MoChiBA: Probabilistic LTL Model Checking Using Limit-Deterministic Büchi Automata

Salomon Sickert[(✉)] and Jan Křetínský

Technische Universität München, Munich, Germany
{sickert,jan.kretinsky}@in.tum.de

Abstract. The limiting factor for quantitative analysis of Markov decision processes (MDP) against specifications given in linear temporal logic (LTL) is the size of the generated product. As recently shown, a special subclass of limit-deterministic Büchi automata (LDBA) can replace deterministic Rabin automata in quantitative probabilistic model checking algorithms. We present an extension of PRISM for LTL model checking of MDP using LDBA. While existing algorithms can be used only with minimal changes, the new approach takes advantage of the special structure and the smaller size of the obtained LDBA to speed up the model checking. We demonstrate the speed up experimentally by a comparison with other approaches.

1 Introduction

Linear temporal logic (LTL) [30] is a prominent specification language and has been proven useful in industrial practice. The key to efficient LTL model checking is the automata-theoretic approach [38]: first, a given LTL formula is translated into an automaton; second, a product of the system and the automaton is constructed and analysed. Since real systems are huge, it is crucial to construct *small automata* in order to avoid a large size increase of the product.

LTL is typically translated into non-deterministic Büchi automata (NBA) [2,8,10,11,14,15,17,18,36]. However, for probabilistic models such as Markov decision processes (MDP) non-deterministic automata are not applicable [3] and the standard solution is to *determinise* them using Safra's construction [22,29, 32,33,37]. This approach is implemented in the most widespread probabilistic model checker PRISM [27]. However, the determinisation step is costly and often increases the size of automata dramatically. Therefore, *direct translations of LTL to deterministic automata* have been proposed [1,13,25,26], implemented [5,16,24], and shown to be more efficient for probabilistic model checking [6]. Nevertheless, despite more sophisticated acceptance conditions, such as generalized Rabin [26], the imposed determinism inevitably increases the size of the automata.

This naturally raises the question whether fully deterministic automata are necessary or whether restricted forms of determinism are sufficient. For instance, in the setting of games where NBA are not applicable either, a weaker notion of determinism called *good-for-games automata* is sufficient [20]. It has been proven sufficient also for probabilistic model checking, but practically "did not improve on the standard approach" [23]. Further, *unambiguous automata* can be

used [9] for model checking Markov chains, but not for MDPs. Moreover, *limit-deterministic* Büchi automata (LDBA) [7,38] can be used for probabilistic model checking MDPs in the qualitative case (deciding whether a property holds with probability 1). This idea has been further explored also in the quantitative setting (computing the probability of satisfaction) and an algorithm constructing products with several (limit-)deterministic automata proposed [19]. Although LDBA cannot in general be used for probabilistic model checking, a recent translation [35] of LTL produces LDBA, which can be used in the standard algorithm based on the construction of a single product. It also discusses the subclass of LDBA that can be used for this task. Note that there is also an exponentially better translation [21] (based on [25]) of a fragment of LTL called $\text{LTL}_{\backslash\mathbf{GU}}$ into LDBA and that there is also an efficient complementation procedure for LDBA [4].

In this paper, we provide the first implementation of the probabilistic model checking procedure proposed in [35] based on LDBA. Apart from smaller sizes of LDBA, another advantage of the Büchi acceptance condition is a faster analysis of maximal end components (MECs), compared to the standard repetitive recomputation for each Rabin pair. We also present several crucial optimizations, which make our implementation outperform other approaches on many formulas. We illustrate this on experimental results. The tool as well as the explanation of its name can be found on https://www7.in.tum.de/~sickert/projects/mochiba/.

2   Overview of the Algorithm

In order to present our implementation and optimizations, we have to sketch how an MDP $\mathcal{M}$ is checked against an LTL formula φ by the algorithm of [35]. First, φ is translated into an LDBA $\mathcal{A}(\varphi)$. Second, $\mathcal{M}$ is checked against $\mathcal{A}(\varphi)$ using a straightforward extension of the standard algorithm.

