

# The Maude LTL Model Checker and Its Implementation<sup>\*</sup>

Steven Eker<sup>1</sup>, José Meseguer<sup>2</sup>, and Ambarish Sridharanarayanan<sup>2</sup>

<sup>1</sup> Computer Science Laboratory, SRI International  
Menlo Park, CA 94025  
eker@csl.sri.com

<sup>2</sup> CS Department, University of Illinois at Urbana-Champaign  
Urbana, IL 61801  
{meseguer, srdhrnry}@cs.uiuc.edu

## 1 Introduction

A model checker typically supports two different levels of specification: (1) a *system specification* level, in which the concurrent system to be analyzed is formalized; and (2) a *property specification* level, in which the properties to be model checked—for example, temporal logic formulae—are specified. The Maude LTL model checker has been designed with the goal of combining a very expressive and general system specification language (Maude [1]) with an advanced on-the-fly explicit-state LTL model checking engine.

Maude specifications are *executable logical theories* in rewriting logic [9], a logic that is a flexible logical framework for expressing a very wide range of concurrency models and distributed systems [9]. A rewrite theory is a triple  $\mathcal{R} = (\Sigma, E, R)$ , with  $(\Sigma, E)$  an equational theory specifying a system's *distributed state structure* and with  $R$  a collection of rewrite rules specifying the *concurrent transitions* of the system. Since no domain-specific model of concurrency is built into the logic, the range of applications that can be naturally specified is indeed very wide. For example, besides conventional distributed system specifications, properties of signalling pathways in mammalian cells have been model checked [3]. Another advantage of Maude as the system specification language is that integration of model checking with theorem proving techniques becomes quite seamless. The same rewrite theory  $\mathcal{R} = (\Sigma, E, R)$  can be the input to the LTL model checker and to several other proving tools in the Maude environment [2]. For a lengthier discussion of the Maude LTL model checker and its LTL satisfiability and tautology procedures, see the companion paper [4].

---

<sup>\*</sup> The first two authors' work has been partially supported by DARPA through Air Force Research Laboratory Contract F30602-97-C-0312, NSF grants CCR-9900326 and CCR-9900334, ONR Contract N00012-99-C-0198, and DARPA through Air Force Research Laboratory Contract F30602-02-C-0130. The last two authors' work is also supported in part by the ONR Grant N00014-02-1-0715.

## 2 LTL Model Checking of Maude Specifications

A Maude module is a rewrite theory  $\mathcal{R} = (\Sigma, E, R)$ . Fixing a distinguished sort *State*, the initial model  $\mathcal{T}_{\mathcal{R}}$  of the rewrite theory  $\mathcal{R} = (\Sigma, E, R)$  has an underlying Kripke structure  $\mathcal{K}(\mathcal{R}, \textit{State})$  given by the total binary relation extending its one-step sequential rewrites. To the initial algebra of states  $T_{\Sigma/E}$  we can likewise associate equationally-defined *computable state predicates* as atomic predicates for such a Kripke structure. In this way we obtain a language of LTL *properties* of the rewrite theory  $\mathcal{R}$ . Maude 2.0 supports on-the-fly LTL model checking for initial states  $[t]$ , say of sort *State*, of a rewrite theory  $\mathcal{R} = (\Sigma, E, R)$  such that the set  $\{[u] \in T_{\Sigma/E} \mid \mathcal{R} \vdash [t] \rightarrow [u]\}$ , of all states *reachable* from  $[t]$  is *finite*. The rewrite theory  $\mathcal{R}$  should satisfy reasonable executability requirements, such as the confluence and termination of the equations  $E$  and coherence of the rules  $R$  relative to  $E$  [1]. In Maude the rewrite theory  $\mathcal{R}$  is specified as a module, say  $M$ . Then, given an initial state, say `init` of sort `StateM`, we can *model check* different LTL properties beginning at this initial state by doing the following:

