quad \mid \text{while } \text{exp} \text{ do } i^* \mid \text{term} \end{aligned}$$

Fig. 1: Syntax of concurrent programs.

the (unique) instruction to which the process execution moves in case the condition in the statement of i evaluates to *true* (resp. *false*).

Configurations. We will assume an infinite set \mathcal{E} of *events*, and will use an event to represent a single execution of an instruction in a process. A given instruction may be executed several times during a run of the program (for instance, when it is in the body of a loop). In such a case, the different executions are represented by different events. An event e is executed in several steps, namely it is *fetch*ed, *initial*ized, and then *comm*itted. Furthermore, a write event may be propagated to the other processes. A *configuration* c is a tuple $\langle \mathbb{E}, \prec, \text{ins}, \text{status}, \text{rf}, \text{Prop}, \prec_{\text{co}} \rangle$, defined as follows.

Events. $\mathbb{E} \subseteq \mathcal{E}$ is a finite set of *events*, namely the events that have been created up to the current point in the execution of the program. $\text{ins} : \mathbb{E} \mapsto \mathcal{I}$ is a function that maps an event e to the instruction $\text{ins}(e)$ that e is executing. We partition the set \mathbb{E} into disjoint sets \mathbb{E}_p , for $p \in \mathcal{P}$, where $\mathbb{E}_p := \{e \in \mathbb{E} \mid \text{ins}(e) \in \mathcal{I}_p\}$, i.e., for a process $p \in \mathcal{P}$, the set \mathbb{E}_p contains the events whose instructions belong to p . For an event $e \in \mathbb{E}_p$, we define $\text{proc}(e) := p$. We say that e is a *write* event if $\text{ins}(e)$ is a write instruction. We use \mathbb{E}^{W} to denote the set of write events. Similarly, we define the set \mathbb{E}^{R} of *read* events, and the set \mathbb{E}^{ACI} of *aci* events whose instructions are either *assume*, *conditional*, or *iterative*. We define \mathbb{E}_p^{W} , \mathbb{E}_p^{R} , and $\mathbb{E}_p^{\text{ACI}}$, to be the restrictions of the above sets to \mathbb{E}_p . For an event e where $\text{stmt}(\text{ins}(e))$ is of the form $x \leftarrow \text{exp}$ or $\$r \leftarrow x$, we define $\text{var}(e) := x$. If e is neither a read nor a write event, then $\text{var}(e) := \perp$.

Program Order. The *program-order* relation $\prec \subseteq \mathbb{E} \times \mathbb{E}$ is an irreflexive partial order that describes, for a process $p \in \mathcal{P}$, the order in which events are fetched from the code of p . We require that (i) $e_1 \not\prec e_2$ if $\text{proc}(e_1) \neq \text{proc}(e_2)$, i.e., \prec only relates events belonging to the same process, and that (ii) \prec is a total order on \mathbb{E}_p .

Status. The function $\text{status} : \mathbb{E} \mapsto \{\text{fetch}, \text{init}, \text{com}\}$ defines, for an event e , the current *status* of e , i.e., whether it has been fetched, initialized, or committed.

Propagation. The function $\text{Prop} : \mathcal{P} \times \mathcal{X} \mapsto \mathbb{E}^{\text{W}} \cup \mathcal{E}^{\text{init}}$ defines, for a process $p \in \mathcal{P}$ and variable $x \in \mathcal{X}$, the latest write event on x that has been propagated to p . Here $\mathcal{E}^{\text{init}} := \{e_x^{\text{init}} \mid x \in \mathcal{X}\}$ is a set disjoint from the set of events \mathcal{E} , and will be used to define the initial values of the variables.

Read-From. The function $\text{rf} : \mathbb{E}^{\text{R}} \mapsto \mathbb{E}^{\text{W}} \cup \mathcal{E}^{\text{init}}$ defines, for a read event $e \in \mathbb{E}^{\text{R}}$, the write event $\text{rf}(e)$ from which e gets its value.

Coherence Order. All processes share a global view about the order in which write events are propagated. This is done through the *coherence order* \prec_{co} that is a partial order on \mathbb{E}^{W} s.t. $e_1 \prec_{\text{co}} e_2$ only if $\text{var}(e_1) = \text{var}(e_2)$, i.e., it relates only events that write to identical variables. If a write event e_1 is propagated to a process before another write event e_2 and both events write to the same variable, then $e_1 \prec_{\text{co}} e_2$ holds. Furthermore, the events cannot be propagated to any other process in the reverse order. However, it might be the case that a write event is never propagated to a given process.

Dependencies. We introduce a number of dependency orders on events that we will use in the definition of the semantics. We define the *per-location program-order* $\prec_{\text{poloc}} \subseteq \mathbb{E} \times \mathbb{E}$ such that $e_1 \prec_{\text{poloc}} e_2$ if $e_1 \prec e_2$ and $\text{var}(e_1) = \text{var}(e_2)$, i.e., it is the restriction of \prec to events with identical variables. We define the *data dependency* order \prec_{data} s.t. $e_1 \prec_{\text{data}} e_2$ if (i) $e_1 \in \mathbb{E}^{\text{R}}$, i.e., e_1 is a read event; (ii) $e_2 \in \mathbb{E}^{\text{W}} \cup \mathbb{E}^{\text{ACI}}$, i.e., e_2 is either a write or an aci event; (iii) $e_1 \prec e_2$; (iv) $\text{stmt}(\text{ins}(e_1))$ is of the form $\$r \leftarrow x$; (v) $\$r \in \mathcal{R}(\text{ins}(e_2))$; and (vi) there is no $e_3 \in \mathbb{E}^{\text{R}}$ such that $e_1 \prec e_3 \prec e_2$ and $\text{stmt}(\text{ins}(e_3))$ is of the form $\$r \leftarrow y$. Intuitively, the loaded value by e_1 is used to compute the value of the expression in $\text{ins}(e_2)$. We define the *control dependency* order \prec_{ctrl} such that $e_1 \prec_{\text{ctrl}} e_2$ if $e_1 \in \mathbb{E}^{\text{ACI}}$ and $e_1 \prec e_2$.

We say that \mathfrak{c} is *committed* if $\text{status}(e) = \text{com}$ for all $e \in \mathbb{E}$. The *initial configuration* $\mathfrak{c}_{\text{init}}$ is defined by $\langle \emptyset, \emptyset, \lambda e. \perp, \lambda e. \perp, \lambda e. \perp, \lambda p. \lambda x. e_x^{\text{init}}, \emptyset \rangle$. We use \mathbb{C} to denote the set of all configurations.

Transition Relation. We define the transition relation as a relation $\rightarrow \subseteq \mathbb{C} \times \mathcal{P} \times \mathbb{C}$. For configurations $\mathfrak{c}_1, \mathfrak{c}_2 \in \mathbb{C}$ and a process $p \in \mathcal{P}$, we write $\mathfrak{c}_1 \xrightarrow{p} \mathfrak{c}_2$ to denote that $\langle \mathfrak{c}_1, p, \mathfrak{c}_2 \rangle \in \rightarrow$. Intuitively, this means that p moves from the current configuration \mathfrak{c}_1 to \mathfrak{c}_2 . The relation \rightarrow is defined through the set of inference rules shown in Fig. 2.