LDBA Construction. An LDBA is a (possibly generalised) Büchi automaton partitioned into an *initial* and an *accepting* part, where the initial part contains no accepting transitions and the accepting part is deterministic. Moreover, the construction of [35] produces LDBA with the initial part deterministic except for ε-transitions into the accepting part.

We illustrate the translation on $\varphi = a \wedge \mathbf{X}(\mathbf{FG}a \vee \mathbf{FG}b)$. Each state in the initial part is labeled with a formula. The words accepted from a state are exactly those satisfying the formula. Observe that $\mathbf{FG}a \vee \mathbf{FG}b$ holds iff eventually we reach a point where $\mathbf{G}a$ holds or $\mathbf{G}b$ holds. We non-deterministically guess this point and take the ε-transition to the accepting part, where we check the guess.


Fig. 1. LDBA $\mathcal{A}(\varphi)$ with the initial part on the left and the accepting on the right.

For this formula `spot` (2.0) produces a deterministic Rabin automaton with 4 states, too, but adding two more disjuncts $\mathbf{FG}c$ and $\mathbf{FG}d$ increases the size to 26 states. In contrast, our LDBA requires only two extra states.

Product Construction and Analysis. We proceed according to the standard algorithm:

1. Construct the product $\mathcal{P} = \mathcal{M} \times \mathcal{A}(\varphi)$.
2. Compute maximal end-components (MECs) of $\mathcal{P}$.
3. Compute the maximum probability to reach winning MECs. A MEC is winning if it satisfies the acceptance condition of $\mathcal{A}(\varphi)$: here, if it contains an accepting transition for each Büchi condition.

Fig. 2. An MDP $\mathcal{M}$.

The standard product of an MDP and a deterministic automaton defines the transitions (in the usual notation) by $P(\langle s, q \rangle, \alpha, \langle s', q' \rangle) = P(s, \alpha, s')$ if $q' = \delta(q, Label(s'))$ and otherwise equals 0. We extend the procedure to handle also non-deterministic ε-transitions by additional actions: let $q_1, \ldots, q_n$ be the successors of q under ε, then for each $i = 1, \ldots, n$ we add a new action called ε_{q_i} and define $P(\langle s, q \rangle, \varepsilon_{q_i}, \langle s, q_i \rangle) = 1$ (note that s does not move here).

Fig. 3. The product $\mathcal{M} \times \mathcal{A}(\varphi)$.

Figure 3 illustrates the construction by a product of the system of Fig. 2 and the automaton of Fig. 1. A crucial optimization used here is that it is sufficient to take ε-transitions only from states in MECs of $\mathcal{M} \times \mathcal{N}(\varphi)$ (which are exactly MECs of the product of $\mathcal{M}$ and the initial part of $\mathcal{A}(\varphi)$). Hence no ε-transitions have to be produced in the initial state here.

3 Implementation and Optimizations

`MoChiBA` [34] replaces the LTL model-checker and the MEC computation in the explicit-state model-checker of PRISM, while other infrastructure (parsing, model construction, probability computation) are inherited from PRISM. The tool cannot be configured — all optimizations are enabled — and does not need to be installed. It reads a model (given as an MDP, `.nm`) and a property specification (`.pctl`) and prints the results to `stdout`:

```
./mochiba.sh model.nm properties.pctl
```

Apart from taking ε-transitions only from states in MECs of $\mathcal{M} \times \mathcal{N}(\varphi)$ as mentioned above, we implement the following optimizations:

Transition-based acceptance leads to smaller automata, compared to state-based acceptance. Consequently, it is used by many translators, for

instance [1,11,24]. However, PRISM translates all automata to state-based, thus increasing the size of the product. Our procedure avoids this and constructs and analyses directly the transition-based product.

Generalised Büchi acceptance condition allows for more efficient analysis than (generalised) Rabin, Streett, or parity conditions. Indeed, for the latter conditions expensive re-computations of MECs are necessary to handle different sets to be visited finitely often. In contrast, we compute MECs only once and check whether each set to be visited infinitely often intersects the MEC.

A single trap state is present in the product. Should the product enter any state from which the automaton component can never accept, the exploration of this part stops and redirects the transition to the single trap state.

