Cubicle-$\mathcal{W}$: Parameterized Model Checking on Weak Memory

Sylvain Conchon[1,2], David Declerck[1,2], and Fatiha Zaïdi[1]

[1] LRI (CNRS & Univ. Paris-Sud), Université Paris-Saclay, F-91405 Orsay
[2] Inria, Université Paris-Saclay, F-91120 Palaiseau

Abstract. We present Cubicle-$\mathcal{W}$, a new version of the Cubicle model checker to verify parameterized systems under weak memory models. Its main originality is to implement a backward reachability algorithm modulo weak memory reasoning using SMT. Our experiments show that Cubicle-$\mathcal{W}$ is expressive and efficient enough to automatically prove safety of concurrent algorithms, for an arbitrary number of processes, ranging from mutual exclusion to synchronization barriers.

Keywords: Parameterized Model Checking, MCMT, SMT, Weak Memory

1 Introduction

Concurrent algorithms are usually designed under the sequential consistency (SC) memory model [19] which enforces a global-time linear ordering of (read or write) accesses to shared memories. However, modern multiprocessor architectures do not follow this SC semantics. Instead, they implement several optimizations which lead to relaxed consistency models on shared memory where read and write operations may be reordered. For instance, in x86-TSO [20, 21] writes can be delayed after reads due to a store buffering mechanism. Other relaxed models (PowerPC [6], ARM) allow even more types of reorderings.

The new behaviors induced by these models may make out-of-the-shelf algorithms incorrect for subtle reasons mixing interleaving and reordering of events. In this context, finding bugs or proving the correctness of concurrent algorithms is very challenging. The challenge is even more difficult if we consider that most algorithms are *parameterized*, that is designed to be run on architectures containing an arbitrary (large) number of processors.

One of the most efficient technique for verifying concurrent systems is model checking. While this technique has been used to verify parameterized algorithms [12, 5, 4, 16, 9, 2, 13] and systems under some relaxed memory assumptions [7, 11, 10, 3, 2], hardly any state-of-the-art model checker support *both* parameterized verification and weak memory models [2].

In this paper, we present Cubicle-$\mathcal{W}$ [1], the new version of the model checker Cubicle [13–15] for verifying safety properties of parameterized array-based transition systems on weak memory. Cubicle-$\mathcal{W}$ is a conservative extension which allows the user to manipulate both SC and weak variables. Its relaxed consistency model is similar to x86-TSO : each process has a FIFO buffer of pending store operations whose side effect is to delay the outcome of its memory writes to all processes.

The core of Cubicle-$\mathcal{W}$ extends Cubicle's SMT-based backward reachability procedure with a new pre-image computation which takes into account the delays between write and read operations. In order to consider only coherent read/write pairs, Cubicle-$\mathcal{W}$ relies on a buffer-free memory model inspired by the logical framework of [8] which is implemented as a new theory in its SMT solver. Cubicle-$\mathcal{W}$ is an open-source software freely available at `http://cubicle.lri.fr/cubiclew`.

2 Tool Presentation

The syntax of Cubicle-$\mathcal{W}$ extends Cubicle's with new constructs for manipulating weak memories. The reader can refer to [13] for the description of Cubicle's input language.

Variable and array declarations can now be prefixed by the keyword `weak` for defining weak memories.

```
weak var X : int
weak array A[proc] : bool
```

Transitions in Cubicle-$\mathcal{W}$ have the same syntactic guard/action form as in Cubicle and they are also supposed to be executed atomically. The new feature is that they must now have *at least* one parameter which represents the process that performs the operations. This parameter is identified using the `[.]` notation. For instance, in the following example, the parameter `[i]` of transition `t1` represents the process performing all read/write operations on `X`, `A[i]` and `A[j]` when `t1` is triggered.

```
transition t1 ([i] j)
requires { X = 42 && A[i] = False }
{ A[j] := False }
```

Even if there is no use of parameter `[i]` in transitions' guards and actions, this parameter is still mandatory, as in the transition `t2` below, to indicate which process performs the operations.

```
transition t2 ([i]) { X := 42 }
```

Note that, as Cubicle-$\mathcal{W}$'s transitions are atomic, having several processes performing reads or writes operations in the same transition would require an unrealistic powerful synchronization mechanism between processes.

The main aspect of our relaxed memory semantics is that, from a global viewpoint, the effect of a write operation on a `weak` memory is not immediately visible to all processes. It is only locally visible to the process that performs it. For instance, if some process `i` executes the transition `t2` above, then `X = 42` is true for `i` after the transition (as the effect of the assignment is immediately locally visible), while all other processes can still read a different value for `X`.