The rule `Fetch` chooses the next instruction to be executed in the code of a process $p \in \mathcal{P}$. This instruction should be a possible successor of the instruction that was last executed by p . To satisfy this condition, we define $\text{MaxI}(\mathfrak{c}, p)$ to be the set of instructions as follows: (i) If $\mathbb{E}_p = \emptyset$ then define $\text{MaxI}(\mathfrak{c}, p) := \{i_p^{\text{init}}\}$, i.e., the first instruction fetched by p is i_p^{init} . (ii) If $\mathbb{E}_p \neq \emptyset$, let e' be the maximal event of p (wrt. \prec) in the configuration \mathfrak{c} and then define $\text{MaxI}(\mathfrak{c}, p) := \text{next}(\text{ins}(e'))$. In other words, we consider the instruction $i' = \text{ins}(e') \in \mathcal{I}_p$, and take its possible successors. The possibility of choosing any of the (syntactically) possible successors corresponds to *speculatively* fetching statements. As seen below, whenever we commit an aci event, we check whether the made speculations are correct or not. We create a new event e , label it by $i \in \text{MaxI}(\mathfrak{c}, p)$, and make it larger than all the other events of p wrt. \prec . In such a way, we maintain the property that the order on the events of p reflects the order in which they are fetched in the current run of the program.

There are two ways in which read events get their values, namely either from *local* write events that are performed by the process itself, or from write events that are *propagated* to the process. The first case is covered by the rule `Local-Read` in which the process p initializes a read event $e \in \mathbb{E}^{\text{R}}$ on a variable (say x), where e has already been fetched. Here, the event e is made to read its value from a local write event $e' \in \mathbb{E}_p^{\text{W}}$ on x such that (i) e' has been initialized but not yet committed, and such that (ii) e' is the closest write event that precedes e in the order \prec_{poloc} . Notice that, by condition (ii), e' is unique if it exists. To formalize this, we define the *Closest Write* function $\text{CW}(\mathfrak{c}, e) := e'$ where e' is the unique event such that (i) $e' \in \mathbb{E}^{\text{W}}$, (ii) $e' \prec_{\text{poloc}} e$, and (iii) there is no event e'' such that $e'' \in \mathbb{E}^{\text{W}}$ and $e' \prec_{\text{poloc}} e'' \prec_{\text{poloc}} e$. Notice that e' may not exist, i.e., it may be the case that $\text{CW}(\mathfrak{c}, e) = \perp$. If e' exists and it has been initialized but not committed, we initialize e and update the read-from relation appropriately. On the other hand, if such an event does not exist, i.e., if there is no write event on x before e in p , or if the closest write event on x before e in p has already been

$$\begin{array}{c}
\frac{e \notin \mathbb{E}, \prec' = \prec \cup \{\langle e', e \rangle \mid e' \in \mathbb{E}_p\}, i \in \text{MaxI}(c, p)}{c \xrightarrow{p} \langle \mathbb{E} \cup \{e\}, \prec', \text{ins}[e \leftarrow i], \text{status}[e \leftarrow \text{fetch}], \text{rf}, \text{Prop}, \prec_{\text{co}} \rangle} \text{Fetch} \\
\frac{e \in \mathbb{E}_p^{\text{R}}, \text{status}(e) = \text{fetch}, \text{CW}(c, e) = e', \text{status}(e') = \text{init}}{c \xrightarrow{p} \langle \mathbb{E}, \prec, \text{ins}, \text{status}[e \leftarrow \text{init}], \text{rf}[e \leftarrow e'], \text{Prop}, \prec_{\text{co}} \rangle} \text{Local-Read} \\
\frac{e \in \mathbb{E}_p^{\text{R}}, \text{status}(e) = \text{fetch}, (\text{CW}(c, e) = \perp) \vee (\text{CW}(c, e) = e' \wedge \text{status}(e') = \text{com})}{c \xrightarrow{p} \langle \mathbb{E}, \prec, \text{ins}, \text{status}[e \leftarrow \text{init}], \text{rf}[e \leftarrow \text{Prop}(p, \text{var}(e))], \text{Prop}, \prec_{\text{co}} \rangle} \text{Prop-Read} \\
\frac{e \in \mathbb{E}_p^{\text{R}}, \text{status}(e) = \text{init}, \text{ComCnd}(c, e), \text{RdCnd}(c, e)}{c \xrightarrow{p} \langle \mathbb{E}, \prec, \text{ins}, \text{status}[e \leftarrow \text{com}], \text{rf}, \text{Prop}, \prec_{\text{co}} \rangle} \text{Com-Read} \\
\frac{e \in \mathbb{E}_p^{\text{W}}, \text{status}(e) = \text{fetch}, \text{WrInitCnd}(c, e)}{c \xrightarrow{p} \langle \mathbb{E}, \prec, \text{ins}, \text{status}[e \leftarrow \text{init}], \text{rf}, \text{Prop}, \prec_{\text{co}} \rangle} \text{Init-Write} \\
\frac{e \in \mathbb{E}_p^{\text{W}}, \text{status}(e) = \text{init}, \text{ComCnd}(c, e), \prec'_{\text{co}} = \prec_{\text{co}} \cup \{\langle e', e \rangle \mid e' \preceq_{\text{co}} \text{Prop}(p, \text{var}(e))\}}{c \xrightarrow{p} \langle \mathbb{E}, \prec, \text{ins}, \text{status}[e \leftarrow \text{com}], \text{rf}, \text{Prop}[\langle p, \text{var}(e) \rangle \leftarrow e], \prec'_{\text{co}} \rangle} \text{Com-Write} \\
\frac{q \in \mathcal{P}, e \in \mathbb{E}_p^{\text{W}}, \text{status}(e) = \text{com}, \text{Prop}(q, \text{var}(e)) \prec_{\text{co}} e, \prec'_{\text{co}} = \prec_{\text{co}} \cup \{\langle e', e \rangle \mid e' \preceq_{\text{co}} \text{Prop}(q, \text{var}(e))\}}{c \xrightarrow{p} \langle \mathbb{E}, \prec, \text{ins}, \text{status}, \text{rf}, \text{Prop}[\langle q, \text{var}(e) \rangle \leftarrow e], \prec'_{\text{co}} \rangle} \text{Prop} \\
\frac{e \in \mathbb{E}_p^{\text{ACI}}, \text{status}(e) = \text{fetch}, \text{ComCnd}(c, e), \text{ValidCnd}(c, e)}{c \xrightarrow{p} \langle \mathbb{E}, \prec, \text{ins}, \text{status}[e \leftarrow \text{com}], \text{rf}, \text{Prop}, \prec_{\text{co}} \rangle} \text{Com-ACI}
\end{array}$$

Fig. 2: Inference rules defining the relation \xrightarrow{p} where $p \in \mathcal{P}$.

committed, then we use the rule **Prop-Read** to let e fetch its value from the latest write event on x that has been propagated to p . Notice this event is the value of $\text{Prop}(p, x)$.

To commit an initialized read event $e \in \mathbb{E}_p^{\text{R}}$, we use the rule **Com-Read**. The rule can be performed if e satisfies two conditions in c . The first condition is defined as $\text{RdCnd}(c, e) := \forall e' \in \mathbb{E}^{\text{R}} : (e' \prec_{\text{po1oc}} e) \implies (\text{rf}(e') \preceq_{\text{co}} \text{rf}(e))$. It states that for any read event e' such that e' precedes e in the order \prec_{po1oc} , the write event from which e' reads its value is equal to or precedes the write event for e in the coherence order \prec_{co} . The second condition is defined by $\text{ComCnd}(c, e) := \forall e' \in \mathbb{E} : (e' \prec_{\text{data}} e) \vee (e' \prec_{\text{ctrl}} e) \vee (e' \prec_{\text{po1oc}} e) \implies (\text{status}(e') = \text{com})$. It states that all events $e' \in \mathbb{E}$ that precede e in one of the orders \prec_{data} , \prec_{ctrl} , or \prec_{po1oc} should have already been committed.