Primitive data structures such as arrays are used instead of the more flexible Java collections, since they are more memory efficient, as boxing into objects is not necessary.

Sparse bit sets have proven more memory efficient for our approach than plain bit sets with a mapping table.

MEC decomposition is performed locally on disconnected accepting parts (corresponding to different ε-transitions). Together with the use of sparse bit sets, MECs are computed faster and using less memory.

4   Experimental Evaluation

We evaluate our novel approach in the setting of [6,19]: we consider the Pneuli-Zuck randomised mutual exclusion protocol [31] of the PRISM benchmark suite [28] and also the same previously considered formulas (see lines 1–10 of Table 1). Additionally, lines 11–14 consider the deeply nested formulas of [35]. Finally, complementary to the **GF**-, **FG**- and fairness-like properties, lines 15–16 include simple reachability properties, which lie in the focus of the traditional methods.

The experiments were performed on a 2.5 GHz Intel Core i7 (I7-4870HQ) and granted 12 GB RAM and 1 h computing time for model checking each property (given the model already in the memory). We denote time-outs and mem-outs by "-". We compare the following tools

- `MoChiBA` (1.0) [34] is our implementation based on the LDBA translation of [35] and the explicit model checker of PRISM.
- `PRISM` (4.3) [27] with the symbolic engine, which is the fastest here, and with the following translators:
  - Built-in LTL to deterministic Rabin automaton translation, re-implementing ltl2dstar [22].
  - `Rabinizer` (3.1) [24] using the Safra-less direct translation into generalised Rabin automata, which are now supported by PRISM.
- `IscasMC` (unofficial, unversioned) implements the lazy approach of [19], using SPOT 1.2.6 [12] to translate LTL to non-deterministic Büchi automaton. We used the two fastest configurations as listed in [19]:
  - Multi-breakpoint (`BP`) construction with the explicit engine.
  - Rabin (`R`) construction with the explicit engine.

Table 1. Runtime comparison on model checking these properties on the Pneuli-Zuck randomised mutual exclusion protocol [31].

