

# Parameterized Model Checking Modulo Explicit Weak Memory Models

Sylvain Conchon

LRI (CNRS & Univ. Paris-Sud),  
Université Paris-Saclay, F-91405 Orsay  
Inria, Université Paris-Saclay, F-91120 Palaiseau  
sylvain.conchon@lri.fr

David Declerck

LRI (CNRS & Univ. Paris-Sud),  
Université Paris-Saclay, F-91405 Orsay  
Inria, Université Paris-Saclay, F-91120 Palaiseau  
david.declerck@u-psud.fr

Fatiha Zaïdi

LRI (CNRS & Univ. Paris-Sud),  
Université Paris-Saclay, F-91405 Orsay  
fatiha.zaidi@lri.fr

We present a modular framework for model checking parameterized array-based transition systems with explicit access operations on weak memory. Our approach extends the MCMT (Model Checking Modulo Theories) framework of Ghilardi and Ranise [10] with explicit weak memory models. We have implemented this new framework in Cubicle- $\mathcal{W}$ , an extension of the Cubicle model checker. The modular architecture of our tool allows us to change the underlying memory model seamlessly (TSO, PSO...). Our first experiments with a TSO-like memory model look promising.

## 1 Introduction

With the emerging of multi-core architectures, concurrent (or multi-threading) programming is becoming a standard practice for boosting the efficiency of an application. To be as efficient as possible, concurrent programs are designed to be run for an arbitrary number of cores. Unfortunately, in practice, the conception and programming of such parameterized programs is error-prone and hard to debug.


The situation is even worse if we consider that modern computer architectures feature weak memory models in which the different processes of a program may not perceive memory operations to happen in the same order. For instance, under the TSO memory model, write operations made by a process might not be immediately visible to all other processes, while they are instantly visible to the process that performs them. The new behaviors induced by these memory models make it hard to design concurrent programs as one has now to take into account both interleavings and reordering of events.

To help debugging such applications, one can use model checking [7, 4, 3, 9, 6, 2], an efficient formal technique used for verifying safety of parameterized concurrent programs [7]. For instance, one can use Cubicle [8], a model checker for array-based transition systems [10], a restricted class of parameterized systems where states are represented as logical formulas manipulating (unbounded) arrays indexed by process identifiers.

However, the MCMT [11] (Model Checking Modulo Theories) framework underlying Cubicle *implicitly* assumes a sequentially consistent (SC) memory model: the semantics of read and write operations is simply given by the order in which the operations are executed, and the process actually performing the operation is irrelevant (all processes share the same view of the memory).

In this paper, we propose an extension of MCMT [11] with *explicit* read and write operations for weak memories. Our weak memory reasoning is based on the axiomatic framework of Alglave *et al.* [5],

which describes the semantics of different weak memory models using *events* and *relations* over these events. More specifically, read and write operations on weak variables give rise to *events*, and according to the dependencies between these events, we build a *global-happens-before* relation (*ghb*) over these events, which orders these events in a global timeline, as depicted in the schema below.



In order to build this happens-before relation (*ghb*), we instrument the backward reachability analysis by (1) generating events for read and write operations on weak memory and (2) building on the fly a *ghb* relation on this events. The coherence of this relation is checked by an SMT solver, modulo a weak memory theory.

In the rest of the paper, we assume a TSO-like memory model[12], though this framework is modular: by changing how the *ghb* relation is built, other weak memory models can be expressed.

## 2 Preliminaries : axiomatic memory models

Our approach relies on an axiomatic model of weak memory. In this section, we give a brief presentation of this kind of models, based on the formalism of Alglave *et al.* [5]. Our presentation will be oriented towards a TSO-like model, but other models can be expressed using the same formalism.

We consider a concurrent program  $P$  composed of  $n$  processes  $i_1 \dots i_n$ , each executing a sequence of instructions  $s_1 \dots s_n$ . For simplicity, we consider instructions to be either read or write operations on weak variables.

These instructions are mapped to events, which are given a unique event identifier. Note that when an instruction is executed several times (for instance in loops), it will be given a new event identifier each time it is considered. A read event is described by a literal of the form  $Rd_\alpha(e, i)$ , where  $\alpha$  is the weak variable being read,  $e$  is the event identifier, and  $i$  is the identifier of the process performing the operation. Similarly, a write event is described by a literal of the form  $Wr_\alpha(e, i)$ , where  $\alpha$ ,  $e$  and  $i$  have the same meaning as before. The value associated to an event  $e$  on a variable  $\alpha$  is given by a function  $Val_\alpha(e)$ .

Then, depending on the properties of these events, different relations are built, and a *global-happens-before* relation (*ghb*) is derived from them. All these relations order the events using their unique identifiers. The set of events together with the different relations constitutes a candidate execution. If the *ghb* relation is a valid partial order (*i.e.* is acyclic), then the execution is considered valid.

The first relation is implied by the program's source code:

- *program order* (*po*) is a total order on all events of a process ; it orders the events in the same order as the source code

Under our TSO-like memory model, this relation allows to derive two new relations:

- *preserved program order (ppo)* is a partial order on the events of a process which represents the events that remain ordered under the weak memory semantics ; it is defined as the subset of event pairs in *po* minus the write-read pairs
- *fence* indicates which write-read pairs in *po* are separated by a fence instruction ; it allows to maintain the order between events that would otherwise not be ordered in *ppo*

The next two relations depend on the actual execution of the program:

- *coherence (co)* is a total order on all writes to the same variable ; it represents the order in which the writes are made globally visible
- *read-from (rf)* orders each write with the reads it provides its value to

These two relations allow to derive two new relations:

- *from-read (fr)* indicates reads that occur before some write becomes globally visible ; is is defined as follows:  $\forall e_1, e_2, e_2 \cdot rf(e_1, e_2) \wedge co(e_1, e_3) \rightarrow fr(e_2, e_3)$
- *external read-from (rfe)* is defined as the subset of event pairs in *rf* that belong to different processes

Finally, the *ghb* relation is defined as the transitive closure of the union of some of these relations:

$$ghb = (ppo \cup fence \cup co \cup rfe \cup fr)^+$$

This relation represents the order in which events appear to be ordered, from a global viewpoint. The key to the process of finding a feasible execution is thus to determine a *co* and *rf* relation that make the derived *ghb* relation acyclic.

Note that the axiomatic model of Alglave *et al.* specifies an auxiliary check (sequential consistency per variable), that we do not mention here. Indeed, for TSO, the exploration strategy we use (described in the next section) makes this check unnecessary.

### Extensions for atomicity

In order to use this framework with array-based transition systems, that may manipulate several different variables in a single transition, we must make some adjustments.