To initialize a fetched write event $e \in \mathbb{E}_p^{\text{W}}$, we use the rule **Init-Write** that requires all events that precede e in the order \prec_{data} should have been initialized. This condition is formulated as $\text{WrInitCnd}(c, e) := \forall e' \in \mathbb{E}^{\text{R}} : (e' \prec_{\text{data}} e) \implies (\text{status}(e') = \text{init})$. When a write event in a process $p \in \mathcal{P}$ is committed, it is also immediately

propagated to p itself. To maintain the coherence order, the semantics keeps the invariant that the latest write event on a variable $x \in \mathcal{X}$ that has been propagated to a process $p \in \mathcal{P}$ is the largest in coherence order among all write events on x that have been propagated to p up to now in the run. This invariant is maintained in `Com-Write` by requiring that the event e (that is being propagated) is strictly larger in coherence order than the latest write event on the same variable as e that has been propagated to p .

Write events are propagated to other processes through the rule `Prop`. A write event e on a variable x is allowed to be propagated to a process q only if it has a coherence order that is strictly larger than the coherence of any event that has been propagated to q up to now. Notice that this is given by coherence order of `Prop(q, x)` which is the latest write event on x that has been propagated to q .

When committing an aci event through the rule `Com-ACI`, we also require that we verify any potential speculation that have been made when fetching the subsequent events. We assume that we are given a function $\text{Val}(c, e)$ that takes as input an aci event e and returns the value of the expression of the conditional statement in the instruction of e when evaluated in the configuration c . The $\text{Val}(c, e)$ is only defined when all events that precede e in the order \prec_{data} should have been initialized.

To that end, we define predicate $\text{ValidCnd}(c, e) := (\exists e' \in \mathbb{E} : e \prec e' \wedge \nexists e'' \in \mathbb{E} : e \prec e'' \prec e') \implies ((\text{Val}(c, e) = \text{true} \wedge \text{ins}(e') = \text{Tnext}(\text{ins}(e))) \vee (\text{Val}(c, e) = \text{false} \wedge \text{ins}(e') = \text{Fnext}(\text{ins}(e))))$. The rule intuitively finds the event e' that was fetched immediately after e . Notice that such an event may not exist and it is unique if it exists. The predicate requires the choice of e' is consistent with the value $\text{Val}(c, e)$ of the expression of the statement in the instruction of e .

Bounded Reachability. A run π is a sequence of transitions $c_0 \xrightarrow{p_1} c_1 \xrightarrow{p_2} c_2 \cdots c_{n-1} \xrightarrow{p_n} c_n$. In such a case, we write $c_0 \xrightarrow{\pi} c_n$. We define $\text{last}(\pi) := c_n$. We define $\pi \uparrow := p_1 p_2 \cdots p_n$, i.e., it is the sequence of processes performing the transitions in π . For a sequence $\sigma = p_1 p_2 \cdots p_n \in \mathcal{P}^*$, we say that σ is a *context* if there is a process $p \in \mathcal{P}$ such that $p_i = p$ for all $i : 1 \leq i \leq n$. We say that π is *committed* (resp. *k-bounded*) if $\text{last}(\pi)$ is committed (resp. if $\pi \uparrow = \sigma_1 \cdot \sigma_2 \cdots \sigma_k$ where σ_i is a context for all $i : 1 \leq i \leq k$).

For $c \in \mathbb{C}$ and $p \in \mathcal{P}$, we define the set of *reachable labels* of the configuration c as follows. (i) If $c = c_{\text{init}}$ then $\text{lbl}(c) := \{\perp\}$, i.e. process p does not reach to any label in the initial configuration. (ii) If $c \neq c_{\text{init}}$, let e be the maximal event of p (wrt. \prec) in c . We define $\text{lbl}(c) := \{\text{lbl}(\text{ins}(e))\}$, i.e. process p reaches to the label of the maximal event e of p (wrt. \prec) in the configuration c . In the *reachability problem*, we are given a label λ and asked whether there is a committed run π and a configuration c such that $c_{\text{init}} \xrightarrow{\pi} c$ where $\lambda \in \text{lbl}(c)$. For a natural number \mathbb{K} , the *\mathbb{K} -bounded reachability problem* is defined by requiring that π in the above definition is \mathbb{K} -bounded.

3 Translation

In this section, we introduce an algorithm that reduces, for a given number \mathbb{K} , the \mathbb{K} -bounded reachability problem for POWER to the corresponding problem for SC. Given an input concurrent program Prog , the algorithm constructs an output concurrent program Prog^\bullet whose size is polynomial in Prog and \mathbb{K} , such that for each \mathbb{K} -bounded

run π in $Prog$ under the POWER semantics there is a corresponding \mathbb{K} -bounded run π^\bullet of $Prog^\bullet$ under the SC semantics that reaches the same set of process labels. Below, we first present a scheme for the translation of $Prog$, and mention some of the challenges that arise due to the POWER semantics. Then, we give a detailed description of the data structures we use in $Prog^\bullet$. Finally, we describe the codes of the processes in $Prog^\bullet$.

Scheme. Our construction is based on code-to-code translation scheme that transforms the program $Prog$ into the program $Prog^\bullet$ following the map function $\llbracket \cdot \rrbracket_{\mathbb{K}}$ given in Fig. 3. Let \mathcal{P} and \mathcal{X} be the sets of processes and (shared) variables in $Prog$. The map $\llbracket \cdot \rrbracket_{\mathbb{K}}$ replaces the variables of $Prog$ by $(|\mathcal{P}| \cdot (\mathbb{K} + 1))$ copies of the set \mathcal{X} , in addition to a finite set of *finite-data* structures (which will be formally defined in the **Data Structures** paragraph). The map function then declares two additional processes `iniProc` and `verProc` that will be used to initialize the data structures and to check the reachability problem at the end of the run of $Prog^\bullet$. The formal definition of `iniProc` (resp. `verProc`) will be given in the **Initial-izing process** (resp. **Verifier process**)

paragraph. Furthermore, the map function $\llbracket \cdot \rrbracket_{\mathbb{K}}$ transforms the code of each process $p \in \mathcal{P}$ to a corresponding process p^\bullet that will simulate the moves of p . The processes p and p^\bullet will have the same set of registers. For each instruction i appearing in the code of the process p , the map $\llbracket i \rrbracket_{\mathbb{K}}^p$ transforms it to a sequence of instructions as follows: First, it adds the code defined by `activeCnt` to check if the process p is active during the current context, then it transforms the statement s of the instruction i into a sequence of instructions following the map $\llbracket s \rrbracket_{\mathbb{K}}^p$, and finally it adds the sequence of instructions defined by `closeCnt` to guess the occurrence of a context-switch. The translation of `aci` statements keeps the same statements and adds `control` to guess the contexts when the corresponding events will be committed. The terminating statement remain identical by the map function $\llbracket \text{term} \rrbracket_{\mathbb{K}}^p$. The translations of write and read statements will be described in the **Write Instructions** and **Read Instructions** paragraphs respectively.