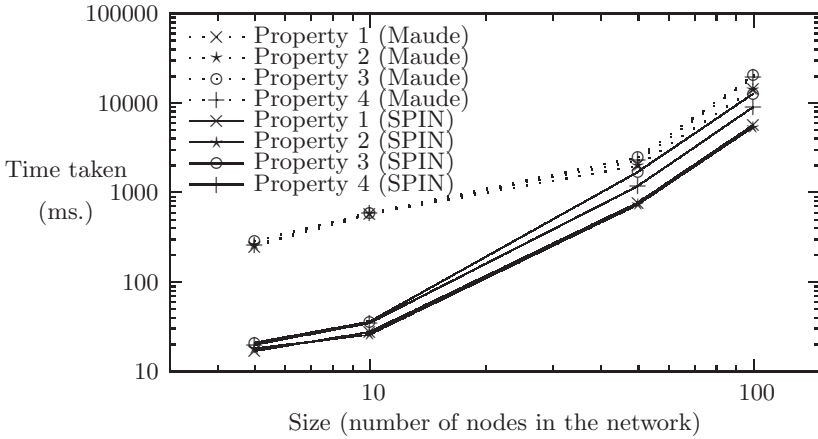
- defining a new module, say `CHECK-M`, that includes the modules  $M$  and the predefined module `MODEL-CHECKER` as submodules;
- giving a *subsort declaration*, `subsort StateM < State .`, where `State` is one of the key sorts in the module `MODEL-CHECKER`;
- defining the *syntax* of the *state predicates* we wish to use (which can be parametric) by means of constants and operators of sort `Prop`, a subsort of the sort `Formula` (i.e., LTL formulas) in the module `MODEL-CHECKER`;
- defining the *semantics* of the state predicates by means of equations.

Once the semantics of each of the state predicates has been defined, we are then ready, given an initial state `init`, to model check any LTL formula, say `form`, involving such predicates. We do so by evaluating in Maude, the expression `modelCheck(init,form) .` Two things can then happen: if the property `form` holds, then we get the result `true`; if it doesn't, we get a counterexample expressed as a finite path followed by a cycle.

## 3 Model Checking Algorithms and Implementation

On-the-fly LTL model checking is performed by constructing a Büchi automaton from the negation of the property formula and lazily searching the synchronous product of the Büchi automaton and the system state transition diagram for a reachable accepting cycle.

**Büchi Automaton Construction.** The negated LTL formula is converted to negative normal form and heuristically simplified by a set of Maude equations, mostly derived from the simplification rules in [5, 10]. Rather than the classical tableaux construction [7], we use a newer technique proposed in [6] based on very weak alternating automata, comprising three basic steps: (1) construct a very weak alternating automaton from the formula, (2) convert the very weak

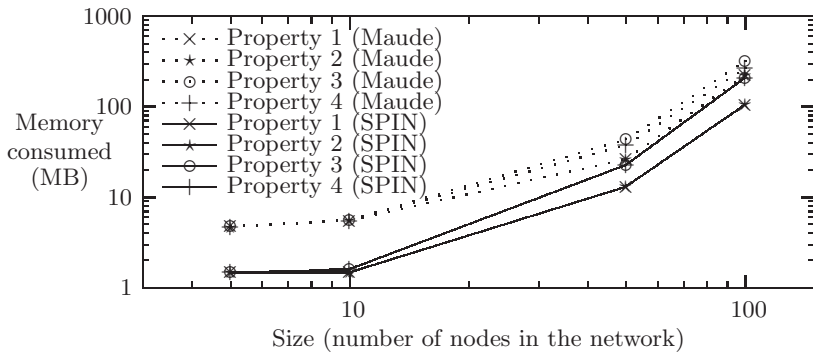


**Fig. 1.** Execution time for the leader-election problem - Maude vs. SPIN

alternating automaton into a generalized Büchi automaton (with multiple fairness conditions on arcs) and (3) convert the generalized Büchi automaton into a regular Büchi automaton. Optimizations and simplifications are performed after each step and we add some strongly connected component optimizations adapted from those in [10]. Throughout the computation, the pure propositional subformulae labelling the arcs of the various automata are stored as BDDs to allow computation of conjunctions, elimination of contradictions, and combination of parallel arcs by disjunction.