First, we allow several events of the same kind and by the same process to share the same event identifier. This is useful for instance, if we want to have two reads by a process to occur simultaneously, without any other event from another process interfering. Similarly, a process may write to two variables at the same time. This means that the writes will be made globally visible to the other processes simultaneously. This does not require any particular change to the model: the form of literal we use already allows this, and building the relations considers the events independently. For instance, let's consider three processes *i*, *j* and *k*, two event identifiers *e*<sub>1</sub> and *e*<sub>2</sub>, and the four events  $Wr_\alpha(e_1, i)$ ,  $Wr_\beta(e_1, i)$ ,  $Rd_\alpha(e_2, j)$  and  $Rd_\beta(e_3, k)$ . The two write events use the same identifier *e*<sub>1</sub> and belong to the same process *i*, but write to different variables. Then, we may have both  $rf(e_1, e_2)$  and  $rf(e_1, e_3)$ , even if *e*<sub>2</sub> and *e*<sub>3</sub> do not read from the same variable.

Another extension we need is that we must be able to (optionally) specify that a read followed by a write from the same process occur atomically, without any other event from another process interfering. This means that, from a global point of view, the events *appear* to happen simultaneously. For this purpose, we add a symmetric, reflexive and transitive *ghb-equal* relation, and redefine the *ghb* relation as follows:

$$ghb = (ppo \cup fence \cup co \cup rfe \cup fr \cup ghb-equal)^+ \setminus ghb-equal$$

This means that *ghb-equal* is only used to expand the transitive closure of *ghb*, however events that are *ghb-equal* are removed from the actual *ghb* relation (otherwise we wouldn't be able to tell whether *ghb* is acyclic).

### 3 Weak Memory Array-Based Transition Systems

In our approach, programs are described by parameterized transition systems, *i.e.* systems manipulating variables and process-indexed arrays using guard-action transitions. From the programmer's point of view, the notion of event is irrelevant: accesses to variables and arrays are literally understood as direct accesses. However, during our analysis, we try to build a *ghb* relation on-the-fly, hence, we also need to be able to represent events and the relations over these events. To accommodate for these different points of view, we define two different logic languages: a description language  $\mathcal{L}_{\mathcal{D}}$ , in which the events are implicit, and a language  $\mathcal{L}_{\mathcal{E}}$  that makes these events explicit. We use translation functions to translate a system from the description language  $\mathcal{L}_{\mathcal{D}}$  to the explicit language  $\mathcal{L}_{\mathcal{E}}$ . To factor out the common parts between these two languages, we define a base language  $\mathcal{L}_{\mathcal{B}}$ , and we have  $\mathcal{L}_{\mathcal{B}} \subset \mathcal{L}_{\mathcal{D}}$  and  $\mathcal{L}_{\mathcal{B}} \subset \mathcal{L}_{\mathcal{E}}$ .

#### Base language

We define the base language  $\mathcal{L}_{\mathcal{B}}$  as follows:

$$\begin{aligned} \text{const, } c &::= \text{constants} \\ \text{proc, } i, j, k &::= \text{process variables} \\ x, y, z &::= \text{regular (= non-weak) arrays} \\ \alpha, \beta, \gamma &::= \text{weak variables and arrays} \\ \text{op} &::= = | \neq | < | \leq \\ \text{term, } t &::= c | i | x[j] \\ \text{atom, } a &::= t \text{ op } t | \text{true} | \text{false} \\ \text{literal, } l &::= a | \neg a \\ \text{qf\_form, } qff &::= l | qff \wedge qff \end{aligned}$$

This language defines quantifier free formulas, which are conjunctions of literals (or their negation). A literal is either true, false, or a comparison between two terms. A term is either a constant, a process variable, or the access to a regular array cell.

#### Description language

We define the description language  $\mathcal{L}_{\mathcal{D}}$  as a superset of  $\mathcal{L}_{\mathcal{B}}$ :

$$\begin{aligned} \text{term, } t &::= \dots | \alpha | \alpha[j] | i @ \alpha | i @ \alpha[j] \\ \text{atom, } a &::= \dots | \text{fence}() \\ \text{uformula, } uf &::= \forall \vec{j} : \text{proc. } qff \\ \text{efformula, } ef &::= \exists \vec{j} : \text{proc. } qff \end{aligned}$$

The description language includes the base language, and defines new terms for accessing the weak variables.  $i @ \alpha[j]$  and  $i @ \alpha$  represent accesses to weak variables by a specific process  $i$ , while  $\alpha$  and  $\alpha[j]$  do not explicitly specify the accessing process. The context dictates which of these two forms of access must be used. The language also defines a `fence()` predicate, which is true for some process when its writes become globally visible to all other processes. Finally, we define formulas prefixed by an existentially or universally quantified process variable.

## Explicit language

We define the explicit language  $\mathcal{L}_{\mathcal{E}}$  as a superset of  $\mathcal{L}_{\mathcal{B}}$ :

$$\begin{aligned}
 & eid, e ::= \text{event identifiers} \\
 & term, t ::= \dots \mid \text{Val}_{\alpha}(e) \mid \text{Val}_{\alpha}(e, j) \\
 & atom, a ::= \dots \mid \text{Rd}_{\alpha}(e, i) \mid \text{Rd}_{\alpha}(e, i, j) \\
 & \quad \quad \quad \mid \text{Wr}_{\alpha}(e, i) \mid \text{Wr}_{\alpha}(e, i, j) \\
 & \quad \quad \quad \mid \text{fence}(e, i) \mid \text{ghb}(e, e) \mid \text{ghb-equal}(e, e) \\
 & qf\_form, qff ::= \dots \mid qff \vee qff \\
 & formula, f ::= qff \mid \forall \vec{x} : \text{type}. f \mid \exists \vec{x} : \text{type}. f
 \end{aligned}$$

The new  $\text{Rd}_{\alpha}$  terms allow to represent the read events, while the  $\text{Wr}_{\alpha}$  terms represent the write events. The  $\text{Val}_{\alpha}$  terms specify the value associated to events.  $\text{ghb}(e, e)$  and  $\text{ghb-equal}(e, e)$  allow to encode the *ghb* and *ghb-equal* relations. The fence predicate indicates that there is a fence before some event  $e$  (it is not to be confused with the fence literal in  $\mathcal{L}_{\mathcal{G}}$ , although the two are related). We also allow for more general formulas by adding disjunctions and quantification over types other than *proc*.

Note that  $\text{Val}_{\alpha}$ ,  $\text{Rd}_{\alpha}$  and  $\text{Wr}_{\alpha}$  are defined for every weak variable or array  $\alpha$ . Also,  $\text{Val}_{\alpha}$  may be considered as regular Cubicle arrays.

## Convenience notations

For simplicity, we use the following notations to represent the different sets of variables that we often use:

- $X$ : the set of all regular (*i.e.* non-weak) arrays  $(x, y, \dots)$
- $\hat{X}$ : the set of all regular arrays  $(x, y, \dots)$  and event values  $(\text{Val}_{\alpha}, \text{Val}_{\beta}, \dots)$
- $A^0$ : the set of all weak variables
- $A^1$ : the set of all weak arrays
- $A$ : the set of all weak variables and arrays  $(A^0 \cup A^1)$
- $A_t^0, A_t^1$  and  $A_t$ : similar to the  $A^0, A^1$  and  $A$  sets, but restricted to the variables and arrays manipulated by a transition  $t$

From the programmer's point of view, only  $X$  and the different  $A$ 's are relevant. However, the translated formulas in  $\mathcal{L}_{\mathcal{E}}$  mainly use  $\hat{X}$ , since they explicitly manipulate events and their values.