Challenges. There are two *aspects* of the POWER semantics (cf. Section 2) that make it difficult to simulate the run π under the SC semantics, namely *non-atomicity* and *asynchrony*. First, events are not executed atomically. In fact, an event is first fetched and initialized before it is committed. In particular, an event may be fetched in one context and be initialized and committed only in later contexts. Since there is no bound on the number of events that may be fetched in a given context, our simulation should

$$\begin{aligned}
\llbracket Prog \rrbracket_{\mathbb{K}} &\stackrel{\text{def}}{=} \text{var } \times (\text{addvars})_{\mathbb{K}}; \langle \text{iniProc} \rangle_{\mathbb{K}} \\
&\quad \langle \text{verProc} \rangle_{\mathbb{K}} (\llbracket \text{proc } p \text{ reg } \$r^* i^* \rrbracket_{\mathbb{K}})^* \\
(\text{addvars})_{\mathbb{K}} &\stackrel{\text{def}}{=} \mu(|\mathcal{P}|, |\mathcal{X}|, \mathbb{K}) \mu^{\text{init}}(|\mathcal{P}|, |\mathcal{X}|, \mathbb{K}) \\
&\quad \alpha(|\mathcal{P}|, |\mathcal{X}|, \mathbb{K}) \alpha^{\text{init}}(|\mathcal{P}|, |\mathcal{X}|, \mathbb{K}) \\
&\quad \nu(|\mathcal{P}|, |\mathcal{X}|) \text{iR}(|\mathcal{P}|, |\mathcal{X}|) \text{cR}(|\mathcal{P}|, |\mathcal{X}|) \\
&\quad \text{iW}(|\mathcal{P}|, |\mathcal{X}|) \text{cW}(|\mathcal{P}|, |\mathcal{X}|) \text{iReg}(|\mathcal{R}|) \\
&\quad \text{cReg}(|\mathcal{R}|) \text{ctrl}(|\mathcal{P}|) \text{active}(\mathbb{K}) \text{cntxt} \\
(\text{iniProc})_{\mathbb{K}} &\stackrel{\text{def}}{=} \llbracket \text{iniProc} \rrbracket_{\mathbb{K}} \\
(\text{verProc})_{\mathbb{K}} &\stackrel{\text{def}}{=} \llbracket \text{verProc} \rrbracket_{\mathbb{K}} \\
\llbracket \text{proc } p \text{ reg } \$r^* i^* \rrbracket_{\mathbb{K}} &\stackrel{\text{def}}{=} \text{proc } p \text{ reg } \$r^* (\llbracket i \rrbracket_{\mathbb{K}}^p)^* \\
\llbracket i \rrbracket_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \lambda: \langle \text{activeCnt} \rangle_{\mathbb{K}}^p \llbracket s \rrbracket_{\mathbb{K}}^p \langle \text{closeCnt} \rangle_{\mathbb{K}}^p \\
\langle \text{activeCnt} \rangle_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \text{assume}(\text{active}(\text{cntxt}) = p) \\
\langle \text{closeCnt} \rangle_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \text{cntxt} \leftarrow \text{cntxt} + \text{gen}([0..|\mathbb{K}| - 1]); \\
&\quad \text{assume}(\text{cntxt} \leq |\mathbb{K}|) \\
\llbracket \$r \leftarrow x \rrbracket_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \llbracket \$r \leftarrow x \rrbracket_{\mathbb{K}}^{\text{Read}} \\
\llbracket x \leftarrow \text{exp} \rrbracket_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \llbracket x \leftarrow \text{exp} \rrbracket_{\mathbb{K}}^{\text{Write}} \\
\llbracket \text{assume } \text{exp} \rrbracket_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \text{assume } \text{exp}; \langle \text{control} \rangle_{\mathbb{K}}^p \\
\llbracket \text{if } \text{exp} \text{ then } i^* &\stackrel{\text{def}}{=} \text{if } \text{exp} \text{ then } (\llbracket i \rrbracket_{\mathbb{K}}^p)^* \\
\text{else } i^* \rrbracket_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \text{else } (\llbracket i \rrbracket_{\mathbb{K}}^p)^*; \langle \text{control} \rangle_{\mathbb{K}}^p \\
\llbracket \text{while } \text{exp} \text{ do } i^* \rrbracket_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \text{while } \text{exp} \text{ do } (\llbracket i \rrbracket_{\mathbb{K}}^p)^*; \langle \text{control} \rangle_{\mathbb{K}}^p \\
\langle \text{control} \rangle_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \text{ctrl}(p) \leftarrow \text{ctrl}(p) + \text{gen}([0..|\mathbb{K}| - 1]); \\
&\quad \text{assume}(\text{ctrl}(p) \leq |\mathbb{K}|) \\
\llbracket \text{term} \rrbracket_{\mathbb{K}}^p &\stackrel{\text{def}}{=} \text{term}
\end{aligned}$$

Fig. 3: Translation map $\llbracket \cdot \rrbracket_{\mathbb{K}}$. We omit the label of an intermediary instruction when it is not relevant.

be able to handle unbounded numbers of pending events. Second, write events of one process are propagated in an *asynchronous* manner to the other processes. This implies that we may have unbounded numbers of “traveling” events that are committed in one context and propagated to other processes only in subsequent contexts. This creates two *challenges* in the simulation. On the one hand, we need to keep track of the coherence order among the different write events. On the other hand, since write events are not distributed to different processes at the same time, the processes may have different views of the values of a given variable at a given point of time.

Since it is not feasible to record the initializing, committing, and propagating contexts of an unbounded number of events in an SC run, our algorithm will instead predict the *summary* of effects of arbitrarily long sequences of events that may occur in a given context. This is implemented using an intricate scheme that first *guesses* and then *checks* these summaries. Concretely, each event e in π is simulated by a sequence of instructions in π^\bullet . This sequence of instructions will be executed atomically (without interruption from other processes and events). More precisely if e is fetched in a context $k : 1 \leq k \leq \mathbb{K}$, then the corresponding sequence of instructions will be executed in the same context k in π^\bullet . Furthermore, we let π^\bullet *guess (speculate)* (i) the contexts in which e will be initialized, committed, and propagated to the other processes, and (ii) the values of variables that are seen by read operations. Then, we *check* whether the guesses made by π^\bullet are valid w.r.t. the POWER semantics. As we will see below, these checks are done both on-the-fly during π^\bullet , as well as at the end of π^\bullet . To implement the guess-and-check scheme, we use a number of data structures, described below.

Data Structures. We will introduce the data structures used in our simulation in order to deal with the above asynchrony and non-atomicity challenges.

Asynchrony. In order to keep track of the coherence order, we associate a *time stamp* with each write event. A time stamp τ is a mapping $\mathcal{P} \mapsto \mathbb{K}^\otimes$ where $\mathbb{K}^\otimes := \mathbb{K} \cup \{\otimes\}$. For a process $p \in \mathcal{P}$, the value of $\tau(p)$ represents the context in which the given event is propagated to p . In particular, if $\tau(p) = \otimes$ then the event is never propagated to p . We use \mathbb{T} to denote the set of time stamps. We define an order \sqsubseteq on \mathbb{T} such that $\tau_1 \sqsubseteq \tau_2$ if, for all processes $p \in \mathcal{P}$, either $\tau_1(p) = \otimes$, or $\tau_2(p) = \otimes$, or $\tau_1(p) \leq \tau_2(p)$. Notice that if $\tau_1 \sqsubseteq \tau_2$ and there is a process $p \in \mathcal{P}$ such that $\tau_1(p) \neq \otimes$, $\tau_2(p) \neq \otimes$, and $\tau_1(p) < \tau_2(p)$ then $\tau_1(q) \leq \tau_2(q)$ whenever $\tau_1(q) \neq \otimes$ and $\tau_2(q) \neq \otimes$. In such a case, $\tau_1 \sqsubset \tau_2$. On the other hand, if either $\tau_1(p) = \otimes$ or $\tau_2(p) = \otimes$ for all $p \in \mathcal{P}$, then both $\tau_1 \sqsubseteq \tau_2$ and $\tau_2 \sqsubseteq \tau_1$. The coherence order \prec_{co} on write events will be reflected in the order \sqsubseteq on their time stamps. In particular, for events e_1 and e_2 with time stamps τ_1 and τ_2 respectively, if $\tau_1 \sqsubset \tau_2$ then e_1 precedes e_2 in coherence order. The reason is that there is at least one process p to which both e_1 and e_2 are propagated, and e_1 is propagated to p before e_2 . However, if both $\tau_1 \sqsubseteq \tau_2$ and $\tau_2 \sqsubseteq \tau_1$ then the events are never propagated to the same process, and hence they need not to be related by coherence order.