Property	n	MoChiBA	PRISM	Rabinizer	IscasMC-BP	IscasMC-R
				Time (rounded, in seconds)		
(1) $\mathbb{P}_{max=?}[\mathbf{GF}p1=10\wedge\mathbf{GF}p2=10 \atop \wedge\mathbf{GF}p3=10]$	4	< 1	16	< 1	< 1	< 1
	5	2	230	< 1	12	11
(2) $\mathbb{P}_{max=?}[{\mathbf{GF}p1=10\wedge\mathbf{GF}p2=10 \atop \wedge\mathbf{GF}p3=10\wedge\mathbf{GF}p4=10}]$	4	< 1	26	< 1	1	< 1
	5	2	345	< 1	12	12
(3) $\mathbb{P}_{min=?}[{\mathbf{GF}p1=10\wedge\mathbf{GF}p2=10 \atop \wedge\mathbf{GF}p3=10\wedge\mathbf{GF}p4=10}]$	4	1	3552	33	1	22
	5	11	-	572	18	641
(4) $\mathbb{P}_{max=?}[{(\mathbf{GF}p1=0\vee\mathbf{FG}p2\neq0) \atop \wedge(\mathbf{GF}p2=0\vee\mathbf{FG}p3\neq0)}]$	4	1	684	18	2	4
	5	15	-	293	19	50
(5) $\mathbb{P}_{max=?}[{(\mathbf{GF}p1=0\vee\mathbf{FG}p1\neq0) \atop \wedge(\mathbf{GF}p2=0\vee\mathbf{FG}p2\neq0)}]$	4	< 1	< 1	23	1	4
	5	1	< 1	403	17	59
(6) $\mathbb{P}_{max=?}[{(\mathbf{GF}p1=0\vee\mathbf{FG}p2\neq0) \atop \wedge(\mathbf{GF}p2=0\vee\mathbf{FG}p3\neq0) \atop \wedge(\mathbf{GF}p3=0\vee\mathbf{FG}p1\neq0)}]$	4	< 1	78	9	3	10
	5	10	1293	137	29	143
(7) $\mathbb{P}_{max=?}[{(\mathbf{GF}p1=0\vee\mathbf{FG}p1\neq0) \atop \wedge(\mathbf{GF}p2=0\vee\mathbf{FG}p2\neq0) \atop \wedge(\mathbf{GF}p3=0\vee\mathbf{FG}p3\neq0)}]$	4	< 1	< 1	61	2	18
	5	1	< 1	1077	27	277
(8) $\mathbb{P}_{min=?}[{(\mathbf{GF}p1\neq10\vee\mathbf{GF}p1=0\vee\mathbf{FG}p1=1) \atop \wedge\mathbf{GF}p1\neq0\wedge\mathbf{GF}p1=1}]$	4	< 1	8	8	1	1
	5	1	145	190	16	21
(9) $\mathbb{P}_{max=?}[{(\mathbf{G}p1\neq10\vee\mathbf{G}p2\neq10\vee\mathbf{G}p3\neq10) \atop \wedge(\mathbf{FG}p1\neq1\vee\mathbf{GF}p2=1\vee\mathbf{GF}p3=1) \atop \wedge(\mathbf{FG}p2\neq1)\vee\mathbf{GF}p1=1\vee\mathbf{GF}p3=1)}]$	4	5	-	1195	8	871
	5	99	-	-	125	-
(10) $\mathbb{P}_{min=?}[{\mathbf{FG}p1\neq0\vee\mathbf{FG}p2\neq0 \atop \vee\mathbf{GF}p3=0\vee(\mathbf{FG}p1\neq10 \atop \wedge\mathbf{GF}p2=10\wedge\mathbf{GF}p3=10)}]$	4	1	728	33	79	6
	5	24	-	486	-	77
(11) $\mathbb{P}_{min=?}[f_{0,0}] = \mathbb{P}_{min=?}[{(\mathbf{GF}p1=10)\mathbf{U} \atop (p2=10)}]$	4	< 1	17	40	2	2
	5	11	257	715	23	54
(12) $\mathbb{P}_{max=?}[f_{0,4}] = \mathbb{P}_{max=?}[{(\mathbf{GF}p1=10)\mathbf{U} \atop (\mathbf{XXXX}p2=10)}]$	4	< 1	3	< 1	1	15
	5	5	20	2	20	2381
(13) $\mathbb{P}_{min=?}[f_{1,0}] = \mathbb{P}_{min=?}[{(\mathbf{GF}p1=10)\mathbf{U} \atop (\mathbf{G}((\mathbf{GF}p2=10)\mathbf{U} \atop (p3=10)))}]$	4	< 1	909	22	314	4
	5	13	-	436	-	59
(14) $\mathbb{P}_{max=?}[f_{1,4}] = \mathbb{P}_{max=?}[{(\mathbf{GF}p1=10)\mathbf{U} \atop (\mathbf{G}((\mathbf{GF}p2=10)\mathbf{U} \atop (\mathbf{XXXX}p3=10)))}]$	4	< 1	-	18	2	2
	5	12	-	285	24	25
(15) $\mathbb{P}_{max=?}[p1 = 0 \ \ \mathbf{U} \ \ p2 = 10]$	4	< 1	< 1	< 1	< 1	< 1
	5	< 1	< 1	< 1	7	7
(16) $\mathbb{P}_{max=?}[\mathbf{XXXXXX}p1 = 0]$	4	< 1	< 1	< 1	1	< 1
	5	3	< 1	< 1	19	16

5   Conclusion

We have implemented a novel approach for probabilistic LTL model checking using a subclass of non-deterministic Büchi automata. Since the experimental results for the explicit state-space implementation are encouraging, we plan to extend the approach to a symbolic one. Further, a parellelisation of the product construction and MECs analysis, as well as dedicated constructions for the *Release*-operator or various LTL fragments could lead to further speed ups.