We also often use the two following notations as shortcuts to represent common expressions:

$$\Delta(\vec{e}) = \bigwedge_{(e_1, e_2) \in \vec{e}} e_1 \neq e_2 \quad \textit{i.e. all elements of } \vec{e} \textit{ are different}$$

$$\diamond(\vec{e}) = \bigwedge_{(e_1, e_2) \in \vec{e}} \text{ghb-equal}(e_1, e_2) \quad \textit{i.e. all elements of } \vec{e} \textit{ are ghb-equal}$$

## Interpretation

The explicit language  $\mathcal{L}_\varepsilon$  is to be interpreted as follows:

$$\begin{array}{ll}
\mathcal{M}[c] & = \mathcal{M}(c) \\
\mathcal{M}[i] & = \mathcal{M}(i) \\
\mathcal{M}[e] & = \mathcal{M}(e) \\
\mathcal{M}[X[j]] & = X^{\mathcal{M}}(\mathcal{M}[j]) \\
\mathcal{M}[\text{Val}_\alpha(e)] & = \text{Val}_\alpha^{\mathcal{M}}(\mathcal{M}[e]) \\
\mathcal{M}[\text{Val}_\alpha(e, j)] & = \text{Val}_\alpha^{\mathcal{M}}(\mathcal{M}[e], \mathcal{M}[j]) \\
\mathcal{M} \models t_1 \text{ op } t_2 & = \mathcal{M}[t_1] \text{ op } \mathcal{M}[t_2] \\
\mathcal{M} \models \text{ghb}(e_1, e_2) & = (\mathcal{M}[e_1], \mathcal{M}[e_2]) \in \text{ghb}^{\mathcal{M}} \\
\mathcal{M} \models \text{ghb-equal}(e_1, e_2) & = (\mathcal{M}[e_1], \mathcal{M}[e_2]) \in \text{ghb-equal}^{\mathcal{M}} \\
\mathcal{M} \models \text{Rd}_\alpha(e, i) & = (\mathcal{M}[e], \mathcal{M}[i]) \in \text{Rd}_\alpha^{\mathcal{M}} \\
\mathcal{M} \models \text{Rd}_\alpha(e, i, j) & = (\mathcal{M}[e], \mathcal{M}[i], \mathcal{M}[j]) \in \text{Rd}_\alpha^{\mathcal{M}} \\
\mathcal{M} \models \text{Wr}_\alpha(e, i) & = (\mathcal{M}[e], \mathcal{M}[i]) \in \text{Wr}_\alpha^{\mathcal{M}} \\
\mathcal{M} \models \text{Wr}_\alpha(e, i, j) & = (\mathcal{M}[e], \mathcal{M}[i], \mathcal{M}[j]) \in \text{Wr}_\alpha^{\mathcal{M}} \\
\mathcal{M} \models \text{fence}(e, i) & = (\mathcal{M}[e], \mathcal{M}[i]) \in \text{fence}^{\mathcal{M}} \\
\mathcal{M} \models \neg a & = \mathcal{M} \not\models a \\
\mathcal{M} \models \text{qff}_1 \wedge \text{qff}_2 & = \mathcal{M} \models \text{qff}_1 \text{ and } \mathcal{M} \models \text{qff}_2 \\
\mathcal{M} \models \text{qff}_1 \vee \text{qff}_2 & = \mathcal{M} \models \text{qff}_1 \text{ or } \mathcal{M} \models \text{qff}_2 \\
\mathcal{M} \models \forall x : \text{type}. f & = \mathcal{M}\{x \mapsto v\} \models f & \text{for all } v \in \mathcal{D}^{\text{type}} \\
\mathcal{M} \models \exists x : \text{type}. f & = \mathcal{M}\{x \mapsto v\} \models f & \text{for some } v \in \mathcal{D}^{\text{type}}
\end{array}$$

The domain of the model is partitioned according to the types *proc*, representing the process identifiers, *eid*, representing the event identifiers, and *val*, for the different values of variables and arrays. We have  $\mathcal{D}_{\mathcal{M}} = \mathcal{D}^{\text{proc}} \uplus \mathcal{D}^{\text{eid}} \uplus \mathcal{D}^{\text{val}}$ .

We also define  $\mathcal{D}^{\text{ur}} \subseteq \mathcal{D}^{\text{eid}}$  the subset of event identifiers that contains only unsatisfied read events, *i.e.* read events that are not connected to any write, and we have:

$$\begin{aligned}
& \forall e_r \in \mathcal{D}^{\text{ur}}. \text{unsat\_read}(e_r) \\
\text{unsat\_read}(e_r) & = \forall i, j, k \in \mathcal{D}^{\text{proc}}, \forall e_w \in \mathcal{D}^{\text{eid}}. e_r = e_w \vee \\
& \left( \left( \bigvee_{\alpha \in A^0} (e_r, i) \in \text{Rd}_\alpha \wedge (e_w, j) \in \text{Wr}_\alpha \right) \vee \right. \\
& \left. \left( \bigvee_{\alpha \in A^1} (e_r, i, k) \in \text{Rd}_\alpha \wedge (e_w, j, k) \in \text{Wr}_\alpha \right) \rightarrow (e_r, e_w) \in \text{ghb} \right)
\end{aligned}$$

## Initial state

The initial state describes the constraints on regular arrays and weak variables and arrays. It is described by a formula in  $\mathcal{L}_{\mathcal{D}}$ , parameterized by a universally quantified process variable. Accesses to weak variables must use the  $\alpha$  or  $\alpha[j]$  form ( $i @ \alpha$  and  $i @ \alpha[j]$  are not allowed). Also, the fence() literal may not be used.

We consider an initial state formula  $I$ , parameterized by a set of regular variables  $X$  and written as follows in the description language  $\mathcal{L}_{\mathcal{D}}$ :

$$I(X) = \forall j : \text{proc}. \mathcal{I}(j, X)$$

To obtain the equivalent formula  $\tilde{I}$  in  $\mathcal{L}_{\mathcal{E}}$ , we apply the transformation function  $\llbracket \cdot \rrbracket_I$ :

$$\tilde{I}(\hat{X}) = \llbracket I(X) \rrbracket_I = \forall i, j : \text{proc}, \forall \vec{e}_\alpha : \text{ur}. \bigwedge_{\alpha \in A^0} \text{Rd}_\alpha(e_\alpha, i) \wedge \bigwedge_{\alpha \in A^1} \text{Rd}_\alpha(e_\alpha, i, j) \wedge \llbracket \mathcal{I}(j, X) \rrbracket_I$$

This function generates a read event for every weak variable or array of the system, not only those actually used in  $I$  (however, only those will also have a value associated to). We have one event identifier  $e_\alpha$  per weak variable or array  $\alpha$ . The event identifiers are chosen in the domain  $\mathcal{D}^{\text{ur}}$ , which restricts the events to the reads that must take their value from the initial state. The process performing the operation is represented by the universally quantified variable  $i$ . This means that, in the initial state, any process reading a weak variable or array will obtain the same value. Note that since the resulting formula makes use of event value terms ( $\text{Val}_\alpha$ ), the formula  $\tilde{I}$  is parameterized by the set  $\hat{X}$  (while  $I$  was parameterized by  $X$ ).