If $\tau_1 \sqsubseteq \tau_2$ then we define the *summary* of τ_1 and τ_2 , denoted $\tau_1 \oplus \tau_2$, to be the time stamp τ such that $\tau(p) = \tau_1(p)$ if $\tau_2(p) = \otimes$, and $\tau(p) = \tau_2(p)$ otherwise. For a sequence $\sigma = \tau_0 \sqsubseteq \tau_1 \sqsubseteq \dots \sqsubseteq \tau_n$ of time stamps, we define the summary $\oplus \sigma := \tau'_n$ where τ'_i is defined inductively by $\tau'_0 := \tau_0$, and $\tau'_i := \tau'_{i-1} \oplus \tau_i$ for $i : 1 \leq i \leq n$. Notice that, for $p \in \mathcal{P}$, we have $\oplus \sigma(p) = \tau_i(p)$ where i is the largest $j : 1 \leq j \leq n$ s.t. $\tau_j(p) \neq \otimes$.

Our simulation observes the sequence of write events received by a process in each context. In fact, the simulation will initially *guess* and later *verify* the summaries of the time stamps of such a sequence. This is done using data structures α^{init} and α . The mapping $\alpha^{init} : \mathcal{P} \times \mathcal{X} \times \mathbb{K} \mapsto [\mathcal{P} \mapsto \mathbb{K}^{\otimes}]$ stores, for a process $p \in \mathcal{P}$, a variable $x \in \mathcal{X}$, and a context $k : 1 \leq k \leq \mathbb{K}$, an *initial guess* $\alpha^{init}(p, x, k)$ of the summary of the time stamps of the sequence of write events on x propagated to p up to the *start* of context k . Starting from a given initial guess for a given context k , the time stamp is updated successively using the sequence of write events propagated in k . The result is stored using the mapping $\alpha : \mathcal{P} \times \mathcal{X} \times \mathbb{K} \mapsto [\mathcal{P} \mapsto \mathbb{K}^{\otimes}]$. More precisely, we initially set the value of α to α^{init} . Each time a new write event e on x is created by p in the context k , we guess the time stamp β of e , and then update $\alpha(p, x, k)$ by computing its summary with β . Thus, given a point in a context k , $\alpha(p, x, k)$ contains the summary of the time stamps of the whole sequence of write events on x that have been propagated to p up to that point. At the end of the simulation, we *verify*, for each context $k : 1 \leq k < \mathbb{K}$, that the value of α for a context k is equal to the value of α^{init} for the next context $k + 1$.

Furthermore, we use three data structures for storing the values of variables. The mapping $\mu^{init} : \mathcal{P} \times \mathcal{X} \times \mathbb{K} \mapsto \mathcal{D}$ stores, for a process $p \in \mathcal{P}$, a variable $x \in \mathcal{X}$, and a context $k : 1 \leq k \leq \mathbb{K}$, an *initial guess* $\mu^{init}(p, x, k)$ of the value of the latest write event on x propagated to p up to the *start* of context k . The mapping $\mu : \mathcal{P} \times \mathcal{X} \times \mathbb{K} \mapsto \mathcal{D}$ stores, for a process $p \in \mathcal{P}$, a variable $x \in \mathcal{X}$, and a point in a context $k : 1 \leq k \leq \mathbb{K}$, the value $\mu(p, x, k)$ of the latest write event on x that has been propagated to p up to that point. Moreover, the mapping $v : \mathcal{P} \times \mathcal{X} \mapsto \mathcal{D}$ stores, for a process $p \in \mathcal{P}$ and a variable $x \in \mathcal{X}$, the latest value $v(p, x)$ that has been written to x by p .

Non-atomicity. In order to satisfy the different dependencies between events, we need to keep track of the contexts in which they are initialized and committed. One aspect of our translation is that it only needs to keep track of the *context* in which the *latest* read or write event on a given variable in a given process is initialized or committed. The mapping $iW : \mathcal{P} \times \mathcal{X} \mapsto \mathbb{K}$ defines, for $p \in \mathcal{P}$ and $x \in \mathcal{X}$, the context $iW(p, x)$ in which the latest write event on x in p is initialized. The mapping $cW : \mathcal{P} \times \mathcal{X} \mapsto \mathbb{K}$ is defined in a similar manner for committing (rather than initializing) write events. Furthermore, we define similar mappings iR and cR for read events. The mapping $iReg : \mathcal{R} \mapsto \mathbb{K}$ gives, for a register $\$r \in \mathcal{R}$, the initializing context $iReg(\$r)$ of the latest read event loading a value to $\$r$. For an expression exp , we define $iReg(exp) := \max \{iReg(\$r) \mid \$r \in \mathcal{R}(exp)\}$. The mapping $cReg : \mathcal{R} \mapsto \mathbb{K}$ gives the context for committing (rather than initializing) of the read events. We extend $cReg$ from registers to expressions in a similar manner to $iReg$. Finally, the mapping $ctrl : \mathcal{P} \mapsto \mathbb{K}$ gives, for a process $p \in \mathcal{P}$, the committing context $ctrl(p)$ of the latest aci event in p .

Initializing Process. Alg. 1 shows the initialization process. The for-loop of lines 1, 3 and 5 define the values of the initializing and committing data structures for the variables and registers together with $v(p, x)$, $\mu(p, x, 1)$, $\alpha(p, x, 1)$ and $ctrl(p)$ for all $p \in \mathcal{P}$ and $x \in \mathcal{X}$. The for-loops of line 7 define the initial values of α and μ at the start of each context $k \geq 2$ (as described above). The for-loop of line 10 chooses an *active* process to execute in each context. The *current context* variable `cntxt` is initialized to 1.