References

1. Babiak, T., Blahoudek, F., Křetínský, M., Strejček, J.: Effective translation of LTL to deterministic Rabin automata: beyond the (F, G)-fragment. In: ATVA, pp. 24–39 (2013)
2. Babiak, T., Křetínský, M., Řehák, V., Strejček, J.: LTL to Büchi automata translation: fast and more deterministic. In: Flanagan, C., König, B. (eds.) TACAS 2012. LNCS, vol. 7214, pp. 95–109. Springer, Heidelberg (2012). doi:10.1007/978-3-642-28756-5_8
3. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press, Cambridge (2008)
4. Blahoudek, F., Heizmann, M., Schewe, S., Strejček, J., Tsai, M.-H.: Complementing semi-deterministic Büchi automata. In: Chechik, M., Raskin, J.-F. (eds.) TACAS 2016. LNCS, vol. 9636, pp. 770–787. Springer, Heidelberg (2016). doi:10.1007/978-3-662-49674-9_49
5. Blahoudek, F., Křetínský, M., Strejček, J.: Comparison of LTL to deterministic Rabin automata translators. In: McMillan, K., Middeldorp, A., Voronkov, A. (eds.) LPAR 2013. LNCS, vol. 8312, pp. 164–172. Springer, Heidelberg (2013). doi:10.1007/978-3-642-45221-5_12
6. Chatterjee, K., Gaiser, A., Křetínský, J.: Automata with generalized Rabin pairs for probabilistic model checking and LTL synthesis. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 559–575. Springer, Heidelberg (2013). doi:10.1007/978-3-642-39799-8_37
7. Courcoubetis, C., Yannakakis, M.: The complexity of probabilistic verification. J. ACM **42**(4), 857–907 (1995)
8. Couvreur, J.-M.: On-the-fly verification of linear temporal logic. In: Wing, J.M., Woodcock, J., Davies, J. (eds.) FM 1999. LNCS, vol. 1708, pp. 253–271. Springer, Heidelberg (1999). doi:10.1007/3-540-48119-2_16
9. Couvreur, J.-M., Saheb, N., Sutre, G.: An optimal automata approach to LTL model checking of probabilistic systems. In: Vardi, M.Y., Voronkov, A. (eds.) LPAR 2003. LNCS (LNAI), vol. 2850, pp. 361–375. Springer, Heidelberg (2003). doi:10.1007/978-3-540-39813-4_26
10. Daniele, M., Giunchiglia, F., Vardi, M.Y.: Improved automata generation for linear temporal logic. In: Halbwachs, N., Peled, D. (eds.) CAV 1999. LNCS, vol. 1633, pp. 249–260. Springer, Heidelberg (1999). doi:10.1007/3-540-48683-6_23
11. Duret-Lutz, A.: Manipulating LTL formulas using Spot 1.0. In: Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 442–445. Springer, Heidelberg (2013). doi:10.1007/978-3-319-02444-8_31
12. Duret-Lutz, A.: LTL translation improvements in Spot 1.0. IJCCBS **5**(1–2), 31–54 (2014)
13. Esparza, J., Křetínský, J.: From LTL to deterministic automata: a safraless compositional approach. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 192–208. Springer, Heidelberg (2014). doi:10.1007/978-3-319-08867-9_13

14. Etessami, K., Holzmann, G.J.: Optimizing Büchi automata. In: Palamidessi, C. (ed.) CONCUR 2000. LNCS, vol. 1877, pp. 153–168. Springer, Heidelberg (2000). doi:10.1007/3-540-44618-4_13

15. Fritz, C.: Constructing Büchi automata from linear temporal logic using simulation relations for alternating Büchi automata. In: Ibarra, O.H., Dang, Z. (eds.) CIAA 2003. LNCS, vol. 2759, pp. 35–48. Springer, Heidelberg (2003). doi:10.1007/3-540-45089-0_5

16. Gaiser, A., Křetínský, J., Esparza, J.: Rabinizer: small deterministic automata for LTL(F, G). In: ATVA, pp. 72–76 (2012)

17. Gastin, P., Oddoux, D.: Fast LTL to Büchi automatatranslation. In: CAV, pp. 53–65 (2001). http://www.lsv.ens-cachan.fr/~gastin/ltl2ba/

18. Giannakopoulou, D., Lerda, F.: From states to transitions: improving translation of LTL formulae to Büchi automata. In: Peled, D.A., Vardi, M.Y. (eds.) FORTE 2002. LNCS, vol. 2529, pp. 308–326. Springer, Heidelberg (2002). doi:10.1007/3-540-36135-9_20