The transformation function  $\llbracket \cdot \rrbracket_I : \mathcal{L}_{\mathcal{G}} \rightarrow \mathcal{L}_{\mathcal{E}}$  is defined as follows:

$$\begin{aligned} \llbracket \text{aff}_1 \wedge \text{aff}_2 \rrbracket_I &= \llbracket \text{aff}_1 \rrbracket_I \wedge \llbracket \text{aff}_2 \rrbracket_I \\ \llbracket \neg a \rrbracket_I &= \neg \llbracket a \rrbracket_I \\ \llbracket \text{true} \rrbracket_I &= \text{true} \\ \llbracket \text{false} \rrbracket_I &= \text{false} \\ \llbracket t_1 \text{ op } t_2 \rrbracket_I &= \llbracket t_1 \rrbracket_I \text{ op } \llbracket t_2 \rrbracket_I \\ \llbracket \alpha \rrbracket_I &= \text{Val}_\alpha(e_\alpha) \\ \llbracket \alpha[j] \rrbracket_I &= \text{Val}_\alpha(e_\alpha, j) \\ \llbracket t \rrbracket_I &= t \quad \text{when } t \neq \alpha \text{ and } t \neq \alpha[j] \end{aligned}$$

## States

States represent the constraints on regular arrays, events, events values and relations. They are not meant to be manipulated directly by the programmer, so they are expressed in  $\mathcal{L}_{\mathcal{E}}$ , and are parameterized by a set  $\hat{X}$ .

A state is represented by a formula of the following form:

$$\varphi(\hat{X}) = \exists \vec{e} : \text{eid}. \Delta(\vec{e}) \wedge \Phi(\vec{e}, \hat{X})$$

This means that a state uses a set of event identifiers that are all different. The second part of the formula,  $\Phi$ , is of the following form:

$$\Phi(\vec{e}, \hat{X}) = \exists \vec{j} : \text{proc}. \Delta(\vec{j}) \wedge \phi(\vec{e}, \vec{j}, \hat{X})$$

Similarly to events, the state uses a set of process identifiers that are all different.  $\phi(\vec{e}, \vec{j}, \hat{X})$  is a conjunction of literals that actually describes the constraints.

## Bad states

Bad states allow to describe the dangerous states of the system in terms of constraints on regular arrays and weak variables and arrays. They are described by a formula in  $\mathcal{L}_{\mathcal{G}}$ , parameterized by a set of regular variables  $X$ . They make use of a set of existentially quantified process variables. Contrary to the initial state, different processes may have different views of some weak variable, so accesses to weak variables must use the  $i @ \alpha$  and  $i @ \alpha[j]$  form ( $\alpha$  and  $\alpha[j]$  are not allowed). Also, fence() may not be used.

We consider a bad state formula  $\Theta$ , written as follows in the description language  $\mathcal{L}_{\mathcal{Q}}$ :

$$\Theta(X) = \exists \vec{j} : \text{proc.} \Delta(\vec{j}) \wedge \vartheta(\vec{j}, X)$$

To obtain the equivalent formula  $\tilde{\Theta}$  in  $\mathcal{L}_{\mathcal{E}}$ , we apply the transformation function  $\llbracket \cdot \rrbracket_u$ :

$$\tilde{\Theta}(\hat{X}) = \llbracket \Theta(X) \rrbracket_u = \exists \vec{e} : \text{eid.} \Delta(\vec{e}) \wedge \diamond(\vec{e}) \wedge \theta(\vec{e}, X)$$

Where:

$$\theta(\vec{e}, X) = \exists \vec{j} : \text{proc.} \Delta(\vec{j}) \wedge \llbracket \vartheta(\vec{j}, X) \rrbracket_u^{\vec{e}}$$

The translation function transforms every weak variable or array access to a read event, and generates one event identifier  $e_i$  per process.

The translation function  $\llbracket \cdot \rrbracket_u : \mathcal{L}_{\mathcal{Q}} \times \text{eid} \rightarrow \mathcal{L}_{\mathcal{E}}$  is defined as follows:

$$\begin{aligned} \llbracket \text{aff}_1 \wedge \text{aff}_2 \rrbracket_u^{\vec{e}} &= \llbracket \text{aff}_1 \rrbracket_u^{\vec{e}} \wedge \llbracket \text{aff}_2 \rrbracket_u^{\vec{e}} \\ \llbracket a \rrbracket_u^{\vec{e}} &= \llbracket a \rrbracket_{u_E}^{\vec{e}} \wedge \llbracket a \rrbracket_{u_V}^{\vec{e}} \\ \llbracket \neg a \rrbracket_u^{\vec{e}} &= \llbracket a \rrbracket_{u_E}^{\vec{e}} \wedge \neg \llbracket a \rrbracket_{u_V}^{\vec{e}} \\ \llbracket t_1 \text{ op } t_2 \rrbracket_{u_E}^{\vec{e}} &= \llbracket t_1 \rrbracket_{u_E}^{\vec{e}} \wedge \llbracket t_2 \rrbracket_{u_E}^{\vec{e}} \\ \llbracket t_1 \text{ op } t_2 \rrbracket_{u_V}^{\vec{e}} &= \llbracket t_1 \rrbracket_{u_V}^{\vec{e}} \text{ op } \llbracket t_2 \rrbracket_{u_V}^{\vec{e}} \\ \llbracket i @ \alpha \rrbracket_{u_E}^{\vec{e}} &= \text{Rd}_{\alpha}(e_i, i) && e_i \in \vec{e} \\ \llbracket i @ \alpha \rrbracket_{u_V}^{\vec{e}} &= \text{Val}_{\alpha}(e_i) && e_i \in \vec{e} \\ \llbracket i @ \alpha[j] \rrbracket_{u_E}^{\vec{e}} &= \text{Rd}_{\alpha}(e_i, i, j) && e_i \in \vec{e} \\ \llbracket i @ \alpha[j] \rrbracket_{u_V}^{\vec{e}} &= \text{Val}_{\alpha}(e_i, j) && e_i \in \vec{e} \\ \llbracket t \rrbracket_{u_E}^{\vec{e}} &= \text{true} && \text{when } t \neq i @ \alpha \text{ and } t \neq i @ \alpha[j] \\ \llbracket t \rrbracket_{u_V}^{\vec{e}} &= t && \text{when } t \neq i @ \alpha \text{ and } t \neq i @ \alpha[j] \end{aligned}$$

Note that the function  $\llbracket \cdot \rrbracket_u$  makes use of two sub-functions  $\llbracket \cdot \rrbracket_{u_E}$  and  $\llbracket \cdot \rrbracket_{u_V}$ . The first one is used to build the literals describing the events, while the second one is used to build the event value terms.

## Transitions

Transitions describe the changes made to regular arrays and weak variables and arrays. They are composed of a guard, representing the conditions that must be met for the transition to be executed, and actions, which may be updates of regular arrays or and updates of weak variables and arrays. They are expressed in the description language  $\mathcal{L}_{\mathcal{Q}}$ , with the restriction that accesses two weak variables must use the  $\alpha$  or  $\alpha[j]$  form ( $i @ \alpha$  and  $i @ \alpha[j]$  are not allowed). Also, the fence() predicate may only be used in the guard.