Alg. 1: Translating $\llbracket \text{iniProc} \rrbracket_{\mathbb{K}}$.

```

1 for  $p \in \mathcal{P} \wedge x \in \mathcal{X}$  do
2    $\text{iR}(p, x) \leftarrow 1$ ;  $\text{cR}(p, x) \leftarrow 1$ ;  $\text{iW}(p, x) \leftarrow 1$ ;
    $\text{cW}(p, x) \leftarrow 1$ ;  $\nu(p, x) \leftarrow 0$ ;  $\mu(p, x, 1) \leftarrow 0$ ;
    $\alpha(p, x, 1) \leftarrow \otimes^{\mathcal{P}}$ ;
3 for  $p \in \mathcal{P}$  do
4    $\text{ctrl}(p) \leftarrow 1$ ;
5 for  $\$r \in \mathcal{R}$  do
6    $\text{iReg}(\$r) \leftarrow 1$ ;  $\text{cReg}(\$r) \leftarrow 1$ ;
7 for  $p \in \mathcal{P} \wedge x \in \mathcal{X} \wedge k \in [2..|\mathbb{K}|]$  do
8    $\alpha(p, x, k) \leftarrow \alpha^{\text{init}}(p, x, k)$ ;
9    $\mu(p, x, k) \leftarrow \mu^{\text{init}}(p, x, k)$ ;
10 for  $k \in [1..|\mathbb{K}|]$  do
11    $\text{active}(k) \leftarrow \text{gen}(\mathcal{P})$ ;
12  $\text{cntxt} \leftarrow 1$ ;

```

Alg. 2: Translating $\llbracket \$r \leftarrow x \rrbracket_{\mathbb{K}}^{\text{p,Read}}$.

```

// Guess
1 old-iR  $\leftarrow \text{iR}(p, x)$ ;
2 iReg( $\$r$ )  $\leftarrow \text{iR}(p, x) \leftarrow \text{gen}([1..|\mathbb{K}|])$ ;
3 old-cR  $\leftarrow \text{cR}(p, x)$ ;
4 cReg( $\$r$ )  $\leftarrow \text{cR}(p, x) \leftarrow \text{gen}([1..|\mathbb{K}|])$ ;
// Check
5 assume( $\text{iR}(p, x) \geq \text{cntxt}$ );
6 assume( $\text{active}(\text{iR}(p, x)) = p$ );
7 assume( $\text{iR}(p, x) \geq \text{iW}(p, x)$ );
8 assume( $\text{iR}(p, x) \geq \text{cW}(p, x) \implies$ 
    $\text{iR}(p, x) \geq \alpha(p, x, \text{old-iR}(p))$ );
9 assume( $\text{cR}(p, x) \geq \text{iR}(p, x)$ );
10 assume( $\text{active}(\text{cR}(p, x)) = p$ );
11 assume( $\text{cR}(p, x) \geq$ 
    $\max\{\text{ctrl}(p), \text{old-cR}, \text{cW}(p, x)\}$ );
// Update
12 if  $\text{iR}(p, x) < \text{cW}(p, x)$  then  $\$r \leftarrow \nu(p, x)$ ;
13 else  $\$r \leftarrow \mu(p, x, \text{iR}(p, x))$ ;

```

Alg. 3: Translating $\llbracket x \leftarrow \text{exp} \rrbracket_{\mathbb{K}}^{\text{p,Write}}$.

```

// Guess
1  $\text{iW}(p, x) \leftarrow \text{gen}([1..|\mathbb{K}|])$ ;
2 old-cW  $\leftarrow \text{cW}(p, x)$ ;
3  $\text{cW}(p, x) \leftarrow \text{gen}([1..|\mathbb{K}|])$ ;
4 for  $q \in \mathcal{P}$  do
5    $\beta(q) \leftarrow \text{gen}(\mathbb{K}^{\otimes})$ ;
// Check
6 assume( $\text{iW}(p, x) \geq \text{cntxt}$ );
7 assume( $\text{active}(\text{iW}(p, x)) = p$ );
8 assume( $\text{iW}(p, x) \geq \text{iReg}(\text{exp})$ );
9 assume( $\text{cW}(p, x) \geq \text{iW}(p, x)$ );
10 assume( $\text{cW}(p, x) \geq$ 
    $\max\{\text{cReg}(\text{exp}), \text{ctrl}(p), \text{cR}(p, x), \text{old-cW}\}$ );
11 for  $q \in \mathcal{P}$  do
12   if  $q = p$  then
13      $\beta(q) \leftarrow \text{cW}(p, x)$ ;
14   if  $q \neq p$  then
15      $\beta(q) \neq \otimes \implies \beta(q) \geq \text{cW}(p, x)$ ;
16   if  $\beta(q) \neq \otimes$  then
17     assume( $\alpha(q, x, \beta(q)) \sqsubseteq \beta$ );
18     assume( $\text{active}(\beta(q)) = p$ );
// Update
19 for  $q \in \mathcal{P}$  do
20   if  $\beta(q) \neq \otimes$  then
21      $\alpha(q, x, \beta(q)) \leftarrow \alpha(q, x, \beta(q)) \oplus \beta$ ;
22      $\mu(q, x, \beta(q)) \leftarrow \text{exp}$ ;
23  $\nu(p, x) \leftarrow \text{exp}$ ;

```

Alg. 4: Translating $\llbracket \text{verProc} \rrbracket_{\mathbb{K}}$.

```

1 for  $p \in \mathcal{P} \wedge x \in \mathcal{X} \wedge k \in [1..|\mathbb{K}| - 1]$  do
2   assume( $\alpha(p, x, k) = \alpha^{\text{init}}(p, x, k + 1)$ );
3   assume( $\mu(p, x, k) = \mu^{\text{init}}(p, x, k + 1)$ );
4 if  $\lambda$  is reachable then error;

```

Write Instructions. Consider a write instruction i in a process $p \in \mathcal{P}$ whose statement is of the form $x \leftarrow \text{exp}$. The translation of i is shown in Alg. 3. The code simulates an event \ominus executing i , by encoding the effects of the inference rules *Init-Write*, *Com-Write* and *Prop* that initialize, commit, and propagate a write event respectively. The translation consists of three parts, namely, *guessing*, *checking*, and *update*.

Guessing. We guess the initializing and committing contexts for the event \ominus , together with its time stamp. In line 1, we guess the context in which the event \ominus will be initialized, and store the guess in $\text{iW}(p, x)$. Similarly, in line 3, we guess the context in which the event \ominus will be committed, and store the guess in $\text{cW}(p, x)$ (having stored its old value in the previous line). In the for-loop of line 4, we guess a time stamp for \ominus and store it in β . This means that, for each process $q \in \mathcal{P}$, we guess the context in which the event \ominus will be propagated to q and we store this guess in $\beta(q)$.

Checking. We perform sanity checks on the guessed values in order to verify that they are consistent with the POWER semantics. Lines 6–8 perform the sanity checks for $\text{iW}(p, x)$. In lines 6–7, we verify that the initializing context of the event \ominus is not smaller

than the current context. This captures the fact that initialization happens after fetching of \mathfrak{e} . It also verifies that initialization happens in a context in which p is active. In line 8, we check whether WrInitCnd in the rule Init-Write is satisfied. To do that, we verify that the data dependency order \prec_{data} holds. More precisely, we find, for each register $\$r$ that occurs in exp , the initializing context of the latest read event loading to $\$r$. We make sure that the initializing context of \mathfrak{e} is later than the initializing contexts of all these read events. By definition, the largest of all these contexts is stored in $\text{iReg}(\text{exp})$.

Lines 9–10 perform the sanity checks for $\text{cW}(p, x)$. In line 9, we check the committing context of the event \mathfrak{e} is at least as large as its initializing context. In line 10, we check that ComCnd in the rule Com-Write is satisfied. To do that, we check that the committing context is larger than (i) the committing context of all the read events from which the registers in the expression exp fetch their values (to satisfy the data dependency order \prec_{data} , in a similar manner to that described for initialization above), (ii) the committing contexts of the latest read and write events on x in p , i.e., $\text{cR}(p, x)$ and $\text{cW}(p, x)$ (to satisfy the per-location program order \prec_{poloc}), and (iii) the committing context of the latest aci event in p , i.e., $\text{ctrl}(p)$ (to satisfy the control order \prec_{ctrl}).