To enforce the global visibility of a write operation, one has to use a *memory barrier*. In Cubicle-$\mathcal{W}$, barriers are provided as a new built-in predicate `fence()`. When used in the guard of a transition, `fence` is true only when the FIFO buffer of the parameter `[i]` of the transition is empty. For instance, if a process executes `t2` then the following transition `t3`:

transition `t3` **([i]) requires** `{ fence() }{ ... }`

the `fence` predicate in `t3`'s guard ensures that the effect of all previous assignments done by `i` are visible to all processes after `t3`. Note that `fence` is not an action: it does not force buffers to be flushed on memory, but just *waits* for a buffer to be empty. As a consequence, it can only be used in a guard.

Implicit memory barriers are also activated when a transition contains both a read and a write to weak variables (not necessarily the same). For instance, the execution of the following transition `t4` guarantees that the buffer of process `i` is empty *before* and *after* `t4`.

transition `t4` **([i])**
requires `{ A[i] = False }`
`{ X := 1 }`

Because there is no unique view of the contents of weak variables, one can not talk about *the value* of `X`, but rather the value of `X` from the *point of view* of a process `i`, denoted `i@X` in Cubicle-$\mathcal{W}$. This notation is used when describing unsafe states. For instance, in the following formula, a state is defined as unsafe when there exist two (distinct) processes `i` and `j` reading respectively `42` and `0` in the weak variable `X`:

`unsafe (i j) { i@X = 42 && j@X = 0 }`

This notation is not used for describing initial states as Cubicle-$\mathcal{W}$ implicitly assumes that *all* processes have the same view of each weak variable in those states. For instance, the following formula defines initial states where, for all processes, `X` equals `0` and all cells of array `A` contain `False`.

`init (i) { X = 0 && A[i] = False }`

Finally, it is important to note that non weak arrays are restricted to be used only locally by processes: given a non weak array `T`, only `i` can read or write to `T[i]`.

3 Backward Reachability Modulo Weak Memory

The core of Cubicle-$\mathcal{W}$ is an extension of Cubicle's symbolic backward reachability algorithm [13, 14]. We first briefly recall how the original Cubicle works, then we give details about our new algorithm.

States in Cubicle are represented by cubes, i.e., formulas of the form $\exists \bar{i}.(\Delta \wedge F)$, where $\bar{i}$ is a set of process variables, Δ is the conjunction of all disequations between the variables in $\bar{i}$ and F is a conjunction of literals. Each literal in F is a comparison $(=, \neq, <, \leq)$ between two terms. A term can be a constant (integer, boolean, real, constructor), a process variable (i), a variable (X) or an array access $(A[i]$, where i is a process variable). All process variables in a state are implicitly existentially quantified. Initial states are represented by a universally quantified formula I of the form $\forall \bar{i}.(\Delta \wedge F)$, where Δ and F are as described above.

The core of Cubicle is a symbolic backward reachability loop that maintains two collections of states: $\mathcal{Q}$ contains the states to visit (it is initialized with the states declared as `unsafe`), and $\mathcal{V}$ is filled with the visited states (initially empty). Each iteration of the loop performs the following operations:

1. (*pop*) retrieve and remove a formula φ from $\mathcal{Q}$
2. (*safety test*) check the satisfiability of $\varphi \wedge I$, i.e. determine if the states described by φ intersect with the initial states I. If so, the system is declared as *unsafe*
3. (*fixpoint test*) check if $\varphi \models \mathcal{V}$ is valid, i.e. determine if the states described by φ have already been visited. If so, discard φ and go back to 1
4. (*pre-image computation*) compute the pre-image $pre(\varphi, t)$ of φ for all instances of transitions t, i.e. determine the set of states that can reach φ in one step by applying t with the processes identifiers `#1, ..., #n` as parameters, add these states to $\mathcal{Q}$ and add φ to $\mathcal{V}$.

If $\mathcal{Q}$ is empty at step 1, then all the states space has been explored and the system is declared *safe*. Note that the (non-trivial) fixpoint and safety tests are discharged to an embedded SMT solver.

Cubicle-$\mathcal{W}$ uses the same procedure but some operations have been extended to reason modulo an axiomatic description of our weak memory model. This axiomatization uses the notion of *events* to describe weak memory accesses and a global-happens-before (*ghb*) relation defined as a partial order relation over these events. This relation is used to determine if an execution is valid.