We consider a transition  $t$ , written as follows in the description language  $\mathcal{L}_{\mathcal{Q}}$ :

$$\begin{aligned} t(X, X') &= \exists i, \vec{j} : \text{proc.} \Delta(\vec{j}) \wedge \gamma(i, \vec{j}, X) \wedge \\ &\quad \bigwedge_{x \in X} x'[i] = \delta_x(i, \vec{j}, X) \wedge \\ &\quad \bigwedge_{\alpha \in A_i^0} \alpha' = \delta_{\alpha}(i, \vec{j}, X) \wedge \\ &\quad \bigwedge_{\alpha \in A_i^1} \bigwedge_{k \in \vec{j}} \alpha'[k] = \delta_{\alpha}(i, \vec{j}, X) \end{aligned}$$

The existentially quantified process  $i$  is the process performing the actions. This also means non-weak array terms are restricted to  $x[i]$ . Note that process  $i$  may be equal to some process in  $\vec{j}$ .



To obtain the equivalent transition  $\tilde{t}$  in  $\mathcal{L}_{\mathcal{E}}$ , we apply the transformation function  $\llbracket \cdot \rrbracket_t$ , which must be given two fresh event identifiers  $e_r$  and  $e_w$ :

$$\tilde{t}(\hat{X}, \hat{X}', e_r, e_w) = \llbracket t(X, X') \rrbracket_t^{e_r, e_w}$$

Where:

$$\begin{aligned} \llbracket t(X, X') \rrbracket_t^{e_r, e_w} &= \exists i, \vec{j} : \text{proc.} \Delta(\vec{j}) \wedge e_r \neq e_w \wedge \text{ghb-equal}(e_r, e_w) \wedge \llbracket \gamma(i, \vec{j}, X) \rrbracket_{\gamma}^{i, e_r} \wedge \\ &\quad \bigwedge_{x \in X} \left( \llbracket x'[i] = \delta_x(i, \vec{j}, X) \rrbracket_{\gamma}^{i, e_r} \wedge (\forall k. k = i \vee x'[k] = x[k]) \right) \wedge \\ &\quad \bigwedge_{\alpha \in A^0} \text{Wr}_{\alpha}(e_w, i) \wedge \llbracket \text{Val}'_{\alpha}(e_w) = \delta_{\alpha}(i, \vec{j}, X) \rrbracket_{\gamma}^{i, e_r} \wedge \\ &\quad \bigwedge_{\alpha \in A^1} \bigwedge_{k \in \vec{j}} \text{Wr}_{\alpha}(e_w, i, k) \wedge \llbracket \text{Val}'_{\alpha}(e_w, k) = \delta_{\alpha}(i, \vec{j}, X) \rrbracket_{\gamma}^{i, e_r} \end{aligned}$$

This translation ensures the two event identifiers  $e_r$  and  $e_w$  are different, and links them in the *ghb-equal* relation. The regular array updates are extended so that all array cells different from  $i$  receive a value equals to the previous one. The weak variable and array updates generate write events. The translation of the guard and updates is further delegated to the function  $\llbracket \cdot \rrbracket_{\gamma} : \mathcal{L}_{\mathcal{D}} \times \text{proc} \times \text{eid} \rightarrow \mathcal{L}_{\mathcal{E}}$ .

$$\begin{aligned} \llbracket \text{aff}_1 \wedge \text{aff}_2 \rrbracket_{\gamma}^{i, e_r} &= \llbracket \text{aff}_1 \rrbracket_{\gamma}^{i, e_r} \wedge \llbracket \text{aff}_2 \rrbracket_{\gamma}^{i, e_r} \\ \llbracket a \rrbracket_{\gamma}^{i, e_r} &= \llbracket a \rrbracket_{\gamma_E}^{i, e_r} \wedge \llbracket a \rrbracket_{\gamma_W}^{i, e_r} \\ \llbracket \neg a \rrbracket_{\gamma}^{i, e_r} &= \llbracket a \rrbracket_{\gamma_E}^{i, e_r} \wedge \neg \llbracket a \rrbracket_{\gamma_W}^{i, e_r} \\ \llbracket \text{fence}() \rrbracket_{\gamma_E}^{i, e_r} &= \text{fence}(e_r, i) \\ \llbracket \text{fence}() \rrbracket_{\gamma_W}^{i, e_r} &= \text{true} \\ \llbracket t_1 \text{ op } t_2 \rrbracket_{\gamma_E}^{i, e_r} &= \llbracket t_1 \rrbracket_{\gamma_E}^{i, e_r} \wedge \llbracket t_2 \rrbracket_{\gamma_E}^{i, e_r} \\ \llbracket t_1 \text{ op } t_2 \rrbracket_{\gamma_W}^{i, e_r} &= \llbracket t_1 \rrbracket_{\gamma_W}^{i, e_r} \text{ op } \llbracket t_2 \rrbracket_{\gamma_W}^{i, e_r} \\ \llbracket \alpha \rrbracket_{\gamma_E}^{i, e_r} &= \text{Rd}_{\alpha}(e_r, i) \\ \llbracket \alpha \rrbracket_{\gamma_W}^{i, e_r} &= \text{Val}_{\alpha}(e_r, i) \\ \llbracket \alpha[j] \rrbracket_{\gamma_E}^{i, e_r} &= \text{Rd}_{\alpha}(e_r, i, j) \\ \llbracket \alpha[j] \rrbracket_{\gamma_W}^{i, e_r} &= \text{Val}_{\alpha}(e_r, i) \\ \llbracket t \rrbracket_{\gamma_E}^{i, e_r} &= \text{true} && \text{when } t \neq \alpha \text{ and } t \neq \alpha[j] \\ \llbracket t \rrbracket_{\gamma_W}^{i, e_r} &= t && \text{when } t \neq \alpha \text{ and } t \neq \alpha[j] \end{aligned}$$

## 4 Backward reachability

Our approach relies on a rather classical backward reachability algorithm (function BWD below), whose objective is to check whether there is a possible path from the initial state to the dangerous states. However, the pre-image computation has been extended to produce events and relations, according to our weak memory semantics. The algorithm takes as input a transition system  $S = (Q, I, \tau)$  and a cube  $\Theta$ , where  $I$  is a formula describing the initial states of the system,  $\tau$  is the set of all transitions, and  $\Theta$  a cube describing the dangerous states. It maintains a set of visited states  $\mathcal{V}$  and a working queue of cubes  $\mathcal{Q}$ .

```

1 function BWD( $\mathcal{S}, \Theta$ ) : begin
2    $\mathcal{V} := \emptyset$ ;
3    $push(\mathcal{Q}, \llbracket \Theta \rrbracket_u)$ ;
4   while not_empty( $\mathcal{Q}$ ) do
5      $\varphi := pop(\mathcal{Q})$ ;
6     if  $\varphi \wedge \llbracket I \rrbracket_I$  satisfiable then
7       return unsafe
8     end
9     else if  $\varphi \notin \mathcal{V}$  then
10       $\mathcal{V} := \mathcal{V} \cup \{\varphi\}$ ;
11       $push(\mathcal{Q}, PRE_\tau(\varphi))$ ;
12    end
13  end
14  return safe