The for-loop of line 11 performs three sanity checks on β . In line 12, we verify that the event \mathfrak{e} is propagated to p in the same context as the one in which it is committed. This is consistent with the rule Com-Write which requires that when a write event is committed then it is immediately propagated to the committing process. In line 14, we verify that if the event \mathfrak{e} is propagated to a process q (different from p), then the propagation takes place in a context later than or equal to the one in which \mathfrak{e} is committed. This is to be consistent with the fact that a write event is propagated to other processes only after it has been committed. In line 17, we check that guessed time stamp of the event \mathfrak{e} does not cause a violation of the coherence order \prec_{co} . To do that, we consider each process $q \in \mathcal{P}$ to which \mathfrak{e} will be propagated (i.e., $\beta(q) \neq \otimes$). The time stamp of \mathfrak{e} should be larger than the time stamp of any other write event \mathfrak{e}' on x that has been propagated to q up to the current point (since \mathfrak{e} should be larger in coherence order than \mathfrak{e}'). Notice that by construction the time stamp of the largest such event \mathfrak{e}' is currently stored in $\alpha(q, x, \beta(q))$. Moreover, in line 18, we check that the event is propagated in the contexts in which p is active.

Updating. The for-loop of line 19 uses the values guessed above for updating the global data structure α . More precisely, if the event \mathfrak{e} is propagated to a process q , i.e., $\beta(q) \neq \otimes$, then we add β to the summary of the time stamps of the sequence of write operations on x propagated to q up to the current point in context $\beta(q)$. In lines 22 and 23, we assign the value exp to $\mu(p, x, \beta(q))$ and $\nu(p, x)$ respectively. Recall that the former stores the value defined by the latest write event on x propagated to q up to the current point in context $\beta(q)$, and the latter stores the value defined by the latest write on x by p .

Read Instructions. Consider a read instruction i in a process $p \in \mathcal{P}$ whose statement is of the form $\$r \leftarrow x$. The translation of i is shown in Alg. 2. The code simulates an event \mathfrak{e} running i by encoding the three inference rules Local-Read , Prop-Read , and Com-Read . In a similar manner to a write instruction, the translation scheme for a read instruction consists of guessing, checking and update parts. Notice however that the initialization of the read event is carried out through two different inference rules.

Guessing. In line 1, we store the old value of $iR(p, x)$. In line 2, we guess the context in which the event e will be initialized, and store the guessed context both in $iR(p, x)$ and $iReg(\$r)$. Recall that the latter records the initializing context of the latest read event loading a value to $\$r$. In lines 3–4, we execute similar instructions for committing (rather than initializing).

Checking. Lines 5–9 perform the sanity checks for $iR(p, x)$. In lines 5–6, we check that the initializing context for the event e is not smaller than the current context. Line 7 makes sure that at least one of the two inference rules `Local-Read` and `Prop-Read` is satisfied, by checking that the closest write event $cW(c, e)$ (if it exists) has been initialized or committed. In line 8, we satisfy `RdCnd` in the rule `Com-Read`. Lines 9–11 perform the sanity checks for $cR(p, x)$ in a similar manner to the corresponding instructions for write events (see above).

Updating. The purpose of the update part (the if-statement of line 12) is to ensure that the correct read-from relation is defined as described by the inference rules `Local-Read` and `Prop-Read`. If $iR(p, x) < cW(p, x)$, then this means that the latest write event e_w on x by p is not committed and hence, according to `Local-Read`, the event e reads its value from that event. Recall that this value is stored in $v(p, x)$. On the other hand, if $iR(p, x) \geq cW(p, x)$ then the event e_w has been committed and hence, according to `Prop-Read`, the event e reads its value from the latest write event on x propagated to p in the context where e is initialized. We notice that this value is stored in $\mu(p, x, iR(p, x))$.

Verifier Process. The verifier process makes sure that the updated value α of the time stamp at the end of a given context k , $1 \leq k \leq \mathbb{K} - 1$, is equal to the corresponding guessed value α^{init} at the start of the next context. It also performs the corresponding test for the values written to variables (by comparing μ and μ^{init}). Finally, it checks whether we reach an error label λ or not.

4 Experimental Results

In order to evaluate the efficiency of our approach, we have implemented a context-bounded model checker for programs under POWER, called `power2sc`³. We use `cbmc` version 5.1 [15] as the backend tool. However, since the code translation is generic and the instrumentation for the different backends only differs in a few lines, backend integration is straightforward and not fundamentally limited to any underlying technology. In the following, we present the evaluation of `power2sc` on 28 C/threads benchmarks collected from `goto-instrument` [7], `nidhugg` [5], `memorax` [4], and the `SV-COMP17` benchmark suit [1]. These are widespread medium-sized benchmarks, and many state-of-the-art analysis tools for weak memory models (e.g. [20, 10, 13, 8, 34, 2]) have been trained on them. We divide our results in two sets. The first set concerns the unsafe programs while the second set concerns the safe ones. In both parts, we compare `power2sc` results to the ones obtained using `goto-instrument` and `nidhugg`, which are, to the best

³ <https://www.it.uu.se/katalog/tuang296/mguess>

Program/size	LB	①	②	③	CB	Program/size	LB	①	②	③	CB
		time	time	time				time	time	time	
bakery/76	8	226	t/o	1	3	bakery/85	8	t/o	t/o	70	3
burns/74	8	t/o	t/o	1	3	burns/79	8	t/o	t/o	1018	3
dekker/82	8	t/o	t/o	1	2	dekker/88	8	t/o	t/o	1158	2
sim dekker/69	8	12	t/o	1	2	sim dekker/73	8	209	t/o	14	2
dijkstra/82	8	t/o	t/o	5	3	dijkstra/88	8	t/o	t/o	t/o	3
szymanski/83	8	t/o	t/o	1	4	szymanski/93	8	t/o	t/o	89	4
fib_bench_0/36	-	2	1101	6	6	fib_bench_1/36	-	9	t/o	5	6
lambert/109	8	t/o	1	1	3	lambert/119	8	t/o	t/o	t/o	3
peterson/76	8	25	1056	1	3	peterson/84	8	928	t/o	7	3
peterson_3/96	8	t/o	1	3	4	peterson_3/111	8	t/o	t/o	348	4
pgsql/69	8	1079	1	1	2	pgsql/73	8	1522	2	38	2
pgsql_bnd/71	-	t/o	1	1	2	pgsql_bnd/75	-	t/o	t/o	10	2
tbar_2/75	8	16	1	1	3	tbar_2/80	8	t/o	332	29	3
tbar_3/94	8	104	1	1	3	tbar_3/103	8	t/o	t/o	138	3

(a)

(b)

Table 1: Comparing ③ power2sc with ① goto-instrument and ② nidhugg on two sets of benchmarks: (a) unsafe and (b) safe (with manually inserted synchronisations). The *LB* column indicates whether the tools were instructed to unroll loops up to a certain bound. The *CB* column gives the context bound for power2sc. The program size is the number of code lines. A *t/o* entry means that the tool failed to complete within 1800 seconds. The superior running time for each benchmark is given in bold font.

of our knowledge, the only two tools supporting C/threads programs under POWER⁴. All experiments were run on a machine equipped with a 2.4 Ghz Intel x86-32 Core2 processor and 4 GB RAM.

Table 1 shows the feasibility and competitiveness of our approach. Table 1a shows that power2sc performs well in detecting bugs compared to the other tools for most of unsafe examples. We observe that nidhugg and goto-instrument time out for several examples while power2sc manages to find all the errors using at most 6 contexts. This confirms that few context switches are sufficient to find bugs. Table 1b demonstrates that our approach is also effective when we run safe programs. power2sc manages to run most of the examples (except dijkstra and lambert) using the same context bounds as in the case of their respective unsafe examples. Observe that nidhugg and goto-instrument time out for several examples but they do not impose any bound on the number of context switches while power2sc does.