Our logic is extended with new literals to represent read and write operations on weak memories. We assume given a (countable) set of events $\mathcal{E}$. A literal of the form $e{:}\mathrm{Rd}_{\mathrm{X}}(i)$ denotes a read access on variable X by a process i labeled with an event identifier $e \in \mathcal{E}$. Similarly, literals of the form $e{:}\mathrm{Wr}_{\mathrm{X}}(i)$ represent write accesses. The value returned by a read (resp. assigned by a write) is given by the term $val(e)$, where e is the event identifier associated to the operation. Operations on weak arrays are represented by literals of the form $e{:}\mathrm{Rd}_{\mathrm{A}}(i, j)$ and $e{:}\mathrm{Wr}_{\mathrm{A}}(i, j)$, which represent an access by a process i to the cell j of an array A. Last, there is also literals of the form $e{:}\mathtt{fence}(i)$ which indicate that a process

i has a memory barrier on the event e, where e is an event identifier associated to a read by the same process.

The reachability loop of Cubicle-$\mathcal{W}$ implements a new pre-image computation. At step 4, $pre(\varphi, t)$ is modified so that read and write operations from t give rise to `Rd` and `Wr` literals labeled with fresh event identifiers. These new events are ordered w.r.t the older ones in the ghb relation expressed by predicates of the form $ghb(e_1, e_2)$, indicating that event e_1 is $ghb\text{-}before$ (i.e., occurs before) event e_2. The ghb-ordering of events is built w.r.t. the following rules:

- New read events are $ghb\text{-}before$ old read and write events from the same process.
- New write events are $ghb\text{-}before$ old write events from the same process, however they are $ghb\text{-}before$ old reads events from the same process only if there is a fence on these reads.
- New write events are $ghb\text{-}before$ all the old write events to the same variable.
- New read events are $ghb\text{-}before$ all the old write events to the same variable.

Finally, when a memory fence is encountered, a literal e:`fence`(i) is added on all old reads events e which belong to the process i executing the transition.

The treatment of write events is also specific when we have to consider the delays introduced by store buffers: when a new write event e is produced, all possible combinations of e with older compatible reads are considered (unlike in SC), as a write operation *may or may not* satisfy subsequent reads. By *compatible read*, we mean a read on the same variable or array cell as the write, though we may also consider the constant values associated to these events in order to obtain a more accurate set of compatible reads. The connection between a write and an older read obeys the following rules:

- When the write event satisfies an old read event from a different process, the write is $ghb\text{-}before$ the read.
- When the write event does not satisfy an old read event from a different process, the read is $ghb\text{-}before$ the write.
- When the write and the read events belong to the same process, none of them is considered $ghb\text{-}before$ the other (unless there is a fence on the read event).

In order to show how our reachability procedure works, we consider the simple parameterized mutual exclusion algorithm and the exploration graph given below. Cubicle-$\mathcal{W}$ starts with the `unsafe` formula in node 1. Then, each node represents the result of a pre-image computation by an instance of a transition (denoted by the label of the edge). Remark that formulas in the graph's nodes are implicitly existentially quantified and that a process identifier i is written $\#_i$.

```
type loc = Idle | Want | Crit          transition t_req ([i])
                                        requires { PC[i] = Idle }
weak array X[proc] : bool               { X[i] := True ; PC[i] := Want }
array PC[proc] : loc
                                        transition t_enter ([i])
init (i) {PC[i] = Idle && X[i] = False} requires { PC[i] = Want && fence() &&
                                                    forall_other k. X[k] = False }
unsafe (i j) {PC[i] = Crit && PC[j] = Crit}  { PC[i] := Crit }

                                        transition t_exit ([i])
                                        requires { PC[i] = Crit }
                                        { X[i] := False ; PC[i] := Idle }
```
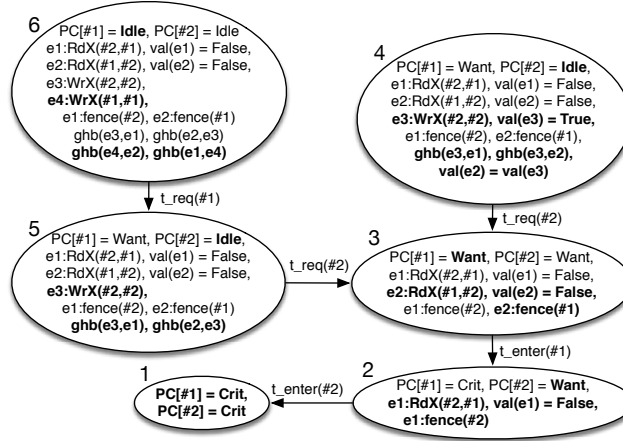