```

We assume the fomulas describing states to be of the form:

$$\varphi(\hat{X}) = \exists \vec{e} : eid. \Delta(\vec{e}) \wedge \Phi(\vec{e}, \hat{X})$$

The pre-image of a formula  $\varphi$  with respect to the set of transitions  $\tau$  (line 11) is given by:

$$PRE_\tau(\varphi)(\hat{X}) = \bigvee_{t \in \tau} PRE_t(\varphi)(\hat{X})$$

The pre-image of a formula  $\varphi$  with respect to a single transition  $t$  is given by:

$$\begin{aligned}
PRE_t(\varphi)(\hat{X}) = & \exists \hat{X}', \exists e_r, e_w, \vec{e} : eid. \Delta(e_r \cdot e_w \cdot \vec{e}) \wedge \\
& \llbracket t(X, X') \rrbracket_t^{e_r, e_w} \wedge \Phi(\vec{e}, \hat{X}') \wedge \\
& extend\_ghb(e_r, e_w, \vec{e}) \wedge rffr(X, X', e_w, \vec{e})
\end{aligned}$$

This pre-image generates two new event identifiers  $e_r$  and  $e_w$ . We ensure these identifiers to be different from those in  $\varphi$  using the expression  $\Delta(e_r \cdot e_w \cdot \vec{e})$ . These identifiers are given as parameters to the translation function  $\llbracket \cdot \rrbracket_t$ , which ensures the transition only manipulates new events.

The *rffr* function determines whether the new writes  $e_w$  from the transition  $t$  satisfy the compatible unsatisfied reads in  $\Phi(\vec{e}, \hat{X}')$ , and if so orders them appropriately in the *ghb* relation. A read and a write are compatible if they refer to the same variable or the same array with the same parameters (the actual value associated to the events is irrelevant to this notion of compatibility). For each unsatisfied read in  $\vec{e}$ , either:

- there is a compatible write  $e_w$  from the *same* process, so the read **MUST** take its value from this write
- there is a compatible write  $e_w$  from a *different* process, in this case the read **MAY** take its value from this write ; if it does,  $e_w$  is before  $e_r$  in *ghb*, otherwise  $e_r$  is before  $e_w$  in *ghb*
- there is no compatible write ; the read remains unsatisfied

This implies the following logical definition of *rffr*:

$$rffr(X, X', e_w, \vec{e}) = \bigwedge_{\alpha \in A} \bigwedge_{e_r \in \vec{e}} \left( \exists i, k : \text{proc. } \text{unsat\_read}_\alpha(e_r, i, k, \vec{e}) \rightarrow \right. \\ \left. \begin{array}{l} \text{internal\_rffr}_\alpha(X, X', e_r, e_w, i, k) \vee \\ \text{external\_rffr}_\alpha(X, X', e_r, e_w, i, k) \vee \\ \text{no\_rffr}_\alpha(X, X', e_r, e_w, k) \end{array} \right)$$

$$\text{internal\_rffr}_\alpha(X, X', e_r, e_w, i, k) = \text{Wr}_\alpha(e_w, i, k) \wedge \text{Val}_\alpha'(e_r, k) = \text{Val}_\alpha'(e_w, k)$$

$$\text{external\_rffr}_\alpha(X, X', e_r, e_w, i, k) = \exists j : \text{proc. } j \neq i \wedge \text{Wr}_\alpha(e_w, j, k) \wedge \\ (\text{Val}_\alpha'(e_r, k) = \text{Val}_\alpha'(e_w, k) \wedge \text{ghb}(e_w, e_r) \vee \\ \text{Val}_\alpha'(e_r, k) = \text{Val}_\alpha(e_r, k) \wedge \text{ghb}(e_r, e_w))$$

$$\text{no\_rffr}_\alpha(X, X', e_r, e_w, i, k) = (\nexists j : \text{proc. } \text{Wr}_\alpha(e_w, j, k)) \wedge \text{Val}_\alpha'(e_r, k) = \text{Val}_\alpha(e_r, k)$$

*rffr* also relies on the *unsat\_read*<sub>α</sub> function, that allows to determine if some read event *e<sub>r</sub>* is not satisfied by any write event in  $\vec{e}$  for some weak variable  $\alpha$ :

$$\text{unsat\_read}_\alpha(e_r, i, k, \vec{e}) = \text{Rd}_\alpha(e_r, i, k) \wedge \bigwedge_{e_w \in \vec{e}} \left( e_r = e_w \vee \left( \forall j : \text{proc. } \text{Wr}_\alpha(e_w, j, k) \rightarrow \text{ghb}(e_r, e_w) \right) \right)$$

The *extend\_ghb* function extends the *ghb* relation by adding literals of the form *ghb*(*e<sub>1</sub>*, *e<sub>2</sub>*) and *ghb-equal*(*e<sub>1</sub>*, *e<sub>2</sub>*) in the formula, according to the dependencies between the new events *e<sub>r</sub>* and *e<sub>w</sub>* and old events in  $\vec{e}$ .

$$\text{extend\_ghb}(e_r, e_w, \vec{e}) = \text{ppo}(e_r \cdot e_w, \vec{e}) \wedge \text{fence}(e_w, \vec{e}) \wedge \text{co}(e_w, \vec{e}) \wedge \text{fr}(e_r, \vec{e})$$

*fr* adds *ghb* pairs between the new read(s) *e<sub>r</sub>* and the old writes.

$$\text{fr}(e_r, \vec{e}) = \bigwedge_{\alpha \in A} \bigwedge_{e_w \in \vec{e}} \left( \exists k : \text{proc. } \text{read\_on}_\alpha(e_r, k) \wedge \text{write\_on}_\alpha(e_w, k) \rightarrow \text{ghb}(e_r, e_w) \right)$$

*co* adds *ghb* pairs between the new write(s) *e<sub>w</sub>* and the old writes.

$$\text{co}(e_{w_1}, \vec{e}) = \bigwedge_{\alpha \in A} \bigwedge_{e_{w_2} \in \vec{e}} \left( \exists k : \text{proc. } \text{write\_on}_\alpha(e_{w_1}, k) \wedge \text{write\_on}_\alpha(e_{w_2}, k) \rightarrow \text{ghb}(e_{w_1}, e_{w_2}) \right)$$

*fence* adds *ghb* pairs between the new writes(s) *e<sub>w</sub>* and the subsequent reads by the same process separated by a fence predicate.

$$\text{fence}(e_w, \vec{e}) = \bigwedge_{e_r \in \vec{e}} \left( \exists i : \text{proc. } \text{fence}(e_r, i) \wedge \text{write\_by}(e_w, i) \wedge \text{read\_by}(e_r, i) \rightarrow \text{ghb}(e_w, e_r) \right)$$