We have also tested the performance of our tool power2sc with respect to the verification of small litmus tests. Our tool power2sc manages to run successfully 913 litmus tests published in [29]. Furthermore, the output result returned by our tool power2sc perfectly matches the one returned by the tool herd [9] in all the litmus tests.

⁴ cbmc previously supported POWER [8], but has withdrawn support in later versions.

References

1. SV-COM17 benchmark suit. <https://sv-comp.sosy-lab.org/2017/benchmarks.php>, 2017.
2. Parosh Aziz Abdulla, Stavros Aronis, Mohamed Faouzi Atig, Bengt Jonsson, Carl Leonards-son, and Konstantinos F. Sagonas. Stateless model checking for TSO and PSO. In *TACAS*, volume 9035 of *LNCS*, pages 353–367. Springer, 2015.
3. Parosh Aziz Abdulla, Mohamed Faouzi Atig, Ahmed Bouajjani, and Tuan Phong Ngo. The benefits of duality in verifying concurrent programs under TSO. In *CONCUR*, volume 59 of *LIPICs*, pages 5:1–5:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
4. Parosh Aziz Abdulla, Mohamed Faouzi Atig, Yu-Fang Chen, Carl Leonardsson, and Ahmed Rezine. Counter-example guided fence insertion under TSO. In *TACAS 2012*, volume 7214 of *LNCS*, pages 204–219. Springer, 2012.
5. Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Carl Leonardsson. Stateless model checking for POWER. In *CAV*, volume 9780 of *LNCS*, pages 134–156. Springer, 2016.
6. 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 *ESOP*, volume 9032 of *LNCS*, pages 308–332. Springer, 2015.
7. J. Alglave, D. Kroening, V. Nimal, and M. Tautschnig. Software verification for weak memory via program transformation. In *ESOP*, volume 7792 of *LNCS*, pages 512–532. Springer, 2013.
8. J. Alglave, D. Kroening, and M. Tautschnig. Partial orders for efficient bounded model checking of concurrent software. In *CAV*, volume 8044 of *LNCS*, pages 141–157, 2013.
9. Jade Alglave, Luc Maranget, and Michael Tautschnig. Herding cats: Modelling, simulation, testing, and data mining for weak memory. *ACM TOPLAS*, 36(2):7:1–7:74, 2014.
10. M. F. Atig, A. Bouajjani, and G. Parlato. Getting rid of store-buffers in TSO analysis. In *CAV*, volume 6806 of *LNCS*, pages 99–115. Springer, 2011.
11. Ahmed Bouajjani, Egor Derevenetc, and Roland Meyer. Checking and enforcing robustness against TSO. In *ESOP*, volume 7792 of *LNCS*, pages 533–553. Springer, 2013.
12. S. Burckhardt, R. Alur, and M. M. K. Martin. CheckFence: checking consistency of concurrent data types on relaxed memory models. In *PLDI*, pages 12–21. ACM, 2007.
13. Sebastian Burckhardt and Madanlal Musuvathi. Effective program verification for relaxed memory models. In *CAV*, volume 5123 of *LNCS*, pages 107–120. Springer, 2008.
14. Jacob Burnim, Koushik Sen, and Christos Stergiou. Testing concurrent programs on relaxed memory models. In *ISSTA*, pages 122–132. ACM, 2011.
15. Edmund M. Clarke, Daniel Kroening, and Flavio Lerda. A tool for checking ANSI-C programs. In *TACAS*, volume 2988 of *LNCS*, pages 168–176. Springer, 2004.
16. A. Marian Dan, Y. Meshman, M. T. Vechev, and E. Yahav. Predicate abstraction for relaxed memory models. In *SAS*, volume 7935 of *LNCS*, pages 84–104. Springer, 2013.
17. Brian Demsky and Patrick Lam. Satcheck: Sat-directed stateless model checking for SC and TSO. In *OOPSLA 2015*, pages 20–36. ACM, 2015.
18. Egor Derevenetc and Roland Meyer. Robustness against Power is PSpace-complete. In *ICALP (2)*, volume 8573 of *LNCS*, pages 158–170. Springer, 2014.
19. Michael Kuperstein, Martin T. Vechev, and Eran Yahav. Automatic inference of memory fences. In *FMCAD*, pages 111–119. IEEE, 2010.
20. Michael Kuperstein, Martin T. Vechev, and Eran Yahav. Partial-coherence abstractions for relaxed memory models. In *PLDI*, pages 187–198. ACM, 2011.
21. Salvatore La Torre, P. Madhusudan, and Gennaro Parlato. Reducing context-bounded concurrent reachability to sequential reachability. In *CAV*, volume 5643 of *LNCS*, pages 477–492. Springer, 2009.

22. Akash Lal and Thomas W. Reps. Reducing concurrent analysis under a context bound to sequential analysis. *FMSD*, 35(1):73–97, 2009.
23. L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Comp.*, C-28(9), 1979.
24. Feng Liu, Nayden Nedev, Nedyalko Prasadnikov, Martin T. Vechev, and Eran Yahav. Dynamic synthesis for relaxed memory models. In *PLDI '12*, pages 429–440. ACM, 2012.
25. Sela Mador-Haim, Luc Maranget, Susmit Sarkar, Kayvan Memarian, Jade Alglave, Scott Owens, Rajeev Alur, Milo M. K. Martin, Peter Sewell, and Derek Williams. An axiomatic memory model for POWER multiprocessors. In *CAV*, volume 7358, pages 495–512. Springer, 2012.
26. Madanlal Musuvathi and Shaz Qadeer. Iterative context bounding for systematic testing of multithreaded programs. In *PLDI*, pages 446–455. ACM, 2007.
27. Scott Owens, Susmit Sarkar, and Peter Sewell. A better x86 memory model: x86-tso. In *TPHOLS*, volume 5674 of *LNCS*, pages 391–407. Springer, 2009.
28. Shaz Qadeer and Jakob Rehof. Context-bounded model checking of concurrent software. In *TACAS*, volume 3440 of *LNCS*, pages 93–107. Springer, 2005.
29. Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. Understanding POWER multiprocessors. In *PLDI*, pages 175–186. ACM, 2011.
30. P. Sewell, S. Sarkar, S. Owens, F. Z. Nardelli, and M. O. Myreen. x86-tso: A rigorous and usable programmer’s model for x86 multiprocessors. *CACM*, 53, 2010.
31. Ermenegildo Tomasco, Truc Nguyen Lam, Bernd Fischer, Salvatore La Torre, and Gennaro Parlato. Embedding weak memory models within eager sequentialization. October 2016. URL: <http://eprints.soton.ac.uk/402285/>.
32. Ermenegildo Tomasco, Truc Nguyen Lam, Omar Inverso, Bernd Fischer, Salvatore La Torre, and Gennaro Parlato. Lazy sequentialization for tso and pso via shared memory abstractions. In *Formal Methods in Computer-Aided Design (FMCAD)*, pages 193–200, 2016.
33. Y. Yang, G. Gopalakrishnan, G. Lindstrom, and K. Slind. Nemos: A framework for axiomatic and executable specifications of memory consistency models. In *IPDPS*. IEEE, 2004.
34. N. Zhang, M. Kusano, and C. Wang. Dynamic partial order reduction for relaxed memory models. In *PLDI*, pages 250–259. ACM, 2015.