**Searching the Synchronous Product.** We use the double depth first method of [8] to lazily generate and search the synchronous product. For each system state generated we keep five bit vectors to record: (1) which propositions have been tested in the state; (2) which propositions were true in the state; (3) which product pairs (with automaton states) have been seen by the first depth first search; (4) which product pairs are currently on the first depth first search stack; and (5) which product pairs have been seen by the second (nested) depth first search. The full term graph representation of each system state is maintained (in order to test propositions) in a separate hash table which also keeps track of rewrites between system states.

**Performance Evaluation.** We compared the performance of the Maude LTL model checker vis-a-vis the SPIN LTL model checker as follows. Given a system specified in PROMELA, we specify it in Maude, and then compare the running times and the memory consumptions of the two model checkers on the respective specifications. The PROMELA specifications used — a solution to the mutual exclusion problem, a solution to the leader election problem for a unidirectional



**Fig. 2.** Total memory consumption for the leader-election problem - Maude vs. SPIN

ring network, and a translation of the  $\pi$ -calculus description of a mobile handoff scenario — are all available on the SPIN web-page.

In all the above situations, only properties satisfied by the corresponding systems were model checked; no generation of counterexamples was attempted. Except in one instance, the default settings for SPIN were used everywhere. The analyses were carried out on a 1.13 GHz Pentium III machine with 384 MB RAM running Red Hat Linux. In most of the cases, both model checkers finished fairly quickly whenever memory was available; lack of memory proved to be the main bottleneck for scalability. The benchmarks showed a comparable performance of SPIN and the Maude LTL model checker, in terms of both speed and memory consumption. The results for the leader election problem are given in Figures 1 and 2. A fuller description of the algorithms and the properties (including properties 1–4 mentioned above), as well as of their respective comparisons, is available in [4].

## References

- [1] M. Clavel and et al. Maude: specification and programming in rewriting logic. *Theoretical Computer Science*, 285:187–243, August 2002. 230, 231
- [2] Manuel Clavel and et al. Building equational proving tools by reflection in rewriting logic. In *CAFE: An Industrial-Strength Algebraic Formal Method*. Elsevier, 2000. 230
- [3] Steven Eker and et al. Pathway logic: Executable models of biological networks. In *Proc. WRLA'02*, volume 71 of *ENTCS*. Elsevier, 2002. 230
- [4] Steven Eker, José Meseguer, and Ambarish Sridharanarayanan. The Maude LTL model checker. In *Proc. WRLA '02*, volume 71 of *ENTCS*. Elsevier, 2002. 230, 233
- [5] Kousha Etessami and Gerard J. Holzmann. Optimizing Büchi automata. In *CONCUR 2000*, number 1877 in *LNCS*, pages 153–167. Springer-Verlag, 2000. 231

- [6] Paul Gastin and Denis Oddoux. Fast LTL to Büchi automata translation. In *CAV '01*, number 2102 in LNCS, pages 53–65. Springer-Verlag, 2001. 231
- [7] Rob Gerth and et al. Simple on-the-fly automatic verification of linear temporal logic. In *Protocol Specification Testing and Verification*, pages 3–18. Chapman & Hall, 1995. 231
- [8] G. J. Holzmann, D. Peled, and M. Yannakakis. On nested depth first search. *Design: An International Journal*, 13(3):289–307, nov 1998. 232
- [9] José Meseguer. Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science*, 96(1):73–155, 1992. 230
- [10] F. Somenzi and R. Bloem. Efficient Büchi automata from LTL formulae. In *CAV '00*, number 1633 in LNCS, pages 247–263. Springer-Verlag, 2000. 231, 232