*ppo* adds *ghb* pairs between the new events *e<sub>a</sub>* and the subsequent events *e<sub>b</sub>* by the same process if they are read-read, read-write or write-write pairs.

$$\text{ppo}(\vec{e}_a, \vec{e}_b) = \bigwedge_{e_1 \in \vec{e}_a} \bigwedge_{e_2 \in \vec{e}_b} \left( \text{ppo\_RR}(e_1, e_2) \wedge \text{ppo\_RW}(e_1, e_2) \wedge \text{ppo\_WW}(e_1, e_2) \right)$$

$$\text{ppo\_RR}(e_1, e_2) = \exists i : \text{proc. } \text{read\_by}(e_1, i) \wedge \text{read\_by}(e_2, i) \rightarrow \text{ghb}(e_1, e_2)$$

$$ppo\_RW(e_1, e_2) = \exists i : proc.read\_by(e_1, i) \wedge write\_by(e_2, i) \rightarrow ghb(e_1, e_2)$$

$$ppo\_WW(e_1, e_2) = \exists i : proc.write\_by(e_1, i) \wedge write\_by(e_2, i) \rightarrow ghb(e_1, e_2)$$

The following  $read\_on_\alpha$  and  $write\_on_\alpha$  predicates allow to easily check whether some event identifier  $e$  corresponds to a read or a write event on a specific weak variable  $\alpha$  or array cell  $\alpha[k]$ .

$$read\_on_\alpha(e, k) = \exists i : proc.Rd_\alpha(e, i) \vee Rd_\alpha(e, i, k)$$

$$write\_on_\alpha(e, k) = \exists i : proc.Wr_\alpha(e, i) \vee Wr_\alpha(e, i, k)$$


Similarly, the  $read\_by$  and  $write\_by$  predicates allow to check whether some event identifier  $e$  corresponds to a read or a write event performed by a specific process  $i$ .

$$read\_by(e, i) = \exists k : proc. \bigvee_{\alpha \in A^0} Rd_\alpha(e, i) \vee \bigvee_{\alpha \in A^1} Rd_\alpha(e, i, j)$$

$$write\_by(e, i) = \exists k : proc. \bigvee_{\alpha \in A^0} Wr_\alpha(e, i) \vee \bigvee_{\alpha \in A^1} Wr_\alpha(e, i, j)$$

## 5 Example

We illustrate our backward reachability algorithm through a simple example. We consider a simple parameterized mutual exclusion algorithm, where each process executes the automaton below.



This automaton is trivially encoded into the following transition system, where the current state is represented by a regular array  $PC$ , and the shared variables  $X_i$  are expressed as a weak array  $X$ .

$$t\_req = \exists i : proc.PC[i] = Idle \wedge PC'[i] = Want \wedge X'[i] = True$$

$$t\_enter = \exists i : proc.PC[i] = Want \wedge fence() \wedge (\forall k.k = i \vee X[i] = False) \wedge PC'[i] = Crit$$

$$t\_exit = \exists i : proc.PC[i] = Crit \wedge PC'[i] = Idle \wedge X'[i] = False$$

The initial state is simply defined as:

$$I = \forall j : proc.PC[j] = Idle \wedge X[j] = False$$

The property we would like to check is that no pair of processes  $i$  and  $j$  can be in state  $Crit$  at the same time, *i.e.* that the following formula never holds:

$$\Theta = \exists i, j : proc.i \neq j \wedge PC[i] = Crit \wedge PC[j] = Crit$$

By applying the translation functions on the system above, we obtain the following event-explicit transition system:

$$\tilde{I} = \forall i, j : proc. \forall e_X : ur.PC[j] = Idle \wedge Rd_X(e_X, i, j) \wedge Val_X(e_X, j) = False$$

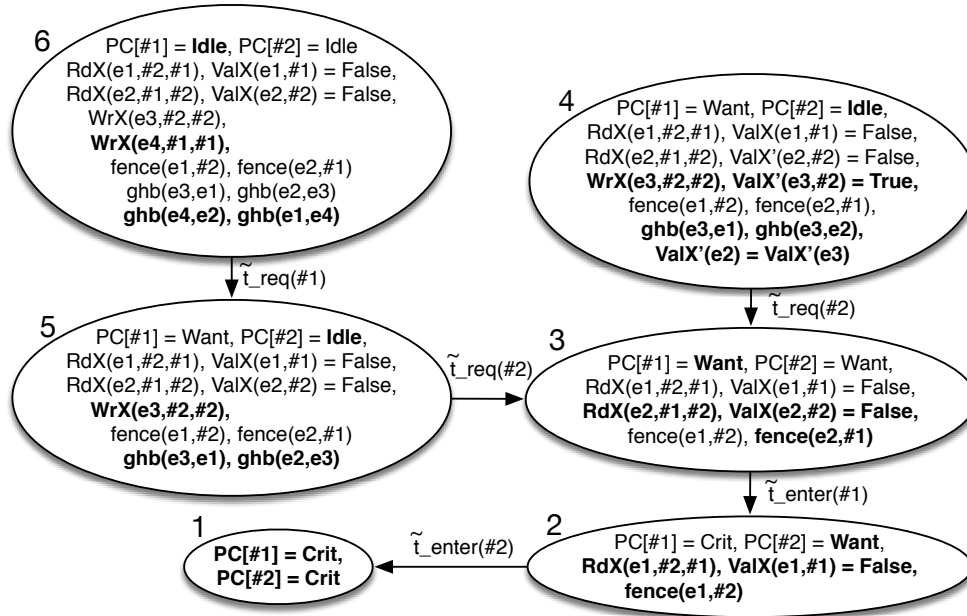
$$\tilde{\Theta} = \exists i, j : \text{proc. } i \neq j \wedge PC[i] = \text{Crit} \wedge PC[j] = \text{Crit}$$

$$\begin{aligned} \tilde{t}_{req}(e_r, e_w) = & \exists i : \text{proc. } e_r \neq e_w \wedge \text{ghb-equal}(e_r, e_w) \wedge \\ & PC[i] = \text{Idle} \wedge \\ & PC'[i] = \text{Want} \wedge Wr_X(e_w, i, i) \wedge Val'_X(e_w, i) = \text{True} \end{aligned}$$

$$\begin{aligned} \tilde{t}_{enter}(e_r, e_w) = & \exists i : \text{proc. } e_r \neq e_w \wedge \text{ghb-equal}(e_r, e_w) \wedge \\ & PC[i] = \text{Want} \wedge \text{fence}(e_r, i) \wedge \\ & (\forall k. k = i \vee Rd_X(e_r, i, k) \wedge Val_X(e_r, i) = \text{False}) \wedge \\ & PC'[i] = \text{Crit} \end{aligned}$$

$$\begin{aligned} \tilde{t}_{exit}(e_r, e_w) = & \exists i : \text{proc. } e_r \neq e_w \wedge \text{ghb-equal}(e_r, e_w) \wedge \\ & PC[i] = \text{Crit} \wedge \\ & PC'[i] = \text{Idle} \wedge Wr_X(e_w, i, i) \wedge Val'_X(e_w, i) = \text{False} \end{aligned}$$

The graph below gives a possible exploration of the system's state space by our algorithm. We start by node 1, which represents the formula describing the dangerous states  $\tilde{\Theta}$ . Then, each node represents the result of a pre-image computation by an instance of a transition. The edges are labeled with the transition name and the process identifier we used to instantiate the existentially quantified process variable in the transition. Remark that formulas in the graph's nodes are implicitly existentially quantified and that a process identifier  $i$  is written  $\#i$ . Also, to avoid cluttering the graph, we omit event identifiers in the edge labels, we remove the unused event identifiers from the nodes (they do not contribute to *ghb*), and we assume all event identifiers to be different.