We focus on node 3 which results from the pre-image of node 1 by `t_enter(#2)` then `t_enter(#1)`. In this state, both processes have read `False` in `X` (events e_1 and e_2). Also, since there is a memory barrier in `t_enter`, both reads are associated to a `fence` literal. The pre-image of node 3 by `t_req(#2)` introduces a new write event `e3:WrX(#2,#2)` with an associated value `val(e3) = True`. Since there is a memory barrier `e1:fence(#2)` on `e1` by the same process `#2`, we add `ghb(e3,e1)` in the formula. Now, 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, event `e3` satisfies `e2`. The equality `val(e2) = val(e3)` is then added to the formula which obviously makes it inconsistent. In node 5, the write `e3` does not satisfy the read `e2`, then the value `val(e3)` is discarded and `ghb(e2,e3)` is added to the formula. Similarly, the pre-image of node 5 by `t_req(#1)` yields the formula in state 6 where the new write `e4` does not satisfy the read `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*.

4 Benchmarks and Conclusion

We have evaluated Cubicle-$\mathcal{W}$ on some classical parameterized concurrent algorithms (available on the tool's webpage [1]). We compared its performances with state-of-the-art model checkers supporting the TSO weak memory model, since our model is similar. The model checkers we used are CBMC [7], Trencher [11, 10], MEMORAX [3] and Dual-TSO [2]. As most of these tools do not support parameterized systems, we used them on fixed-size instances of our benchmarks and increased the number of processes until we obtained a timeout (or until we reached a high number of processes, *i.e.* 11 in our case). Dual-TSO supports a restricted form of parameterized systems, but does not allow process-indexed arrays, which are often needed to express parameterized programs. When it was possible, we used it on both parameterized and non parameterized versions of our benchmarks.

		Cubicle $\mathcal{W}$	Memorax SB	Memorax PB	Trencher	CBMC Unwind 2	CBMC Unwind 3	Dual TSO
naive mutex	US	0.04s [N]	– TO [6] 7m54s [5]	– TO [10] 12m02s [9]	– TO [5] 10.1s [4]	– 23.6s [11] 14.7s [10]	– 5m37s [11] 3m39s [10]	NT [N] TO [6] 1m12s [5]
naive mutex	S	0.30s [N]	– TO [5] 23.3s [4]	– TO [11] 2m28s [10]	– TO [6] 54.8s [5]	– TO [5] 2m24s [4]	– TO [3] 19.4s [2]	NT [N] TO [5] 35.7s [4]
lamport	US	0.10s [N]	– TO [4] 17.4s [3]	– TO [4] 25.4s [3]	– KO [4] 1.73s [3]	– 7m42s [11] 4m29s [10]	– TO [7] 5m12s [6]	NT [N] TO [6] 13m12s [5]
lamport	S	0.60s [N]	– TO [3] 0.14s [2]	– TO [4] 3m02s [3]	– KO [5] 3.37s [4]	– TO [4] 8m39s [3]	– TO [3] 1m55s [2]	NT [N] TO [4] 9.42s [3]
spinlock [21]	S	0.07s [N]	– TO [5] 8m51s [4]	– TO [7] 9m52s [6]	– TO [7] 21.45s [6]	– TO [3] 19.58s [2]	– TO [3] 5m08s [2]	TO [N] TO [6] 1m16s [5]
sense [18] reversing barrier	S	0.06s [N]	– TO [3] 0.34s [2]	– TO [3] 0.09s [2]	– TO [5] 1m58s ☠ [4]	– TO [9] 12m25s [8]	– TO [4] 1m43s [3]	NT [N] TO [3] 0.09s [2]
arbiter v1 [17]	S	0.18s [N]	– TO [1+2]	– TO [1+2]	– KO [1+5] 4.57s [1+4]	– TO [1+6] 12m02s [1+5]	– TO [1+3] 44.3s [1+2]	NT [N] TO [1+6] 2m45s ☠ [1+5]
arbiter v2 [17]	S	13.5s [N]	– TO [1+2]	– TO [1+2]	– KO [1+4] 1.62s [1+3]	– TO [1+4] 2m56s [1+3]	– TO [1+2]	NT [N] TO [1+3] 24.2s [1+2]
two phase commit	S	54.1s [N]	– TO [2]	– TO [4] 39.7s [3]	– TO [4] 7.08s ☠ [3]	– TO [11] 12m39s [10]	– TO [11] 13m41s [10]	NT [N] TO [3] 12.3s [2]