19. Hahn, E.M., Li, G., Schewe, S., Turrini, A., Zhang, L.: Lazyprobabilistic model checking without determinisation. In: CONCUR. LIPIcs, vol. 42, pp. 354–367 (2015)

20. Henzinger, T.A., Piterman, N.: Solving games without determinization. In: Ésik, Z. (ed.) CSL 2006. LNCS, vol. 4207, pp. 395–410. Springer, Heidelberg (2006). doi:10.1007/11874683_26

21. Kini, D., Viswanathan, M.: Limit deterministic and probabilistic automata for LTL \ GU. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 628–642. Springer, Heidelberg (2015)

22. Klein, J.: ltl2dstar - LTL to deterministic Streett and Rabinautomata. http://www.ltl2dstar.de/

23. Klein, J., Müller, D., Baier, C., Klüppelholz, S.: Are good-for-games automata good for probabilistic model checking? In: Dediu, A.-H., Martín-Vide, C., Sierra-Rodríguez, J.-L., Truthe, B. (eds.) LATA 2014. LNCS, vol. 8370, pp. 453–465. Springer, Heidelberg (2014). doi:10.1007/978-3-319-04921-2_37

24. Komárková, Z., Křetínský, J.: Rabinizer 3: safraless translation of LTL to small deterministic automata. In: Cassez, F., Raskin, J.-F. (eds.) ATVA 2014. LNCS, vol. 8837, pp. 235–241. Springer, Heidelberg (2014). doi:10.1007/978-3-319-11936-6_17

25. Křetínský, J., Ledesma-Garza, R.: Rabinizer 2: small deterministic automata for $LTL_{\backslash \mathbf{GU}}$. In: ATVA, pp. 446–450 (2013)

26. Křetínský, J., Esparza, J.: Deterministic automata for the (F, G)-fragment of LTL. In: CAV, pp. 7–22 (2012)

27. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011). doi:10.1007/978-3-642-22110-1_47

28. Kwiatkowska, M.Z., Norman, G., Parker, D.: The PRISM benchmark suite. In: QEST, pp. 203–204 (2012)

29. Piterman, N.: From nondeterministic Büchi and Streett automata to deterministic parity automata. In: LICS, pp. 255–264 (2006)

30. Pnueli, A.: The temporal logic of programs. In: FOCS, pp. 46–57 (1977)

31. Pnueli, A., Zuck, L.D.: Verification of multiprocess probabilistic protocols. Distrib. Comput. **1**(1), 53–72 (1986). doi:10.1007/BF01843570

32. Safra, S.: On the complexity of omega-automata. In: FOCS, pp. 319–327 (1988)

33. Schewe, S.: Tighter bounds for the determinisation of Büchi automata. In: Alfaro, L. (ed.) FoSSaCS 2009. LNCS, vol. 5504, pp. 167–181. Springer, Heidelberg (2009). doi:10.1007/978-3-642-00596-1_13

34. Sickert, S.: MoChiBA. https://www7.in.tum.de/~sickert/projects/mochiba/
35. Sickert, S., Esparza, J., Jaax, S., Křetínský, J.: Limit-deterministic Büchi automata for linear temporal logic. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9780, pp. 312–332. Springer, Heidelberg (2016). doi:10.1007/978-3-319-41540-6_17
36. Somenzi, F., Bloem, R.: Efficient Büchi automata from LTL formulae. In: Emerson, E.A., Sistla, A.P. (eds.) CAV 2000. LNCS, vol. 1855, pp. 248–263. Springer, Heidelberg (2000). doi:10.1007/10722167_21
37. Tsai, M.-H., Tsay, Y.-K., Hwang, Y.-S.: GOAL for games, omega-automata, and logics. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 883–889. Springer, Heidelberg (2013). doi:10.1007/978-3-642-39799-8_62
38. Vardi, M.Y., Wolper, P.: An automata-theoretic approach to automatic program verification (preliminary report). In: LICS, pp. 332–344 (1986)