We focus on node 3 which results from the pre-image of node 1 by  $t\_enter(\#2)$  then  $t\_enter(\#1)$ . In this state, both processes have read False in X (events  $e_1$  and  $e_2$ ). Also, since there is a memory barrier in  $t\_enter$ , both reads are associated to a fence literal. The pre-image of node 3 by  $t\_req(\#2)$  introduces a new write event  $Wr_X(e_3, \#2, \#2)$  with an associated value  $Val_X'(e_3, \#2) = \text{True}$ . Since there is a memory barrier  $fence(e_1, \#2)$  on  $e_1$  by the same process  $\#2$ , the *extend\_ghb* predicate causes

$ghb(e3, e1)$  to be added to the formula. Now, the *rfrr* predicate dictates that this new write event may or may not satisfy the read  $e2$ , so we must consider both cases (node 4 and 5).

In node 4, we consider the case where the write event  $e3$  satisfies the read event  $e2$ . As prescribed by the *external\_rfrr* predicate, the equality  $ValX'(e2, \#2) = ValX'(e3, \#2)$  is added to the formula, which obviously makes it inconsistent. In node 5, the write event  $e3$  does not satisfy the read event  $e2$ , so the value  $ValX'(e3, \#2)$  is discarded and  $ghb(e2, e3)$  is added to the formula, as indicated by the *no\_rfrr* predicate. Similarly, the pre-image of node 5 by  $t\_req(\#1)$  yields the formula in state 6 where the new write event  $e4$  does not satisfy the read event  $e1$ . Now, the *ghb* relation is not a valid partial order as the sequence  $ghb(e2, e3), ghb(e3, e1), ghb(e1, e4), ghb(e4, e2)$  forms a cyclic relation. Therefore, this state is discarded and the program is declared *safe*.

Remark that if we removed the fence predicate in  $t\_enter$ , then we would only have  $ghb(e3, e1), ghb(e4, e2)$  in state 6, which is a valid partial order relation, so the formula would intersect with the initial state and the program would be *unsafe*.

## 6 Implementation

We have implemented this framework in Cubicle- $\mathscr{W}$  [1] and used it to check the correctness of several parameterized concurrent algorithms on a TSO-like model whose source codes are available on the tool's webpage.

Table 1 gives for each benchmark the number of non-weak arrays, the number of weak variables, the number of weak arrays, the number of transitions and the running time. The S/US next to the benchmark name indicates whether the algorithm is correct or not. Incorrect algorithm have a second version that was fixed by using fence predicates.

Case study		Arrays	Weak Var.	Weak Arr.	Transitions	Time
naive mutex	US	1	0	1	4	0.04s
naive mutex	S	1	0	1	4	0.30s
lamport	US	1	4	0	8	0.10s
lamport	S	1	4	0	8	0.60s
spinlock	S	1	1	0	6	0.07s
sense reversing barrier	S	1	0	1	4	0.06s
arbiter v1	S	2	1	1	7	0.18s
arbiter v2	S	2	0	2	8	13.5s
two phase commit	S	1	1	1	5	54.1s

Table 1: Performance of Cubicle- $\mathscr{W}$

The results shown in this table look promising. Although these benchmarks are of modest size, they are already consider as very challenging for state-of-the-art model checkers for weak memories as they combined parametricity, concurrency and non trivial used of weak memories.

## 7 Conclusion & Future Work

We have presented in this paper an extension of Model Checking Modulo Theories (MCMT) for model checking parameterized transitions with *explicit* read and write operations on weak memories.

The core of our procedure is a backward reachability algorithm combined with an SMT axiomatic model for reasoning about weak memory. The explicit relaxed consistency model underlying our framework is similar to x86-TSO where the effect of a store operation by a process is delayed (due to a store buffering) to all processes.

We have implemented this framework in Cubicle- $\mathcal{W}$ , a conservative extension of the Cubicle model checker. Our first experiments show that our implementation is expressive and efficient enough to prove safety of concurrent algorithms, for an arbitrary number of processes, ranging from mutual exclusion to synchronization barriers.

Immediate future work includes the support for other models, such as PSO. We are also working on the extension of Cubicle's invariant generation mechanism for weak memory. With this mechanism, we should gain in efficiency and be able to tackle even more complex programs. Finally, we also plan to investigate how our framework for weak memories could be extended to reason about programs communicating via channels.

## References

- [1] <https://www.lri.fr/~declerck/cubiclew/>.
- [2] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Ahmed Bouajjani & Tuan Phong Ngo (2016): *The Benefits of Duality in Verifying Concurrent Programs under TSO*. In: *CONCUR*.
- [3] Parosh Aziz Abdulla, Giorgio Delzanno, Noomene Ben Henda & Ahmed Rezine (2007): *Regular Model Checking Without Transducers*. In: *TACAS*, Springer.
- [4] Parosh Aziz Abdulla, Giorgio Delzanno & Ahmed Rezine (2007): *Parameterized verification of infinite-state processes with global conditions*. In: *CAV*, Springer.
- [5] Jade Alglave, Luc Maranget & Michael Tautschnig (2014): *Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory*. *ACM Trans. Program. Lang. Syst.* 36(2), pp. 7:1–7:74.
- [6] K R Apt & D C Kozen (1986): *Limits for automatic verification of finite-state concurrent systems*. *Inf. Process. Lett.* 22(6), pp. 307–309.
- [7] E. M. Clarke, O. Grumberg & M. C. Browne (1986): *Reasoning About Networks with Many Identical Finite-state Processes*. In: *PODC*, ACM, New York, NY, USA.
- [8] Sylvain Conchon, Amit Goel, Sava Krstić, Alain Mebsout & Fatiha Zaïdi (2012): *Cubicle: A Parallel SMT-based Model Checker for Parameterized Systems: Tool Paper*. In: *CAV*, Berlin, Heidelberg.
- [9] Steven M. German & A. Prasad Sistla (1992): *Reasoning about systems with many processes*. *J. ACM* 39(3), pp. 675–735.
- [10] Silvio Ghilardi, Enrica Nicolini, Silvio Ranise & Daniele Zucchelli (2008): *Towards SMT Model Checking of Array-Based Systems*. In: *Automated Reasoning*, 5195, pp. 67–82.
- [11] Silvio Ghilardi & Silvio Ranise (2010): *MCMT: A Model Checker Modulo Theories*. In: *IJCAR*, pp. 22–29.
- [12] Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli & Magnus O. Myreen (2010): *X86-TSO: A Rigorous and Usable Programmer's Model for x86 Multiprocessors*. *CACM* 53(7).