Table 1. Performance comparison of Cubicle-$\mathcal{W}$ and other tools

Table 1 gives the running time for each benchmark, with the number of processes between square brackets, where N indicates the parametric case. The second column indicates whether the program is expected to be unsafe (US) or safe (S). Unsafe programs have a second version that was fixed by adding fence predicates. ☠ indicates that a tool gave a wrong answer. **KO** means that a tool crashed. **NT** indicates a benchmark that was not translatable to Dual-TSO.

The tests were run on a MacBook Pro with an Intel Core i7 CPU @ 2,9 Ghz and 8GB of RAM, under OSX 10.11.6. The timeout (TO) was set to 15 minutes.

These results show that state-of-the-art model checkers suffer from scalability issues and justify the use of parameterized techniques. Cubicle-$\mathcal{W}$ is thus a very promising approach to the verification of concurrent programs that are both parameterized and operating under weak memory.

References

1. Cubicle-$\mathcal{W}$, `https://www.lri.fr/~declerck/cubiclew/`
2. Abdulla, P.A., Atig, M.F., Bouajjani, A., Ngo, T.P.: The benefits of duality in verifying concurrent programs under TSO. In: CONCUR (2016)
3. Abdulla, P.A., Atig, M.F., Chen, Y., Leonardsson, C., Rezine, A.: Memorax, a precise and sound tool for automatic fence insertion under TSO. In: TACAS (2013)
4. Abdulla, P.A., Delzanno, G., Henda, N.B., Rezine, A.: Regular model checking without transducers. In: TACAS (2007)
5. Abdulla, P.A., Delzanno, G., Rezine, A.: Parameterized verification of infinite-state processes with global conditions. In: CAV (2007)
6. Alglave, J., Fox, A., Ishtiaq, S., Myreen, M.O., Sarkar, S., Sewell, P., Nardelli, F.Z.: The semantics of power and arm multiprocessor machine code. In: DAMP (2008)
7. Alglave, J., Kroening, D., Nimal, V., Tautschnig, M.: Software verification for weak memory via program transformation. In: ESOP (2013)
8. Alglave, J., Maranget, L., Tautschnig, M.: Herding cats: Modelling, simulation, testing, and data mining for weak memory. ACM Trans. Program. Lang. Syst. (2014)
9. Apt, K.R., Kozen, D.C.: Limits for automatic verification of finite-state concurrent systems. Inf. Process. Lett. (1986)
10. Bouajjani, A., Calin, G., Derevenetc, E., Meyer, R.: Lazy tso reachability. In: FASE (2015)
11. Bouajjani, A., Derevenetc, E., Meyer, R.: Checking and enforcing robustness against tso. In: ESOP (2013)
12. Clarke, E.M., Grumberg, O., Browne, M.C.: Reasoning about networks with many identical finite-state processes. In: PODC (1986)
13. Conchon, S., Goel, A., Krstić, S., Mebsout, A., Zaïdi, F.: Cubicle: A parallel smt-based model checker for parameterized systems: Tool paper. In: CAV (2012)
14. Conchon, S., Goel, A., Krstic, S., Mebsout, A., Zaidi, F.: Invariants for Finite Instances and Beyond. In: FMCAD (2013)
15. Conchon, S., Mebsout, A., Zaïdi, F.: Certificates for parameterized model checking. In: FM (2015)
16. German, S.M., Sistla, A.P.: Reasoning about systems with many processes. J. ACM (1992)
17. Goeman, H.J.M.: The arbiter: an active system component for implementing synchronizing primitives. Fundam. Inform. (1981)
18. Herlihy, M., Shavit, N.: The Art of Multiprocessor Programming (2008)
19. Lamport, L.: How to make a multiprocessor computer that correctly executes multiprocess programs. IEEE Trans. Comput. (1979)
20. Owens, S., Sarkar, S., Sewell, P.: A better x86 memory model: X86-TSO. In: TPHOL (2009)
21. Sewell, P., Sarkar, S., Owens, S., Nardelli, F.Z., Myreen, M.O.: X86-TSO: A rigorous and usable programmer's model for x86 multiprocessors. CACM (2010)